The pricing of volatility risk across asset classes

Zhi Da[†] and Ernst Schaumburg[‡,*]

Version: November 21, 2011

Abstract

In the Merton (1973) ICAPM, state variables that capture the evolution of the investor's opportunity set are necessary to explain observed asset prices. We show that augmenting the CAPM by a measure of market-wide volatility innovation yields a two-factor model that performs well in explaining the cross-section of returns on securities in several asset classes. The consistent pricing of volatility risk (with a negative risk premium) suggests that volatility risk indeed acts as a state variable rather than being just another statistical factor. In addition, we propose a novel method for extracting volatility risk factors from the cross-section and find it help to price assets, especially synthetic volatility swaps.

*We thank Tobias Adrian, Torben Andersen, Mikhail Chernov, Robert Korajczyk, Paul Gao, Robert Whitelaw and Yexiao Xu; seminar participants at University of Notre Dame, the April 2006 Risk Management Conference (Mont Tremblant), the 2007 China International Conference in Finance, the Federal Reserve Bank of Chicago, the Federal Reserve Bank of New York, and the Caesarea Centre 6th Annual Academic Conference; and especially Ravi Jagannathan for numerous suggestions and insights.

† zda@nd.edu, Mendoza College of Business, University of Notre Dame.

‡ Corresponding author: Ernst.Schaumburg@gmail.com, Federal Reserve Bank of New York, Tel: 212-720-2744, Fax: 212-720-1582

1 Introduction

The idea that volatility risk should be priced has received considerable attention in the asset pricing literature.[1] In a discrete-time version of the Merton (1973) intertemporal capital asset pricing model (ICAPM), Campbell (1993) shows that any variable that forecasts future returns or future volatility is a good candidate state variable. This result is simplified considerably in a continuous-time setting (with no jumps), where Nielsen and Vassalou (2006) demonstrate that a single state variable (the instantaneous maximum Sharpe ratio) is a sufficient statistic to describe the investor opportunity set. Empirically, Brennan, Wang, and Xia (2004) show that including a measure of innovations to the maximum Sharpe ratio improves the performance of their pricing model significantly. Without strong parametric assumptions, however, it is not practical to work directly with Sharpe ratio innovations. Moreover, state variables that reliably forecast future returns out-of-sample are hard to come by, as mixed evidence on return predictability shows.[2] This implies that, to the extent Sharpe ratios are predictable, the predictability is likely driven mainly by the predictability of the denominator.[3]

Most of the existing studies of volatility risk have almost exclusively investigated the pricing of risk in portfolios of a single asset class (e.g., stocks). This approach, however, fails to fully leverage the strong implications of the ICAPM framework: If market volatility is truly a state variable in the ICAPM sense rather than just another statistical factor, it should be priced *consistently* across asset classes. Moreover, any valid proxy for the underlying state variable should produce similar results in asset pricing tests. The discipline imposed by cross-asset class pricing improves statistical power against certain alternatives that might otherwise yield a spuriously high volatility risk premium.

It is well known that linear beta pricing may incorrectly price non-linear payoffs, as noted by Wang and Zhang (2006) among others. For instance, in a model where some stocks have higher betas in down markets and lower betas in up markets, a spurious

[1] We review this literature in detail in the next section.

[2] see Boudoukh, Richardson, and Whitelaw (2008), Campbell and Thompson (2008), Cochrane (2008), Lettau and Nieuwerburgh (2008), and Welch and Goyal (2008)

[3] A point Breen, Glosten, and Jagannathan (1989) make in a market timing context.

finding of a negative volatility risk premium is to be expected, as we show in Appendix A. Yet, in such a non-linear beta world, other asset classes, say, static delta-neutral option portfolios, should not show any significant volatility premium in pricing. On the other hand, the finding of a very high negative volatility risk premium in the market for certain stock options may in part reflect a liquidity premium earned by option market makers rather than a risk premium, but this should not be reflected in, say, bond prices.[4]

To our best knowledge, our paper is the first to examine the pricing of volatility risk across multiple asset classes in a coherent framework using a single volatility proxy, which allows us to compare not only the sign but also the magnitude of volatility risk premium across various asset classes.[5] We find that a simple two-factor model, consisting of the market excess return factor and a volatility innovation factor, can price different assets such as portfolios of stocks, stock options and corporate bonds. Consistent with the prior literature, we find a significant negative volatility risk premium especially in the stock and the option market. The volatility risk premium is also negative but not significant for corporate bonds, due to the fact that the corporate bond returns are extremely volatile during the short sample period we examine. Adding to the literature, we establish that the magnitude of the volatility risk premium is similar across asset assets, supporting the notion that volatility is a state-variable in the ICAPM sense and not just another statistical pricing factor derived from an underlying factor structure in asset returns.

We measure the volatility innovations using a number of commonly used broad-based stock indices. These measures all perform similarly in our asset pricing tests, although it is crucial that a measure of *unforecastable* innovations be used rather than levels or first-differences, and it appears that value-weighted indices do better than equal-weighted ones.

One area of some concern is the arbitrariness of the stock index choice, which results in a volatility proxy that may be affected by time-varying portfolio weights and correlations. A novel alternative approach, which we pursue in this paper, is to construct a "non-parametric" volatility proxy by analyzing the cross-section of realized monthly volatility

[4]Deng (2008) for instance finds that option market inefficiencies are first order important.
[5]More recently Bollerslev, Marrone, Xu, and Zhou (2011) have studied volatility risk pricing across countries.

innovations of US equities over a long sample period. We find that the (unbalanced) panel of univariate volatility innovations is well described by a simple factor structure. Moreover, the principal factor is highly correlated with the various market index volatility innovations, although substantially less noisy and yields the best pricing performance of the volatility measures considered.

An additional advantage of the non-parametric principal component analysis is that it allows us to entertain the possibility of more than one priced component of stock market volatility. When we price synthetic volatility swaps, it appears that at least one additional volatility factor may be helpful in explaining the cross-section of swap returns.

We find that our simple two-factor model prices different assets as well as the Fama and French (1993) three-factor model. The fact that the Fama-French three-factor model has any pricing ability at all beyond portfolios of stocks suggests that there is more to the HML and SMB factors than merely sorting stocks into portfolios on the basis of a characteristic that ex-post is correlated with expected stock returns, as McKinlay (1995), and Ferson, Sarkissian, and Simin (1999) suggest. This is confirmed by the close relation between the loadings on the HML and SMB factors and the loading on our volatility factor. The finding is robust across asset classes, and, in fact, augmenting our two-factor model by the HML and SMB factors yields no further improvement in pricing ability.

Liewa and Vassalou (2000), Vassalou (2003), and Petkova (2006) suggest that the Fama-French factors are related to macroeconomic variables that appear to span the SMB and HML factors (or their projection onto the payoff space of size and book-to-market sorted portfolios). Their work is very different from ours, as it is based on arbitrage pricing theory (APT) rather than the ICAPM. The fact that the HML and SMB factors appear related to macroeconomic variables is of course not entirely surprising, as the link between stock market volatility and indicators of economic fundamentals has been well documented (see Schwert (1989) and Hamilton and Lin (1998) among others).

The remainder of the paper is structured as follows. Section 2 briefly reviews the pricing of volatility risk and the most closely related empirical literature. Section 3 discusses the measurement of aggregate volatility innovations while section 4 investigates the pricing implications for portfolios of stocks, stock options, and corporate bonds. Section

5 concludes.

2 The Pricing of Volatility Risk

Aggregate market volatility is a natural state variable that describes the investor's investment opportunity set, and, in the Merton (1973) ICAPM framework, covariance with volatility innovations will therefore be priced. In a discrete time setting with Epstein-Zin utility and time-varying volatility, Campbell (1993) derives an equation for the (log-linearized) stochastic discount factor as a function of the current market return and innovations in expectations about future market return and volatility:

$$\begin{aligned} m_{t+1} &= \gamma E_t \Delta r_{m,t+1} + (1-\gamma)(E_{t+1}-E_t)\sum_{j=1}^{\infty}\rho^j r_{m,t+1+j} \\ &\quad -\frac{\theta}{2\sigma}(E_{t+1}-E_t)\sum_{j=1}^{\infty}\rho^j \mathbb{V}\mathrm{ar}_{t+j}\left[\Delta c_{t+j+1}-\sigma r_{m,t+j+1}\right] \end{aligned} \tag{1}$$

where r_m and Δc represent the (log) market return and aggregate consumption growth, respectively. Since aggregate consumption growth is less volatile than the market return empirically, we have $\mathbb{V}\mathrm{ar}_t\left[\Delta c_{t+1}-\sigma r_{m,t+1}\right] \approx \sigma^2 \mathbb{V}\mathrm{ar}_t\left[r_{m,t+1}\right]$, and any factor that can predict future market return (as in the second term) or predict future market volatility (as in the third term) is therefore a good candidate state variable. Because volatility is persistent, one may therefore argue that the volatility innovations of the stock market index represents a reasonable choice of pricing factor.

We follow the empirical literature and abstract from the literal setting of Campbell (1993). Instead we focus on the specific components of the stochastic discount factor (SDF) made up by the (gross) market excess return (R_m) and volatility innovations (ΔV), by positing the SDF specification:

$$M_{t+1} = \frac{1}{R_f}(\lambda_0 - \lambda_1 R_{m,t+1} - \lambda_2 \Delta V_{t+1}) \tag{2}$$

If we denote the investor's value function by $\mathcal{J}$ and level of wealth by W, we can rewrite the coefficients in (2) as resulting from the investor's first-order conditions:

$$\lambda_1 \;=\; -\mathcal{J}_{WW}W/\mathcal{J}_W \;\;(\equiv \gamma \text{ in } (1)\,) \;\;, \qquad \lambda_2 = -\mathcal{J}_{WV}/\mathcal{J}_W$$

The SDF specification (2) leads to the pricing equation for an arbitrary asset's (gross) return R_{t+1}:

$$\begin{aligned} E_t[M_{t+1}R_{t+1}] = 1 & \qquad\qquad\qquad\qquad \Rightarrow \\ E_t[R_{t+1}] - R_f = & \;\; \lambda_1 \mathbb{C}\mathrm{ov}_t(R_{t+1}, R_{m,t+1}) + \lambda_2 \mathbb{C}\mathrm{ov}_t(R_{t+1}, \Delta V_{t+1}) \end{aligned} \tag{3}$$

The model (3) represents the most parsimonious pricing framework in which to study the relationship between volatility risk and expected returns. The sign of the variance risk premium λ_2 is likely to be negative for at least two reasons. First, all else equal, unexpectedly high volatility worsens the investor's risk-return trade-off and hence corresponds to a bad state of the world. Second, high volatility often coincides with periods of low market returns so that assets that are highly sensitive to market volatility serve as a good hedge ($\mathcal{J}_{WV} > 0$), and therefore should earn a lower expected return.

Examination of the risk-return relation implied in (3) is of fundamental importance to the asset pricing literature. The existing research can be divided into two broad groups, according to focus. Papers in the first group focus on the time-series risk-return relation.

To focus on the market return itself, (3) becomes:

$$E_{t+1}[R_{m,t+1}] - R_f = \lambda_1 \mathbb{V}\mathrm{ar}_t(R_{m,t+1}) + \lambda_2 \mathbb{C}\mathrm{ov}_t(R_{m,t+1}, \Delta V_{t+1}) \tag{4}$$

Many authors either fail to identify a statistically significant intertemporal relation between risk and return of the market portfolio or find a negative relation. Examples include L.R. Glosten (1993), Whitelaw (1994), and Harvey (2001).

More recently, better estimating expected return and expected volatility or explicitly accounting for hedging demands [the second term in (3)], several authors have found a positive risk-return relation in the time series. French, Schwert, and Stambaugh (1987) and Ghysels, Santa-Clara, and Valkanov (2005), for instance, estimate $\mathbb{V}\mathrm{ar}_t(R_{m,t+1})$ using squared daily returns, while Guo and Whitelaw (2006) explicitly model both the risk component and the hedging component. Bollerslev, Tauchen, and Zhou (2009) consider the difference between implied and realized variance, and show that it predicts future market return. Moving away from a single market return, Bali (2008) establishes a positive time series risk-return relation for a large cross-section of stock portfolios using a GARCH estimation procedure. In a recent paper, Smith and Whitelaw (2009) suggest that a model with a (counter-cyclical) time varying risk premium may resolve some of the difficulty in identifying a positive significant λ_1.

The second group of research examines the cross-sectional risk-return relation, particularly the pricing of volatility risk implied by the hedge component [the second term in (3)]. If the volatility risk premium (λ_2) is negative, then the asset with more sensitivity to volatility risk should earn a lower average return in the cross-section. This approach has been used to examine the pricing of volatility risk in the stock market. Ang, Hodrick, Xing, and Zhang (2006), for instance, measure volatility risk using changes in the *VIX* index from the Chicago Board Options Exchange. They document over their sampling period of 1986 - 2000 a negative volatility risk premium, and confirm that stocks that are more sensitive to volatility risk do earn lower returns. Adrian and Rosenberg (2008) decompose the market volatility into separate long-run and short-run components and show that the return covariance with each component is priced, and risk premia on both components are negative.

The pricing of volatility risk has also been examined in stock options markets. The identification strategy involves constructing a set of market-neutral option portfolios that are sensitive only (or at least to first-order) to volatility risk, making it a clean test asset for testing the pricing of volatility risk. Two examples are: (1) delta-hedged index and individual stock options (see Bakshi and Kapadia (2003a), (2003b), and Duarte and Jones (2007)); and (2) synthetic variance swaps (see Bondarenko (2004) and Carr and Wu (2009)).

To see clearly how the identification works in the options setting, let R_o be the (gross) return on a delta-neutral portfolio of options written on stock i, and let V_i be the volatility of the underlying. Using equation (3) and Stein's lemma, we have:

$$\begin{aligned}
\lambda_1 \mathbb{C}\mathrm{ov}_t(R_{o,t+1}, R_m) &= \lambda_1 E_t\left[\frac{\partial R_o}{\partial V_i}\right]\mathbb{C}\mathrm{ov}_t(V_{i,t+1}, R_{m,t+1}),\\
\lambda_2 \mathbb{C}\mathrm{ov}_t(R_{o,t+1}, \Delta V_{t+1}) &= \lambda_2 E_t\left[\frac{\partial R_o}{\partial V_i}\right]\mathbb{C}\mathrm{ov}_t(V_{i,t+1}, \Delta V_{t+1})\\
&= \lambda_2 \beta_i^V E_t\left[\frac{\partial R_o}{\partial V_i}\right]\mathbb{V}\mathrm{ar}_t(\Delta V_{t+1}),\\
\text{where } \beta_i^V &\equiv \mathbb{C}\mathrm{ov}_t(V_{i,t+1}, \Delta V_{t+1})/\mathbb{V}\mathrm{ar}_t(\Delta V_{t+1}),
\end{aligned}$$

and all other partial derivatives are zero by portfolio construction. In this case (3) can be rewritten as:

$$\begin{aligned}
E_t[R_{o,t+1}] - R_f \quad = \quad & \lambda_1 E_t\left[\frac{\partial R_o}{\partial V_i}\right]\mathbb{C}\mathrm{ov}_t(V_{i,t+1}, R_{m,t+1}) + \\
& \lambda_2 \beta_i^V E_t\left[\frac{\partial R_o}{\partial V_i}\right]\mathbb{V}\mathrm{ar}_t(\Delta V_{t+1})
\end{aligned} \tag{5}$$

Most of the recent empirical findings on variance risk in option markets can be understood using the pricing equation (6). First, using index options data, the market price of aggregate variance risk is shown to be negative as in Bakshi and Kapadia (2003a), Bondarenko (2004), and Carr and Wu (2009). Second, for the index option, $V_{i,t+1} = V_{t+1}$ and $\beta_i^V = 1$. The first term, $\lambda_1 E\left[\frac{\partial R_{o,t+1}}{\partial V_i}\right]\mathbb{C}\mathrm{ov}_t(V_m, R_m)$, is usually estimated to be negative (but close to zero) since down markets tend to be associated with above-average levels of volatility. Therefore, if the option portfolios have very large negative returns on average, as found elsewhere, it must be the case that λ_2 is negative. Third, the excess return on the options portfolio could be positive for an individual stock, as found in both Bakshi and Kapadia (2003b) and Carr and Wu (2009). This is consistent with (6), provided that β_i^V is negative, which means that the underlying stock tends to be less volatile than average when aggregate volatility is high. Finally, in a cross-sectional regression, the expected excess return on options portfolios should decline with β_i^V, which is documented by Carr and Wu (2009) using synthetic variance swaps for a sample of 40 stocks and stock indices (although their conclusion arguably may depend on a few index option outliers).

Our research falls squarely within the second group in that we focus on estimating

λ_2, but it differs from the literature in two important aspects. First, rather than focus on the testing the pricing of volatility risk in a single market, we take the implications of the ICAPM seriously and examine pricing performance across multiple asset classes within a coherent pricing framework. While recent papers by Moise (2010) and DeLisle, Price, and Sirmans (2011) have explored the pricing of volatility risk in alternative asset classes such as treasury bonds and REITs, to our best knowledge, our paper is the first to examine the pricing of volatility risk across stocks, bonds, and options in a coherent framework using a single volatility proxy, which allows us to compare not only the sign but also the magnitude of volatility risk premium across various asset classes. Second, we examine the impact of the choice of volatility proxy and propose a non-parametric measure of volatility innovations based on a principal component analysis, which allows us to investigate whether more than one component of volatility risk is priced in a linear beta pricing setting.

3 Measuring Aggregate Volatility Innovations

To examine the volatility risk premium, one needs to measure the aggregate volatility innovations first. In this section, we consider two ways to do that. We first follow standard practice in the literature to proxy the aggregate volatility by the the volatility on equity indices. As a non-parametric alternative, we also extract the aggregate volatility innovation as the first principal component from a cross-section of individual stock volatility innovations. This alternative allows us to extract additional volatility risk factors that are potentially useful in cross-sectional asset pricing.

3.1 Stock index volatility innovations

The fundamental unobservability of the market portfolio, and hence market returns, as pointed out by Roll (1977), applies equally to the measurement of volatility risk. While stocks are a non-trivial element of the overall market portfolio, stock market volatility risk is just one component of aggregate volatility risk. Yet if covariance with aggregate volatility risk is priced, and the assets of interest are, say, equities, equity options, or low-grade corporate bonds, then stock market volatility risk is arguably a component of first-order importance in determining this covariance. It is therefore reasonable to consider stock market volatility as a proxy for aggregate volatility when we price such assets.[6] In fact, much of the literature on the pricing of volatility risk has implicitly followed this logic by choosing stock index volatility innovations as the proxy for volatility risk.

As is the case with the CAPM, theory provides no guidance for which specific stock index to choose, except that it should be broad-based and that value-weighted indices are preferred. Moreover, any noisy proxy ought to lead to similar pricing implications. Consistent with this conjecture, we confirm later that the CRSP value-weighted index and the S&P 500 produce very similar pricing predictions and that the value-weighted indices result in smaller pricing errors. The Dow Jones and the NYSE indices do slightly worse, possibly because of reduced market coverage (e.g., the NYSE misses out on a big part of the technology sector). Consequently, we choose to define our benchmark stock index volatility estimates using the daily returns of the CRSP value-weighted index.

It is important to note that it is the unforecastable volatility innovations and not the level of volatility itself that is priced.[7] In order to identify innovations, one must first take a stand on a reasonable forecasting model. Ang, Hodrick, Xing, and Zhang (2006) use first-differences of stock index volatility as innovations, which implies a random walk forecasting model. In our data, the random walk model is inferior to the ARMA(1,1) in terms of forecasting the log volatility of the CRSP value-weighted index, and the pricing performance of the random walk innovations is significantly poorer than that of ARMA(1,1) innovations. In Table 1, we justify the choice of the ARMA(1,1) model by its superior average out-of-sample volatility forecast performance.

Specifically, for each stock and each month between January 1962 and December 2006, we calculate realized volatility as the sum of squared daily returns. We then compute

[6] It would make a lot less sense to price, say, portfolios of Treasuries. In that case, other components of volatility could be considered. For instance, yield curve volatility innovations can be extracted from options on Treasury futures or from the swap curve based on swaptions, an exercise that we do not undertake.

[7] Figure 1 shows the time series of realized index volatility versus the options-implied index volatility. The latter of course includes a risk premium, but nonetheless it is clear that the unanticipated component is sizable and that investor expectations appear to react with a lag.

one-month-ahead out-of-sample forecast errors of log realized volatilities based on 60-month moving windows. We do not consider more sophisticated higher-frequency volatility estimates since the lack the pre-requisite intra-day data would severely curtail our sample size. The forecast model is fitted according to jump-filtered data to avoid dependence on outliers, but the forecast errors are not filtered. We thus have for each stock a time series of forecast errors under each model. We report the cross-sectional average and median R^2 from regressing the realization of log realized volatility onto its forecast. We also report the fraction of stocks for which the intercept of this regression is insignificant and the slope insignificantly different from one. Finally, we report the cross-sectional average and median mean squared error (MSE). We conclude that ARMA(1,1), which has the highest R^2 and the lowest MSE, seems to be the best model specification (on average) for computing expected future volatility. We therefore use ARMA(1,1) and the corresponding one-period-ahead innovations in (log) realized volatilities as our measures of aggregate volatility innovations, denoted as $mktvol_{inno}$.

An alternative approach, is to impose a parametric assumption about the joint behavior of stock market returns and volatility (e.g., GARCH, EGARCH) and extract the market volatility innovations as the in-sample model residuals. The use of in-sample residuals, however, has the drawback of inducing a potential for look-ahead bias that must be controlled for in subsequent asset pricing tests. Moreover, we do not want to build-in any parametric dependence between first and second moments. We sidestep these issues by instead relying on out-of-sample forecast errors from a parsimonious (univariate) forecasting model. We should emphasize that the choice of the ARMA(1,1) is deliberately not made based on any ex-ante expectation of "optimal" pricing performance but merely represents what in our sample appears to be a reasonable filter for extracting a measure of unforecastable innovations.

3.2 Volatility factors from the cross-section

As an alternative to picking an arbitrary index, one can extract information about aggregate volatility from the cross-section of individual stock volatility innovations. There are at least two reasons for conducting such an analysis. First, it allows us to confirm that the stock index volatility innovations are indeed the principal common driver of individual stock volatility innovations, even though an index represents only a small subset of all stocks (except for the CRSP index) and is subject to time-varying portfolio weights. Second, we can entertain the possibility that more than one component of volatility is priced in a linear beta pricing setting.[8]

We start with a cross-section of one-period-ahead ARMA(1,1) innovations in the (log) realized volatilities of individual stocks. We then extract the principal components using the approximate principal component analysis (APCA) of Connor and Korajczyk (1988). We conduct the APCA in a rolling window 60 months long. We shift the rolling window one month at a time to extract the entire time series of the principal components from January 1967 through December 2006. Appendix B provides details on the estimation procedure. There are of course many alternative methods of extracting systematic volatility innovations from the cross-section of individual stock volatilities. The method we propose is only one of many possible, and is not intended to be optimal in any statistical sense but merely intuitive and easy to implement.

Table 2 shows that the number of significant factors and their explanatory power remain remarkably constant over most of the 1967-2006 sample period. The first component is clearly dominant; it consistently explains at least 15%-25% of the cross-sectional variation in volatility innovations with the notable exception of the early 1990s. The Bai and Ng (2002) information criteria consistently suggest the presence of only one statistically significant factor, so we will focus on the pricing implications of this factor. The second and third factors are clearly less important statistically, accounting for less than 3%-4% of the cross-sectional variation each, but could still be priced in the cross section.[9] We denote these first three principal components as $F1$, $F2$, and $F3$. The time series of the first principal component ($F1$) along with the log volatility innovations of the CRSP

[8]We stress that the factors extracted from the cross section of stock return volatility innovations are very different from statistical factors extracted from the cross section of stock returns.

[9]The mapping from statistical importance to pricing power not obvious. For instance, a factor may be statistically important without carrying any pricing implications (e.g. a common component of idiosyncratic volatility). On the other hand, a factor with low statistical significance can be an important pricing factor (e.g. factors capturing long run risk).

value-weighted market ($mktvol_{inno}$) are shown in Figure 2. All volatility factors and innovations throughout are normalized so that the market excess return has a regression slope of -1 when regressed on the factor. The first factor is clearly similar to, but distinct from, the market innovations, with a correlation of 0.77. In fact, it appears to be a smoothed version of the index volatility innovations.

4 The Pricing of Volatility Risk Across Asset Classes

We test the performance of the two-factor (volatility augmented CAPM) in the cross-section by considering pricing of stock portfolios, portfolios of stock options formed to replicate variance swap contracts, and bond index portfolios, and comparing the estimated risk premia.

4.1 Pricing tests using stock portfolios

We use the 25 Fama-French size/book-to-market portfolios as the test assets. Fama and French (1992 and 1996) show that sorting on size and book-to-market ratio generates cross-sectional variation in expected portfolio returns that is not explained by the CAPM. We examine here whether augmenting the CAPM with our volatility risk factor might help. We use the Fama-French three-factor model (1993) as a benchmark for comparison over the sampling period from January 1967 through December 2006.

We estimate the factor loadings by regressing the monthly value-weighted returns on the 25 portfolios on the market excess return factor (MKT) and a volatility risk factor (either $mktvol_{inno}$ or $F1$). The factor loadings are reported in Panel A of Table 3. In line with the previous literature, we find growth stocks to have higher loadings on the MKT factor or CAPM betas. Because value stocks with higher book-to-market ratios earn higher average returns empirically, the CAPM is unable to explain the value premium.

When we examine the factor loadings on the volatility factor, we find that small and value stocks have lower (more negative) volatility factor loadings than big and growth stocks. In addition, the average volatility factor loadings in the cross-section are negative.

This means that, on average, stock portfolios tend to do poorly when volatility increases (consistent with the "leverage effect"), and this is more so for small and value stocks.[10] The patterns in volatility betas are similar whether we measure the aggregate volatility innovations using $mktvol_{inno}$ or $F1$. In their study of equity option pricing, Coval and Shumway (2001) find a similar result: When regressing the excess return of size sorted portfolios on the excess market return and the return to delta neutral straddles, small stocks load more negatively on the straddle returns.

Once the factor loadings are estimated in the time series regressions, we test their pricing in the cross-section using Fama and MacBeth (1973) cross-sectional regressions. Each month, portfolio returns are regressed on the factor loadings. The regression coefficients are then averaged across time to produce estimates of risk premia on the factors. The corresponding t-values are computed after accounting for the first-step estimation error and potential error autocorrelation using the Newey-West correction with 12 lags. Lewellen, Nagel, and Shanken (2010) argue that, when returns follow factor structures, the OLS R^2 from cross-sectional regression may not be a good model performance measure. Following their prescription, we calculate both OLS R^2 and GLS R^2. The results are presented in Panel B of Table 3.

Across all models, we document a negative risk premium on the MKT factor. This is not too surprising, given that value stocks are associated with higher returns but lower CAPM betas. Petkova (2006) obtains a similar finding. One potential explanation is that the market portfolio acts as a hedge against uncertainty in some missing state variables. Consistent with this interpretation, we find a positive and significant intercept term in the cross-sectional regressions across all models including the Fama-French three-factor model.[11] Another possible explanation is that risk premia are simply time varying, as suggested by Smith and Whitelaw (2009). We do not here take a stand on the source of

[10] We note that a mechanical "leverage effect" does not drive the cross-sectional variations in volatility factor loadings. Orthogonalizing the volatility factors on the market excess return factor and then computing the volatility factor loadings using the residual volatility factors yields very similar results throughout our study.

[11] By contrast, Adrian and Rosenberg (2008) report a positive risk premium on the MKT factor, which they obtain by imposing zero pricing error in the cross-section. Specifically, they set the intercept term to be zero in the second-stage cross-sectional regressions, thereby imposing correct pricing of the risk-free asset. Due to this difference, their volatility risk premia estimates are not directly comparable to ours.

the possible mis-specification, but merely note that the two-factor model is a convenient and parsimonious model for studying the relation between volatility risk and expected returns.

As expected, the CAPM does not seem to explain return variation across the 25 portfolios over our sample period. The adjusted OLS R^2 is a meager 0.174, and the GLS R^2 is even lower at 0.146. After including a volatility factor, the resulting two-factor model does a remarkable job of explaining the returns on the 25 portfolios. When we use $F1$ as the measure of aggregate volatility risk, the adjusted OLS R^2 jumps to 0.841 and the GLS R^2 to 0.405. More important, the risk premium on the volatility factor is negative and significant (-0.0045 per month with a t-value of -3.33). We obtain similar improvements in R^2s and a significantly negative volatility risk premium when we measure aggregate volatility risk using $mktvol_{inno}$. Interestingly, the 25 portfolios do not appear to load on the second and third volatility principal components ($F2$ and $F3$) implying that there is little evidence that additional components of volatility are priced in the stock sample.

The performance of our two-factor model is comparable to that of the Fama-French three-factor model in explaining the returns on the 25 portfolios. Over the same sampling period, the Fama-French three-factor model has a slightly lower adjusted R-square of 0.758. In Figure 3 we graph this result, for simplicity only considering $F1$ as the aggregate volatility risk measure. Both our two-factor model (MKT and $F1$) and the Fama-French three-factor model do a good job in fitting the cross-sectional variation in average excess returns across the 25 portfolios. However, the size and book-to-market factors (SMB and HML) are purely technical factors without clear economic interpretation while the volatility risk factor has a more direct interpretation as a state variable in an ICAPM framework.

4.2 Pricing tests using synthetic variance swaps

In this subsection, we examine the pricing of volatility risk in a cross-section of returns on synthetic variance swaps constructed using portfolios of equity options. One complication of working with synthetic variance swaps is that the panel is generally unbalanced. This occurs both because the number of stocks with liquid option chains increased over our sampling period but also because the variance swap construction requires that a wide range of (liquid) strikes be available which may not be the case in a given month. The result is that there can be large gaps in the return data for a given swap contract. To overcome this problem, we choose to work with log returns which will allow us to calculate factor betas using the innovation in the realized volatility of the underlying stock rather than the (at times unobserved) return on the swap contract itself.

4.2.1 Pricing framework

Let SW_t denote the swap rate determined at time t for a contract that pays an amount RV_{t+1} at time $t+1$, which is equal to the realized variance between t and $t+1$. Applying the pricing formula with the stochastic discount factor (SDF) M_{t+1} and denoting by m_{t+1} the log SDF, we have:

$$\begin{aligned}
SW_t &= E_t[M_{t+1} RV_{t+1}] \Longrightarrow \\
SW_t &= E_t[\exp(m_{t+1} + \log RV_{t+1})] \\
\log SW_t &\approx E_t[m_{t+1}] + E_t[\log RV_{t+1}] + \frac{1}{2}\mathbb{V}\mathrm{ar}_t[m_{t+1}] + \frac{1}{2}\mathbb{V}\mathrm{ar}_t[\log RV_{t+1}] \\
&\quad + \mathbb{C}\mathrm{ov}_t\left[m_{t+1}, \log RV_{t+1}\right]
\end{aligned} \tag{6}$$

Denote the one-period gross risk-free return by $R_{f,t+1}$ and $r_{f,t+1} = \log R_{f,t+1}$, we have:

$$\begin{aligned}
1 &= E_t[M_{t+1} R_{f,t}] \Longrightarrow \\
0 &\approx E_t[m_{t+1}] + r_{f,t} + \frac{1}{2}\mathbb{V}\mathrm{ar}_t[m_{t+1}].
\end{aligned} \tag{7}$$

Combining (6) and (7), we obtain the pricing equation:

$$E_t\left[\log\left(\frac{RV_{t+1}}{SW_t}\right)\right] - r_{f,t} + \frac{1}{2}\mathbb{V}\mathrm{ar}_t[\log RV_{t+1}] = -\mathbb{C}\mathrm{ov}_t\left[m_{t+1}, \log RV_{t+1}\right].$$

If we consider a linear beta pricing model where $m_{t+1} = a_t - b_t' F_{t+1}$ and ignore any potential jump component in the volatility, the excess return of the variance swap (after the convexity adjustment term) is linear in volatility factor betas:

$$E_t\left[\log\left(\frac{RV_{t+1}}{SW_t}\right)\right] - r_{f,t} + \frac{1}{2}\mathbb{V}\mathrm{ar}_t[\log RV_{t+1}] = \beta'\lambda_t, \tag{8}$$

$$\begin{aligned}
\beta_i &= \mathbb{C}\mathrm{ov}_t(F_{i,t+1}, \log RV_{t+1})/\mathbb{V}\mathrm{ar}_t(F_{i,t+1}),\\
\lambda_{i,t} &= \mathbb{V}\mathrm{ar}_t(F_{i,t+1})b_{i,t}.
\end{aligned}$$

Variance swaps are traded mostly in the over-the-counter (OTC) market where prices are not readily available,[12], but price can be accurately replicated using portfolios of calls and puts discussed in Bondarenko (2004) and Carr and Wu (2009). Carr and Wu (2009) shows that the variance swap contract can be replicated by a continuum of positions in out-of-the-money (OTM) calls and puts:

$$\begin{aligned}
SW_i &= E^Q[RV_i]\\
&\doteq \frac{2e^{rt}}{T-t}\int_{-\infty}^{F}\frac{P(K)}{K^2}dK + \int_{F}^{\infty}\frac{C(K)}{K^2}dK.
\end{aligned} \tag{9}$$

The weight on an option with a strike of K is $w(K) = \frac{2e^{rt}}{(T-t)K^2}$. Equation (9) can be estimated accurately by interpolating the implied volatility surface and using numerical integration. In a cross-sectional regression, Carr and Wu (2009) test a similar version of (8) using only one volatility factor — the volatility on the S&P 500 index. Working with a small sample of 35 stocks and 5 indices, they document a negative risk premium. We test (8) directly using aggregate volatility innovations and the volatility factors we extracted using the APCA in a larger sample.

4.2.2 Empirical results

We obtain options data from the OptionMetrics' Ivy database. Between 1997 and 2006, on the Monday after the third Friday in each month, we retain the options that mature in

[12]The market for variance swaps has grown in size dramatically. According to Richard Carson, Deutsche Bank's London-based global head of structured products trading, the market for variance swaps was more than €1 billion in vega in 2005, which represents about €300 billion of equivalent options notional.

the next month that have at least 2 out-of-the-money calls and 2 out-of-the-money puts and positive trading volume.[13] We include 10 stock index options from the major index list in OptionMetrics. The underlying stock indices are listed in Panel A of Table 4. The number of individual stocks included in the sample per month increases from 50 in 1997 to more than 180 in 2006 (see Table 4, Panel B). The number of traded strikes per option for individual stock options average around 6. The range of strikes on which index options are traded is much wider; the average number of traded strikes per option is above 20. For each option, OptionMetrics provides its implied volatility, adjusted for dividends and the American exercise feature. We use these option-implied volatilities to compute the implied variance swap rate SW_i using (9).

As we are dealing with an unbalanced panel, we test the factor pricing model (8) using the Fama and MacBeth (1973) regression approach. In each month t and for stock i with variance swap rate $SW_{i,t}$, we compute the stock's factor betas by regressing $V_{i,t+1}$ on the factors in a five-year rolling window.[14] The conditional variance of realized volatility - $\mathbb{V}\mathrm{ar}_t[\log RV_{i,t+1}]$ is the in-sample ARMA(1,1) residual variance over the 60 months leading up to time t for each individual stock. In each month t, we then run a cross-sectional regression:

$$\log\left(\frac{RV_{i,t+1}}{SW_{i,t}}\right) - r_{f,t} + \frac{1}{2}\mathbb{V}\mathrm{ar}_t[\log RV_{i,t+1}] = \lambda_{0,t} + \widehat{\beta}'_{i,t}\lambda_t + u_{i,t}$$

and finally compute the time series average of $\lambda_{0,t}$, λ_t and the associated t-values. We also compute the Newey-West corrected t-values, which account for the autocorrelation of the estimates with up to 12 lags .

The regression results are provided in Panel A of Table 5. For all volatility factor models, the intercept terms of the regressions do not significantly differ from zero. It follows that we cannot reject the factor models (8). In our two-factor model with the two factors being market excess return (MKT) and the volatility innovation on the CRSP

[13] We choose the Monday after the third Friday because options trading volume is much higher then due to contract rollover. We skip the year 1996 because there are few options in each cross-section, and there are no options data on most stock indices.

[14] We require a stock to have a minimum of 24 months of data to be included in the rolling window regression.

value-weighted stock index ($mktvol_{inno}$), the volatility factor ($mktvol_{inno}$) carries a significant negative risk premium: about -58 basis points per month (t-value = -3.53). When we measure the aggregate volatility risk using the first principal component ($F1$) in the two-factor model, we find that $F1$ also carries a significant negative risk premium. The fact that $F1$ and $mktvol_{inno}$ produce similar results confirms the robustness of our two-factor model to choice of the aggregate volatility risk measure. For brevity, we again focus on $F1$ as the aggregate volatility risk measure for the rest of this subsection.

When we add the second volatility principal component ($F2$), the negative risk premium on $F1$ becomes even more significant. $F2$, while statistically important, is not significantly priced. Finally, when we add the third volatility principal component ($F3$), both $F1$ and $F3$ are priced with negative risk premia. Overall, both the first and the third volatility factor seem to be important in pricing the cross-section of variance swap contracts. This suggests that additional volatility risk factors, while less important in pricing equity returns, could be helpful in pricing assets whose payoffs are directly tied to volatility.

More important, the risk premium on the first volatility factor is still significantly negative. In addition, we show that volatility risk premium of -0.0022 is similar to that obtained from the stock market. In Panel B of Table 5, we estimate the two-factor model in the 25 Fama-French size and book-to-market sorted portfolios during the same sampling period. The volatility risk premium (coefficient on $F1$) is estimated at -0.0027. A paired t-test fails reject the null hypothesis that these two risk premia are different (p-value of the test is 0.5896, and the Newey-West t-value of -0.49 is also close to zero).

Interestingly, the Fama-French (1993) factors SMB and HML also turn out to be useful in pricing the cross-section of variance swaps. As shown in Panel A, HML carries a significantly positive risk premium (t-value = 2.81), and SMB is associated with a negative risk premium that is close to significant (t-value = -1.62). The SMB and HML factors are constructed to capture average return premia in the underlying equity market; the fact that they are also helpful in pricing volatility-based assets suggests that they likely capture state variable risk in the spirit of Merton's ICAPM, confirming the conjecture by Fama and French (1996). We explore this issue further later.

The regression results can be more intuitively represented using a single cross-sectional regression similar to that used in Carr and Wu (2009). From 1997 through 2006, for each stock and index with more than 70 observations over the 10-year sample period, we compute the time-series averages of the actual excess returns on their variance swaps.[15] We then run a cross-sectional regression of these excess returns on their volatility factor betas. The regression results are plotted in Figure 4. The Fama-French three-factor model yields a respectable adjusted R-square (AR^2) of 0.382. If we use the benchmark two-factor model ($MKT + F_1$) instead, the adjusted R-square improves to 0.45. As in our pricing results, the third volatility factor provides additional explanatory power. A four-factor model has an AR^2 of 0.514. Moreover, the stock indices mostly lie below the 45 degree line, indicating that index options in fact are more expensive than individual stock options. The relative expensiveness of index options shrinks with inclusion of the additional volatility factors in the model, and, in contrast to the results in Carr and Wu (2009), the negative estimate of the volatility risk premium is not driven purely by index options.

Finally, we conduct out-of-sample pricing tests by estimating factor loadings and risk premium during the in-sample period of 1997-2004 and then pricing the average variance swap returns during the out-of-sample period from 2005-2006. We keep in sample stock and index portfolios with more than 50 observations during the in-sample period and more than 16 observations during the out-of-sample period. We use the corresponding average variance swap returns and factor loadings during the in-sample period to estimate factor risk premia. For each stock and index, we then compute a predicted excess return on the variance swap as an inner product between factor betas (computed in the rolling window ending in December 2004) and estimated factor risk premia suggested by the pricing equation (8). We then plot the predicted variance swap excess returns against the actual average excess returns during the out-of-sample period in Figure 5. We also report the root mean square errors (RMSE). As in the earlier results, the benchmark two-factor model does better than the Fama-French three-factor model in predicting future returns on variance swaps, while addition of the additional volatility factors do not appear to help

[15] This leaves us with 46 observations in the cross-section.

at the margin.

If a second and/or third volatility factor ($F2$ or $F3$) is indeed priced, it should ideally have an economic interpretation. To access this question, we study the pattern of pricing errors between the two-factor model ($MKT + F1$) and the four-factor model $MKT + F1 + F2 + F3$). Figure 7 shows the change in the relative pricing error for each index option portfolio as well as the average change in absolute relative pricing error across GICS categories for individual stock options. The average absolute two-factor pricing errors are shown as well. It appears that the additional volatility factors do help improve the pricing of index options across the board, but are particularly helpful in reducing the pricing error of the small-cap stock index (RUT). The pattern of changes in pricing errors across sectors one the other hand is less clear-cut, with a somewhat stronger impact on health care and technology. One interpretation of the (weak) evidence of additional priced volatility factors is that an additional small-cap volatility risk factor is needed to explain the cross-section of option portfolio returns. This may also be related to the non-linear relation between size and volatility risk sensitivity documented in Panel A of Table 3.

4.3 Pricing tests using corporate bond index portfolios

Finally, we also examine total returns on Lehman Brothers US corporate bond index portfolios across different maturities and credit rating categories.

Bond index returns from April 1990 through December 2006 are obtained from Datastream. We exclude bond indices with missing returns during the sampling period. This leaves us with 19 corporate bond index portfolios described in Panel A of Table 6. An intermediate bond index portfolio includes bonds with maturities shorter than 10 years and a long bond index portfolio includes bonds with maturities longer than 10 years.

The composition and the duration of the bond index portfolio changes over time. To minimize the impact of time-varying duration on the asset pricing test, we compute the excess returns on the bond index portfolio by taking the difference between the total return on the bond index and the return on a portfolio of Treasury STRIPs constructed with matching duration. Since this duration match is performed at the beginning of each month, the resulting excess returns will be less affected by the term structure of interest rates but are of course still subject to the term structure of credit risk and liquidity.

We estimate the factor loadings by regressing the monthly excess returns on the 19 corporate bond portfolios on the market excess return factor (MKT) and a volatility risk factor (either $mktvol_{inno}$ or $F1$). The average excess returns and factor loadings are reported in Panel A of Table 6. As expected, corporate bonds with lower credit ratings earn higher average returns after adjusting for duration effects. The factor loadings on the volatility risk factor ($mktvol_{inno}$ or $F1$) is negative across all bond portfolios, indicating that the corporate bond return is lower during volatile periods. Corporate bonds with lower credit ratings also have more negative volatility risk betas. This is especially true when we use $F1$ as the measure of aggregate volatility. As a result, the average bond returns are almost perfectly negatively correlated with volatility risk betas in the cross-section, suggesting a negative volatility risk premium.

The negative volatility risk premium is confirmed in cross-sectional regressions. Each month, portfolio excess returns are regressed on the factor loadings. The regression coefficients are then averaged across time to produce estimates of risk premia on the factors. t-values are computed after accounting for estimation error in factor loadings and also error autocorrelation using the Newey-West formula with 12 lags. The results are presented in Panel B of Table 6. The risk premium on the volatility factor is indeed negative (-0.0016 per month using $F1$). We obtain similar results with $mktvol_{inno}$ as the aggregate volatility risk measure, and therefore will focus on $F1$ for the rest of this subsection.

The sampling period is short, and high-yield bond returns were extremely volatile during the period, so the risk premium estimate is not significant. In fact, no factor is significant across all models during this sampling period. In addition, the volatility risk premium of -0.0016 is similar in magnitude to that obtained from the stock market during the same sampling period. In Panel C of Table 6, we estimate the two-factor model in the 25 Fama-French size and book-to-market sorted portfolios for the same sampling period. The volatility risk premium (coefficient on $F1$) is estimated to be -0.0018. A paired t-test fails to reject the null hypothesis that the two risk premia are different (p-value of the test is 0.9609 and the Newey-West t-value of -0.04 is also close to zero).

Most of the models, including the two-factor model, do a good job of explaining the cross-sectional variations in the average excess returns across the 19 bond portfolios. The adjusted OLS R^2 of the two-factor model (MKT and $F1$) is 0.937 (the GLS R^2 is 0.819). Adding two additional volatility risk factors does not improve the R^2 much. The Fama-French three-factor model has an adjusted R^2 of 0.967 (the GLS R^2 is 0.857).[16] Figure 6 plots this result. Both our two-factor model and the Fama-French three-factor model do a good job in fitting the cross-sectional variation in the excess returns across the 25 portfolios.

Overall, we find the two-factor model to be comparable to the Fama-French three factor model in pricing the cross-section of corporate bond returns. The volatility risk premium is also negative in the bond market and of a magnitude comparable to that in the stock market. However, the risk premium is not significant for corporate bonds, due to the fact that the corporate bond returns are extremely volatile during the short-sample period we examine.

4.4 Interpreting the Fama-French factors

Throughout our tests, we find that the Fama-French (1993) three-factor model performs much like our two-factor model in pricing average returns in the stock, option, and bond markets. Why does the Fama-French three-factor model perform so well across asset classes? The SMB and HML factors often construed as purely technical factors designed to capture the cross-sectional variation in returns on book-to-market and size-sorted portfolios, but not to price delta-neutral option portfolios.

Fama and French (1996) conjecture that their SMB and HML factors may be capturing state variable risk in the spirit of Merton's ICAPM. If so, it is not surprising that the Fama and French (1993) three-factor model performs consistently well across asset classes despite the fact that the factors are constructed using stock returns. To the extent that our two-factor model is a more direct implementation of Merton's ICAPM, we would expect

[16]The high R^2 occurs because the returns on the 19 bond portfolios have a strong two-factor structure, with the first factor (accounting for about 70% of the variation) the level of the corporate-Treasury spread, and the second (accounting for roughly 20% of the variation) the investment-grade - non-investment-grade spread.

comparable asset pricing performance between the three-factor model and our two-factor model. Both Fama-French factors and the volatility factor could be capturing the state variables with a significant amount of noise, leading to low correlations among the factors. But for them to have similar asset pricing performance as implied by the CAPM, we would expect factor loadings of test assets on these two set of factors to be closely related.

Panel A of Table 3 appears to indicate that there is a systematic relation between size and book-to-market ratios on the one hand and sensitivity to volatility risk on the other. We saw in particular that small stocks and value stocks are more sensitive to volatility risk than large stocks (and hence should earn a higher return because the loading is negative). This suggests that there may be a relation between volatility betas and HML and SMB betas. This is explored in Figure 8. In Panel (a) we ask whether the volatility betas of the 25 stock portfolios can be explained by their SMB and HML loadings. The answer is yes, the adjusted R-square is around 90%, which explains why the two-factor model and the three-factor model do about equally well in this sample.

Panels (b)-(c) show that the SMB loading is most closely related to volatility beta, but that the relation is not a univariate one; using HML and SMB loadings together yields a significant boost over using the SMB loading alone. This is reinforced by looking at option portfolios in Figure 9. Here the relation is more noisy. The HML and SMB loadings explain only about 28% of the cross-sectional variation in the principal volatility betas. Interestingly, the HML and SMB betas also seem to span the loadings on the third volatility factor ($AR^2 \approx 34\%$) but not the betas of the second volatility factor, which does not help in pricing (compare with Figure 5). Correspondingly, the three-factor model has somewhat less explanatory power than the two-factor model, a respectable adjusted R-square of 38% versus 45% and 51% for the models with one and three volatility factors, respectively. The HML loading is most closely related to the volatility beta, while the SMB loading appears unrelated but helps explain the loadings on the third volatility factor; the opposite was true in the stock sample. Part of the reason for this difference may be that the option portfolios fail to produce a clear spread in the SMB loadings, leading to very imprecise point estimates.[17]

[17]In addition, the betas in the swap sample are based on 5-year rolling window regressions due to the extreme

More formally, we compute the components of SMB and HML factor loadings that are orthogonal to the volatility factor loadings. In the cross-section, we regress the factor loadings on SMB on the factor loadings on our volatility factor $F1$. SMB_{Res} refers to the residuals of this regression. We also regress the factor loadings on HML on the factor loadings on the volatility factor $F1$, and denote the residual as HML_{Res}. We then include these residual Fama-French factor loadings in the Fama-MacBeth cross-sectional regressions. The cross-sectional regressions are computed separately for the sample of Fama-French 25 stock portfolios (1967:01-2006:12), the sample of 19 corporate bond index portfolios (1990:05-2006:12), and the sample of variance swaps (1997:01-2006:12). The results are presented in Table 7.

We find that the residual factor loadings on both SMB and HML add little incremental explanatory power in the cross-section. In the stock sample, neither SMB_{Res} nor HML_{Res} is significant, and the adjusted R^2 of the regression actually drops (from 0.841 to 0.820). In the bond sample, neither SMB_{Res} nor HML_{Res} is significant again, and the adjusted R^2 of the regression improves by less than 3% (from 0.937 to 0.964). In the sample of variance swaps, the increase in adjusted R^2 is slightly higher at 6.8% (from 0.450 to 0.518) and both SMB_{Res} and HML_{Res} are significant. This may be explained by the fact that SMB and HML span the third volatility factor, and indeed the three volatility factor model has a comparable adjusted R^2 of 51.4%.

4.5 Alternative volatility specifications

So far we have used the volatility innovations on the CRSP value-weighted stock index $mktvol_{inno}$ and the principal component $F1$ as the proxies for volatility risk in the pricing tests, and they perform very similarly. We can also evaluate of alternative specifications based on observed stock index volatilities. In Table 8 we show the pricing performance of two-factor models using stock and options portfolios where volatility risk carries a significant negative risk premium.

In the options sample, all the broad-based market indices perform much the same

unbalanced nature of the panel. This leads to fairly noisy beta estimates, which may account for the weaker spanning in the option sample versus the stock sample.

as the volatility factor $F1$, provided innovations are constructed as ARMA(1,1) forecast errors. If first-differences are used instead, the explanatory power is significantly reduced, but the estimated risk premium is negative and significant in all cases. In the stock sample, a slightly different picture emerges. All the indices have roughly the same explanatory power (although the Dow Jones index does slightly more poorly), but all have lower levels of significance than $F1$. One reason for this is evident in a look at Figure 2. The principal volatility factor $F1$ tends to be much less noisy than the market index-based factors, and what matters most for pricing are low-frequency persistent movements rather than high-frequency transitory movements, a point raised in Adrian and Rosenberg (2008). In all cases, however, the point estimate of the volatility risk premium is negative (although not significant) when first-differences are used instead of ARMA(1,1) innovations.

We therefore conclude that the two-factor model (volatility-augmented CAPM) is robust to the choice of the broad-based market index used to proxy for market volatility.[18] It is, however, of first-order importance how innovations are defined. A naive first-difference approach appears to work well in pricing book-to-market and size-sorted stock portfolios but falls short in pricing options. This once again illustrates the importance of considering multiple asset classes.

5 Conclusion

We find that a simple ICAPM-inspired two-factor model combining a measure of volatility risk and the market excess return factor performs well in explaining the cross-section of returns across portfolios of stocks, options, and corporate bonds. Using this coherent pricing framework, we find that volatility risk is priced consistently, with a negative risk premium of similar magnitude, across diverse portfolios of different assets. This suggests that stock market volatility risk, as captured by the range of proxies we consider, acts as a state variable in the ICAPM sense, and is not just yet another statistical pricing factor.

The Fama-French three-factor model performs surprisingly well even in pricing port-

[18]This may not appear surprising given the high correlation between these indices, but high correlation of returns does not guarantee equally high correlation of second moments and does not ex-ante guarantee equal pricing performance.

folios of delta-neutral options. We provide direct evidence suggesting that the SMB and HML factors are proxying for market volatility risk, thus supporting the Fama and French (1996) conjecture that the Fama-French factors are indeed capturing state-variable risk.

We propose a "non-parametric" proxy for aggregate volatility risk by applying principal component analysis to the cross-section of realized monthly volatility innovations of US equities. Our results show that the principal factor is highly correlated with the various market index volatility innovations although it is appear to be substantially less noisy and thus yields the best pricing performance of the volatility measures considered. One additional advantage of the principal component analysis is that it allows us to entertain the possibility that more than one component of stock market volatility is priced. In the pricing of synthetic volatility swaps, for example, it emerges that at least one additional volatility factor may be helpful in explaining the cross-section of returns.

References

Adrian, T. and J. Rosenberg (2008). Stock returns and volatility: pricing the short-run and long-run components of market risk. *Journal of Finance 63*(6), 2997–3030.

Amengual, D. and M. W. Watson (2007). Consistent estimation of the number of dynamic factors in a large n and t panel. *Journal of Business and Economic Statistics 25*(1), 91–96.

Andersen, T. G., T. Bollerslev, P. F. Christoffersen, and F. X. Diebold (2005). Volatility forecasting. *NBER Working Paper* (11188).

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang (2006). The cross-section of volatility and expected returns. *Journal of Finance 61*(1), 259–99.

Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica 70*(1), 191–221.

Bakshi, G. and N. Kapadia (2003a). Delta hedged gains and the negative volatility risk premium. *Review of Financial Studies 16*, 527–566.

Bakshi, G. and N. Kapadia (2003b). Volatility risk premium embedded in individual equity options: Some new insights. *Journal of Derivatives*, 45–54.

Bali, T. G. (2008). The intertemporal relation between expected returns and risk. *Journal of Financial Economics 87*(1), 101–131.

Bollerslev, T., J. Marrone, L. Xu, and H. Zhou (2011). Stock Return Predictability and Variance Risk Premia: Statistical Inference and International Evidence. *SSRN eLibrary*.

Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected stock returns and variance risk premia.

Bondarenko, O. (2004). Market price of variance risk and performance of hedge funds. *Working Paper*.

Boudoukh, J., M. Richardson, and R. F. Whitelaw (2008). The myth of long-horizon predictability. *Review of Financial Studies 21*(4), 1577–1605.

Breen, W., L. Glosten, and R. Jagannathan (1989). Economic significance of predictable variations in stock index returns. *Journal of Finance 44*, 1177–1189.

Brennan, M. J., A. W. Wang, and Y. Xia (2004). Estimation and test of a simple model of intertemporal asset pricing. *Journal of Finance 59*(4), 1743–1776.

Campbell, J. Y. (1993). Intertemporal asset pricing without consumption data. *American Economic Review 83*, 487–512.

Campbell, J. Y. and S. B. Thompson (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies 21*(4), 1509–1531.

Carr, P. and L. Wu (2009). Variance risk premia. *Review of Financial Studies 22*(3), 1311–1341.

Cochrane, J. H. (2008). The dog that did not bark: A defense of return predictability. *Review of Financial Studies 21*(4), 1533–1575.

Connor, G. (1984). A unified beta pricing theory. *Journal of Economic Theory 34*(1), 13–31.

Connor, G. and R. A. Korajczyk (1988). Risk and return in an equilibrium apt: application of a new test methodology. *Journal of Financial Economics 21*, 255–289.

Coval, J. D. and T. Shumway (2001). Expected option returns. *Journal of Finance* (56 (3)), p 983–1009.

Cox, J. C. and S. A. Ross (1976). The valuation of options for alternative stochastic processes. *Journal of Financial Economics 3*(1/2), 145–66.

DeLisle, J., S. M. Price, and C. F. Sirmans (2011). Pricing of volatility risk in reits. *working paper*.

Deng, Q. (2008). Volatility Dispersion Trading. *SSRN eLibrary*.

Duarte, J. and C. S. Jones (2007). The price of market volatility risk. *Working Paper*.

Dybvig, P. H. and J. Ingersoll, Jonathan E. (1982). Mean-variance theory in complete markets. *Journal of Business 55*(2), 233–51.

Fama, E. and K. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*, 3–56.

Fama, E. and K. French (1996). Multifactor explanations of asset pricing anomalies. *Journal of Finance 51*, 55–84.

Fama, E. and J. D. MacBeth (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy 81*, 607–636.

Ferson, W. E., S. Sarkissian, and T. Simin (1999). The alpha pricing factor model: A parable. *Journal of Financial Markets* (2).

French, K. R., G. W. Schwert, and R. F. Stambaugh (1987). Expected stock returns and volatility. *Journal of Financial Economics 19*(1), 3–29.

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

Guasoni, P., G. Huberman, and Z. Wang (2011). Performance maximization of actively managed funds. *Journal of Financial Economics 101*, 574–595.

Guo, H. and R. F. Whitelaw (2006). Uncovering the risk-return relation in the stock market. *Journal of Finance 61*(3), 1433–63.

Hamilton, J. D. and G. Lin (1998). Stock market volatility and the business cycle. *Journal of Applied Econometrics 11*(5).

Harvey, C. (2001). The specification of conditional expectations. *Journal of Empirical Finance 8*, 573–637.

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

Lewellen, J., S. Nagel, and J. Shanken (2010). A skeptical appraisal of asset pricing tests. *Journal of Financial Economics*, 175–194.

Liewa, J. and M. Vassalou (2000). Can book-to-market, size and momentum be risk factors that predict economic growth? *Journal of Financial Economics 57*(2), 221–245.

L.R. Glosten, Ravi Jagannathan, D. R. (1993). On the relation between the expected value and the volatility of the nominal excess returns on stocks. *Journal of Finance 48*, 1779–1801.

McKinlay, A. C. (1995). Multi-factor models do not explain deviations from the capm. *Journal of Financial Economics* (38).

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica 41*(5), 867–87.

Moise, C. E. (2010). Volatility pricing in the stock and treasury markets. *working paper*.

Nielsen, L. T. and M. Vassalou (2006). The instantaneous capital market line. *Economic Theory 28*(3), 651–664.

Petkova, R. (2006). Do the fama-french factors proxy for innovations in predictive variables? *Journal of Finance 61*, 581–612.

Roll, R. (1977). A critique of the asset pricing theory's tests: Part i: On past and potential testability of the theory. *Journal of Financial Economics 4*(2), 129–76.

Schwert, G. W. (1989). Why does stock market volatility change over time? *Journal of Finance 44*(5).

Smith, D. R. and R. F. Whitelaw (2009). Does risk aversion change over time? *working paper*.

Vassalou, M. (2003). News related to future gdp growth as a risk factor in equity returns. *Journal of Financial Economics 68*(1), 47–73.

Wang, Z. and X. Zhang (2006). Empirical evaluation of asset pricing models: arbitrage and pricing errors over contingent claims. (265).

Welch, I. and A. Goyal (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies 21*(4), 1455–1508.

Whitelaw, R. (1994). Time variations and covariations in the expectation and volatility of stock market returns. *Journal of Finance 49*, 515–541.

A Non-linear Payoffs and Linear Beta Pricing

In a standard linear beta-pricing framework, non-linearities in payoffs may lead to an inflated covariance with volatility risk and the finding of a volatility risk premium. The reason linear pricing models have difficulty in pricing non-linear payoffs is that positivity of the implied stochastic discount factor is not guarannteed, as Dybvig and Ingersoll (1982) note in the context of the CAPM and Connor (1984) points out in the context of linear factor models.[19]

To see how this mispricing may be confused with a volatility risk premium, consider pricing a non-linear payoff within the standard CAPM model. The strategy consists of buying one unit of the index and shorting θ out-of-the-money put options with strike $K_{\text{put}} < S_0$, where S_0 is the level of the index. This payoff is shown in Figure 10 for $\theta = 1/2$ and can be interpreted as the payoff of a stock whose CAPM beta increases in down markets (abstracting from idiosyncratic risk).

The put contracts has a negative CAPM alphas, implying that this strategy will have a positive alpha (for $\theta > 0$) and that the alpha increases with the number of contracts θ. Denote by $\Delta RV \equiv \int_0^1 \sigma_u^2 du - E\left[\int_0^1 \sigma_u^2 du\right]$ the volatility innovation estimated by the econometrician. In models with "leverage" effects, $\mathbb{C}\text{ov}(R_{\text{stock}}, \Delta RV)$ is negative, although in practice, this effect tends to be small in magniture. The covariance $\mathbb{C}\text{ov}(R_{\text{put}}, \Delta RV)$ is positive. To see this, note that, all else equal, an unexpected increase in volatility is good news for the existing owner of the option. Secondly, to the extent that the leverage effect is present, it is even better news since the index value tends to fall. Thus, as strategy which is long the index and short the put option will have a negative covariance with volatility innovations, and the more negative the larger is θ. This therefore implies a negative relation between the CAPM alpha and the volatility beta in such models leading to the finding of a negative volatility risk premium.

Consider the constant elasticity of variance (CEV) model of Cox and Ross (1976),

[19] For en elegant exposition, seeWang and Zhang (2006) and Guasoni, Huberman, and Wang (2011) in the context of performance evaluation.

which for $\gamma < 2$ displays the leverage effect:

$$dS_t = \mu S_t dt + \sigma S_t^{\gamma/2} dB_t \tag{10}$$

Since volatility is a deterministic function of the stock price, volatility risk is by construction absent. Nonetheless, a linear factor model will mistakenly imply that volatility risk is priced.

To illustrate this, we conduct a Monte Carlo experiment, generating daily stock prices (S_t) from model (10) with parameters $\mu = 8\%/\text{year}, \sigma = 20\%/\text{year}$, and $\gamma \in \{0.0, 0.5, 1.0, 1.5\}$. We calculate the time series of monthly returns (**??**) for $\theta \in [0; 1]$ and the realized volatility as the sum of squared daily returns during each month. From this monthly time series (setting $r_f = 3\%/\text{year}$) the CAPM alpha and the volatility beta are computed as for each level of optionality (θ) in the return. Figure 11 shows the relation between the CAPM alpha and the loading on ΔRV in a two-factor volatility risk augmented CAPM for θ ranging between 0 and 1 and for various values of elasticities γ. In the CEV model, the result is a clear negative relation between CAPM α and the two-factor model volatility beta, leading to a finding of a negative volatility risk premium.

While the CEV model is hardly a realistic model model of asset prices, it does illustrate how small degrees of optionality in returns may lead to a non-trivial covariance with volatility risk and (in the example) an overstated estimated volatility risk premium in a linear factor pricing framework.

B Constructing a Volatility Innovation Index

Construction of a realized volatility innovation index extracted from the cross-section of individual realized stock volatilities follows a two-step procedure. In the first step, we follow common practice in the volatility forecasting literature and determine the volatility innovations stock by stock from a univariate time series of estimated realized volatility (see Andersen, Bollerslev, Christoffersen, and Diebold (2005)). In the second step, the volatility innovations are then extracted from the large panel of individual stock volatility innovations by principal component analysis.[20]

For a given stock i during a given month t, we estimate the realized volatility by

$$RV_{i,t} = \frac{252}{N_t} \sum_{n=1}^{N_t} r_{it,n}^2 \tag{11}$$

where N_t is the number of trading days in month t and $r_{it,n}$ is the return on stock i on day n during month t. For pricing purposes it is convenient to work with the log realized volatility, which has the additional advantage of having a near-Gaussian marginal distribution. Once a time series of realized (log) volatilities is constructed, the task is to extract the volatility innovations as the unforecastable component of month-to-month volatility changes.

We follow common practice in the volatility forecasting literature and base our one-step-ahead forecasts on univariate models. In order to avoid making any strict parametric assumptions about the dynamics of individual stock volatilities, we focus on simple ARMA filtering rules:[21]

$$\Phi(L) \log(RV_{i,s}) \quad = \quad \Theta(L)\varepsilon_{i,s} \quad s = t - 59, \ldots, t \tag{12}$$

where the model is fitted using a moving 60-month window in order to accommodate possible parameter non-stationarities. This estimation procedure also allows us to compute the volatility innovations $\hat{\varepsilon}_{i,t+1}$ as the one-month-ahead *out-of-sample* forecast errors rather than relying on in-sample residuals from the model with the best *in-sample* fit. Table 1 in the text shows that the ARMA(1,1) model dominates alternative specifications across all forecast evaluation measures, and we use it to extract the volatility innovations

[20] An alternative procedure would be first to extract principal components from a panel of realized volatilities and then filter out the innovations to the extracted factors using a VARMA filter. This procedure has the drawback that many factors are typically found in the first step because of strong serial dependence in the data, and there is no easy way to determine the actual number of (non-dynamic) factors. For this reason, we choose not to pursue this alternative strategy here.

[21] A large number of alternative volatility forecasting models have been considered in the literature. We rule out implied volatility based estimates due to insufficient data and forego ARCH/GARCH models in order to avoid building in any assumption about the joint behavior of first and second moments of returns. We also do not report ARIMA and ARFIMA results since the one-step-ahead forecasting performance was too poor based on estimation windows up to 120 months.

(13) throughout:[22]

$$\hat{\varepsilon}_{i,t+1} \quad \stackrel{\text{def.}}{\equiv} \quad \log(RV_{i,t+1}) - \hat{E}_t[\log(RV_{i,t+1})] \tag{13}$$

We assume that the volatility innovations have a common (systematic) component, F_t:[23]

$$\underset{N\times 1}{\hat{\varepsilon}_t} \quad = \quad \underset{N\times k}{\Gamma}\ \underset{k\times 1}{F_t} + \underset{N\times 1}{\eta_t} \quad t = 1,\ldots,T \tag{14}$$

Assuming that the fitted forecast errors are only weakly dependent across time and across stocks, Bai and Ng (2002) show that the Connor and Korajczyk (1988) approximate principal component analysis (APCA) approach can be used to estimate the number of factors and extract the common components from the cross-section of individual stock volatility innovations.[24] A key assumption underlying the factor structure specification (14) is that the factor loadings Γ remain constant over the estimation period for a given firm.

To address the concern of parameter non-stationarity, we opt to extract factors using a rolling 60-month window as follows: The value of the $k \times 1$ vector of common components in month t, F_t, is found by applying the APCA (14) to the (balanced) panel of log volatility forecast errors $\{\hat{\varepsilon}_{i,s}\}_{s=t-59,\ldots,t}$ for the subset of stocks with a complete record of forecast errors during months $t-59,\ldots,t$. This procedure results in estimates $\{F_s\}_{s=t-59,\ldots,t}$ of which we keep only the time t value, F_t. Our time series of factor realizations is thus the result of a sequence of overlapping window factor extractions. Apart from providing us with robustness to parameter non-stationarity, this construction largely eliminates the otherwise severe survivorship bias and enables us to work with much larger cross-sections, as the ACPA procedure requires a continuous record for each stock. Table 2 shows that

[22]Note that the forecast errors (13) need not be serially uncorrelated, although they in practice exhibit very low autocorrelation

[23]There turns out to be significant serial correlation in the extracted principal factor even though the univariate forecast errors are approximately white. This strong predictability of the common component of individual stock volatility forecast errors has yet to be exploited in the volatility forecasting literature.

[24]This leads to consistent estimation of the common components of the true log volatility innovations (as $\min(N,T) \to \infty$) as long as the measurement errors stemming from the first-step ARMA estimation are small, as argued by Amengual and Watson (2007).

only 149 firms are present every month during our 40-year sample period, much lower than the number available in each 5-year window. The only cost is a likely loss of efficiency compared to running a single factor extraction in the event the data is truly stationary. If, on the other hand, the factor loadings Γ are not constant, running a single factor extraction would allow us to detect of additional spurious factors.

C Figures & Tables

Table 1: Out-of-sample model selection criteria. For each stock and each month between January 1962 and December 2006, we calculate the realized volatility as the sum of squared daily returns (subject to at most 5 trading days missing). We then compute the one-month-ahead out-of-sample forecast errors of log realized volatilities based on 60-month moving windows. The forecast model is fitted using jump-filtered data to avoid dependence on outliers, but the forecast errors are not filtered. We thus have for each stock a time series of forecast errors for each model. We report the cross-sectional average and median R^2 from regressing the realization of log realized volatility onto its forecast. We also report the fraction of stocks for which the intercept of this regression is insignificant and the slope insignificantly different from one. Finally we report the cross-sectional average and median mean squared error (MSE).

Model	Average R^2	Median R^2	Fraction H_0: $\alpha=0$ not rejected	Fraction H_0: $\beta=1$ not rejected	Average MSE	Median MSE
ARMA(1,0)	0.181	0.161	0.733	0.752	0.681	0.531
ARMA(0,1)	0.135	0.108	0.683	0.691	0.750	0.582
ARMA(1,1)	***0.211***	***0.198***	***0.744***	***0.762***	***0.647***	***0.498***
ARMA(2,0)	0.194	0.176	0.709	0.734	0.671	0.516
ARMA(0,2)	0.150	0.127	0.644	0.662	0.727	0.565
ARMA(2,1)	0.204	0.187	0.665	0.685	0.658	0.505
ARMA(1,2)	0.201	0.183	0.622	0.648	0.662	0.508
ARMA(3,0)	0.201	0.183	0.661	0.684	0.663	0.507
ARMA(0,3)	0.168	0.148	0.602	0.621	0.709	0.549
ARMA(2,2)	0.191	0.172	0.486	0.516	0.683	0.523
ARMA(3,1)	0.199	0.180	0.571	0.594	0.667	0.511
ARMA(1,3)	0.194	0.176	0.508	0.529	0.678	0.520
ARMA(3,2)	0.187	0.166	0.399	0.412	0.691	0.534
ARMA(2,3)	0.186	0.166	0.403	0.428	0.692	0.535
ARMA(3,3)	0.184	0.161	0.310	0.328	0.701	0.547

Table 2: Results of the Asymptotic Principal Component Analysis (APCA). From January 1967 through December 2006, we run an asymptotic principal component analysis (APCA) in each of the eight five-year windows. We report: the number of stocks in each period; the optimal number of factors according to the Bai and Ng (2002) information criteria; and the percentage of the cross-sectional variance explained by each of the three first factors.

Period	#Stocks	Factor 1 %explained	Factor 2 %explained	Factor 3 %explained	IC1 #factors	IC2 #factors
January 1967 - December 2006	**149**	**26.5%**	**4.2%**	**1.5%**	**2**	**2**
January 1967 - December 1971	621	15.0%	3.5%	2.8%	1	1
January 1972 - December 1976	710	21.9%	3.4%	2.5%	1	1
January 1977 - December 1981	799	15.2%	3.2%	2.7%	1	1
January 1982 - December 1986	976	12.8%	3.4%	2.9%	1	1
January 1987 - December 1991	922	34.0%	2.5%	2.2%	1	1
January 1992 - December 1996	1159	8.1%	3.7%	3.0%	1	1
January 1997 - December 2001	1105	24.1%	3.1%	2.6%	1	1
January 2002 - December 2006	1220	20.7%	4.8%	3.4%	1	1

Table 3: Cross-sectional Regressions on Fama-French 25 size / book- to-market sorted portfolios. Panel A reports factor loadings and the associated t-values on the Fama-French 25 portfolios sorted on size and book-to-market ratio. Panel B reports results from Fama and MacBeth (1973) regressions. Each month, portfolio returns are regressed on the factor loadings. The regression coefficients are then averaged across-time to produce estimates of risk premia on the factors. t-values are computed after accounting for the estimation error in factor loadings and also error autocorrelation using the Newey-West formula of 12 lags. The goodness of fit measures are based on a regression of average returns on average factor loadings. The GLS R^2 is calculated as in Lewellen, Nagel, and Shanken (2010). The sampling period is 1967-2006.

Panel A: Factor loadings

	Factor loading on *MKT*					*t*-value				
	small	2	3	4	big	small	2	3	4	big
growth	1.459	1.447	1.369	1.264	1.013	28.1	37.5	42.4	52.2	56.6
2	1.224	1.164	1.105	1.074	0.956	27.3	37.7	46.8	53.3	54.8
3	1.069	1.017	0.960	0.969	0.851	29.0	36.2	40.7	43.6	40.0
4	0.981	0.959	0.890	0.902	0.793	28.0	33.9	35.4	38.2	32.1
value	1.003	1.028	0.976	0.976	0.821	26.2	30.7	31.0	32.3	25.4

	Factor loading on $\text{mktvol}_{\text{inno}}$					*t*-value				
	small	2	3	4	big	small	2	3	4	big
growth	-0.627	-0.459	-0.333	-0.131	0.255	-3.5	-3.5	-3.0	-1.6	4.2
2	-0.722	-0.538	-0.295	-0.184	0.050	-4.7	-5.1	-3.6	-2.6	0.8
3	-0.594	-0.551	-0.351	-0.140	0.070	-4.7	-5.7	-4.3	-1.8	1.0
4	-0.640	-0.512	-0.272	-0.162	0.188	-5.4	-5.3	-3.1	-2.0	2.2
value	-0.721	-0.630	-0.450	-0.219	0.080	-5.5	-5.5	-4.2	-2.1	0.7

	Factor loading on F1					*t*-value				
	small	2	3	4	big	small	2	3	4	big
growth	-0.912	-0.780	-0.534	-0.227	0.330	-2.4	-2.7	-2.2	-1.3	2.5
2	-1.160	-0.960	-0.661	-0.414	0.009	-3.5	-4.2	-3.8	-2.8	0.1
3	-0.977	-1.079	-0.783	-0.408	0.080	-3.6	-5.2	-4.5	-2.5	0.5
4	-1.108	-1.000	-0.662	-0.382	0.349	-4.3	-4.8	-3.6	-2.2	1.9
value	-1.288	-1.199	-0.836	-0.491	0.176	-4.5	-4.8	-3.6	-2.2	0.7

Panel B: Cross-sectional regression results

CAPM + vol factors

	Intercept	MKT	$\text{mktvol}_{\text{inno}}$	F1	F2	F3	*Adj* R^2 / *GLS* R^2
Coeff.	0.0136	-0.0058					0.174
t-value	*3.54*	*-1.20*					*0.146*
Coeff.	0.0177	-0.0121	-0.0070				0.821
t-value	*4.16*	*-2.38*	*-2.72*				*0.420*
Coeff.	0.0161	-0.0108		-0.0045			0.841
t-value	*3.90*	*-2.18*		*-3.33*			*0.405*
Coeff.	0.0148	-0.0095		-0.0052	0.0001	0.0000	0.837
t-value	*4.07*	*-1.99*		*-3.88*	*0.96*	*-0.04*	*0.418*

Fama-French Three-factor Model

	Intercept	MKT	SMB	HML	*Adj* R^2 / *GLS* R^2
Coeff.	0.0137	-0.0085	0.0020	0.0049	0.758
t-value	*4.38*	*-1.98*	*0.93*	*2.45*	*0.317*

Table 4: Option sample description. Panel A lists the 10 stock indices in the sample taken from the list of major indices in OptionMetrics. Panel B lists the average number of stocks and indices in our sample, and the average number of options associated with each stock/index used in computing the variance swap rate.

Panel A: Stock indices included

Index Name	Ticker
Dow Jones Industrial Average	DJX
NASDAQ 100 Index	NDX
CBOE Mini	MNX
AMEX Major Market Index	XMI
S&P 500 Index	SPX
S&P 100 Index	OEX
S&P Midcap 400 Index	MID
S&P Smallcap 600 Index	SML
Russell 2000 Index	RUT
PSE Wilshire Smallcap Index	WSX

Panel B: Sample description

	Individual Stock Option		Index Option	
year	Num of obs per month	Num of options per underlying	Num of obs per month	Num of options per underlying
1997	50.0	5.8	5.4	22.5
1998	68.5	5.8	5.7	21.0
1999	96.4	6.5	5.7	23.4
2000	147.9	7.4	5.0	26.9
2001	134.3	6.0	4.5	25.5
2002	133.1	5.6	4.6	25.2
2003	122.3	5.4	5.4	21.8
2004	125.5	5.2	5.3	23.7
2005	143.1	5.5	5.9	23.3
2006	182.7	6.1	6.7	30.2
Average	120.4	5.9	5.4	24.3

Table 5: Cross-sectional regression results on variance swaps. Panel A reports the cross-sectional regression results on variance swaps. $F1$ to $F3$ are the first three volatility factors extracted using Asymptotic Principal Component Analysis (APCA). The Fama and MacBeth (1973) cross-sectional regressions are estimated from 1997 to 2006. We report the Newey-West t-values which account for the autocorrelation of the estimates with a lag of 12. Panel B reports the cross-sectional regression results on the 25 Fama-French stock portfolios during the same sampling period, which allows for a direct comparison of volatility risk premium in these two markets. We direct test for the inequality of the volatility risk premia in the two markets. The goodness of fit measures in Panel B are based on a regression of average returns on average factor loadings. The GLS R^2 is calculated as in Lewellen, Nagel, and Shanken (2010). GLS R^2 are not reported in Panel A as the panel is unbalanced.

Panel A: Sample of variance swaps

CAPM + vol factors							
	Intercept	MKT	$mktvol_{inno}$	F1	F2	F3	Adj R^2
Coeff.	0.0671	0.0065	-0.0058				0.450
t-value	*0.98*	*1.16*	*-3.53*				
Coeff.	0.0740	-0.0010		-0.0019			0.450
t-value	*1.20*	*-0.23*		*-2.63*			
Coeff.	0.0567	0.0033		-0.0022	0.0000		0.449
t-value	*0.75*	*0.62*		*-2.72*	*1.05*		
Coeff.	0.0613	-0.0007		-0.0019	0.0285	0.0004	0.514
t-value	*0.84*	*-0.13*		*-2.19*	*1.85*	*-2.26*	
Fama-French Three-factor Model							
	Intercept	MKT	SMB	HML			Adj R^2
Coeff.	-0.0058	0.0186	-0.0046	0.0114			0.383
t-value	*-0.12*	*3.56*	*-1.62*	*2.81*			

Panel B: Fama-French 25 portfolios in the same sampling period

	Intercept	MKT	F1	*Adj* R^2 / *GLS* R^2
Coeff.	0.0116	-0.0070	-0.0027	0.737
t-value	*1.44*	*-0.73*	*-1.57*	*0.454*
p-value of the equal-mean test:			*0.5896*	
Newey-West t-value:			*-0.49*	

Table 6: Cross-sectional Regressions on 19 bond index portfolios. The sampling period is 1990/05-2006/12. Panel A reports average excess returns (duration-adjusted) and factor loadings on the 19 bond index portfolios. Panel B reports results from Fama and MacBeth (1973) regressions. Each month, portfolio returns are regressed on the factor loadings. The regression coefficients are then averaged across-time to produce estimates of risk premia on the factors. t-values are computed after accounting for the estimation error in factor loadings and also error autocorrelation using the Newey-West formula of 12 lags. Panel C reports cross-sectional regression results on the 25 Fama-French stock portfolios during the same sampling period, which allows for a direct comparison of volatility risk premium in these two markets. We directly test for the inequality of the volatility risk premia in the two markets. The goodness of fit measures are based on a regression of average returns on average factor loadings. The GLS R^2 is calculated as in Lewellen, Nagel, and Shanken (2010).

Panel A: Average excess returns and factor loadings

Investment Grade					Non-investment Grade				
Index	excess ret	MKT beta	$mktvol_{inno}$ beta (x100)	F1 beta	Index	excess ret	MKT beta	$mktvol_{inno}$ beta (x100)	F1 beta
AAA all	0.0006	0.046	-0.049	-0.210	B all	0.0019	0.292	-0.367	-0.911
AAA intermediate	0.0005	0.026	-0.030	-0.126	B intermediate	0.0015	0.279	-0.391	-0.915
AAA long	0.0012	0.074	-0.083	-0.378	Ba all	0.0019	0.206	-0.266	-0.826
AA all	0.0008	0.058	-0.046	-0.244	Ba intermediate	0.0017	0.201	-0.239	-0.771
AA intermediate	0.0006	0.031	-0.021	-0.154	Ba long	0.0014	0.194	-0.221	-0.795
AA long	0.0014	0.078	-0.072	-0.497	Caa all	0.0028	0.391	-0.611	-1.532
A all	0.0007	0.072	-0.100	-0.351	Caaintermediate	0.0028	0.366	-0.637	-1.498
A intermediate	0.0006	0.048	-0.069	-0.258	Caa long	0.0034	0.630	-0.352	-1.846
A long	0.0012	0.106	-0.138	-0.535					
Baa all	0.0008	0.104	-0.162	-0.539					
Baa intermediate	0.0006	0.081	-0.141	-0.411					

Table 6 Cont'ed: Cross-sectional Regressions on 19 bond index portfolios.
Panel B: Cross-sectional regression results

CAPM + vol factors

	Intercept	MKT	$mktvol_{inno}$	F1	F2	F3	*Adj* R^2 / *GLS* R^2
Coeff.	0.0005	0.0048					0.899
t-value	*0.71*	*0.60*					*0.819*
Coeff.	0.0004	0.0047	-0.0716				*0.909*
t-value	*0.81*	*0.56*	*-0.21*				*0.838*
Coeff.	0.0003	0.0002		-0.0016			0.937
t-value	*0.45*	*0.01*		*-0.34*			*0.819*
Coeff.	0.0002	0.0029		-0.0010	0.0000	0.0000	0.941
t-value	*0.34*	*0.29*		*-0.32*	*0.28*	*0.03*	*0.868*

Fama-French Three-factor Model

	Intercept	MKT	SMB	HML			*Adj* R^2 / *GLS* R^2
Coeff.	0.0000	0.0105	-0.0113	0.0062			0.967
t-value	*-0.04*	*0.41*	*-0.32*	*0.46*			*0.857*

Panel C: Fama-French 25 portfolios in the same sampling period

	Intercept	MKT	F1	*Adj* R^2 / *GLS* R^2
Coeff.	0.0144	-0.0080	-0.0018	0.661
t-value	*2.52*	*-1.19*	*-1.36*	*0.333*
p-value of the equal-mean test:			*0.9609*	
Newey-West t-value:			*-0.04*	

Table 7: Performance of residual Fama-French factor loadings. In the cross section, we regress the factor loadings on SMB on the factor loadings on our volatility factor $F1$. SMB_Res refers to the residuals of this regression. Likewise, we regress the factor loadings on HML on the factor loadings on our volatility factor $F1$ and denote the residual as HML_Res. We then include these residual Fama-French factor loadings in the Fama-MacBeth cross-sectional regressions. The cross-sectional regressions are done separately in the sample of Fama-French 25 stock portfolios (196701-200612), the sample of 19 corporate bond index portfolios (199005-200612) and the sample of variance swaps (199701-200612).

Fama-French 25 portfolios, 196701-200612						
	Intercept	MKT	F1	SMB_Res	HML_Res	*Adj* R^2
Coeff.	0.0177	-0.0121	-0.0070			0.841
t-value	*4.16*	*-2.38*	*-2.72*			
Coeff.	0.0145	-0.0092	-0.0042	-0.0001	0.0007	0.820
t-value	*5.11*	*-2.48*	*-3.03*	*-0.05*	*0.44*	
Corporate bond index portfolios (duration adjusted), 199005-200612						
Coeff.	0.0003	0.0002	-0.0016			0.937
t-value	*0.45*	*0.01*	*-0.34*			
Coeff.	0.0008	0.0135	0.0025	-0.0165	0.0134	0.964
t-value	*0.75*	*0.81*	*0.46*	*-1.19*	*0.76*	
Variance swaps, 199701-200612						
Coeff.	0.0567	0.0033	-0.0019			0.450
t-value	*0.75*	*0.62*	*-2.63*			
Coeff.	0.0753	0.0110	-0.0019	-0.0073	0.0077	0.518
t-value	*1.12*	*2.27*	*-2.74*	*-2.86*	*2.16*	

Table 8: Performance of alternative volatility specifications. We replace the volatility factor $F1$ by the innovations in specific stock indices. The innovations are either ARMA(1,1) out-of-sample forecast errors (denoted "inno") or first differences (denoted "diff"). The left panel shows the performance based on the option portfolios and the right panel the performance based on the Fama-French 25 stock portfolios. The goodness of fit measures are based on a regression of average returns on average factor loadings. For the FF25 (which is a balanced panel), the GLS R^2 is calculated as in Lewellen, Nagel, and Shanken (2010).

		Volatility Swaps (1997-2006)				FF25 portfolios (1967-2006)			
Model		intercept	mkt	volfactor	Adj R^2	intercept	mkt	volfactor	Adj R^2/GLS R^2
F_1	coeff	0.0671	0.0065	-0.0019	0.450	0.0161	-0.0108	-0.0045	0.841
	t-value	*0.98*	*1.16*	*-2.63*		*3.90*	*-2.18*	*-3.33*	0.405
$\text{mktvol}_{\text{inno}}$	coeff	0.0739	-0.0013	-0.0058	0.450	0.0177	-0.0121	-0.0070	0.821
	t-value	*1.20*	*-0.23*	*-3.53*		*4.16*	*-2.38*	*-2.72*	0.420
$\text{NYSEvw}_{\text{diff}}$	coeff	0.0429	0.0019	-0.0057	0.317	0.0210	-0.0156	-0.0063	0.824
	t-value	*0.81*	*0.29*	*-3.81*		*2.89*	*-1.91*	*-1.75*	0.402
$\text{sp500}_{\text{diff}}$	coeff	0.0404	0.0021	-0.0050	0.286	0.0198	-0.0144	-0.0064	0.829
	t-value	*0.78*	*0.32*	*-3.93*		*3.13*	*-2.00*	*-1.99*	0.425
dji_{diff}	coeff	0.0195	0.0094	-0.0023	0.235	0.0183	-0.0126	-0.0050	0.771
	t-value	*0.36*	*1.59*	*-2.59*		*2.00*	*-1.23*	*-1.43*	0.340
$\text{NYSEvw}_{\text{inno}}$	coeff	0.0745	-0.0012	-0.0059	0.492	0.0173	-0.0119	-0.0065	0.826
	t-value	*1.30*	*-0.21*	*-4.39*		*4.07*	*-2.33*	*-2.76*	0.420
$\text{sp500}_{\text{inno}}$	coeff	0.0701	-0.0010	-0.0057	0.473	0.0175	-0.0121	-0.0067	0.827
	t-value	*1.25*	*-0.17*	*-4.28*		*4.10*	*-2.36*	*-2.75*	0.433
dji_{inno}	coeff	0.0483	0.0071	-0.0028	0.449	0.0159	-0.0099	-0.0052	0.759
	t-value	*0.76*	*1.31*	*-3.07*		*3.41*	*-1.81*	*-2.74*	0.317

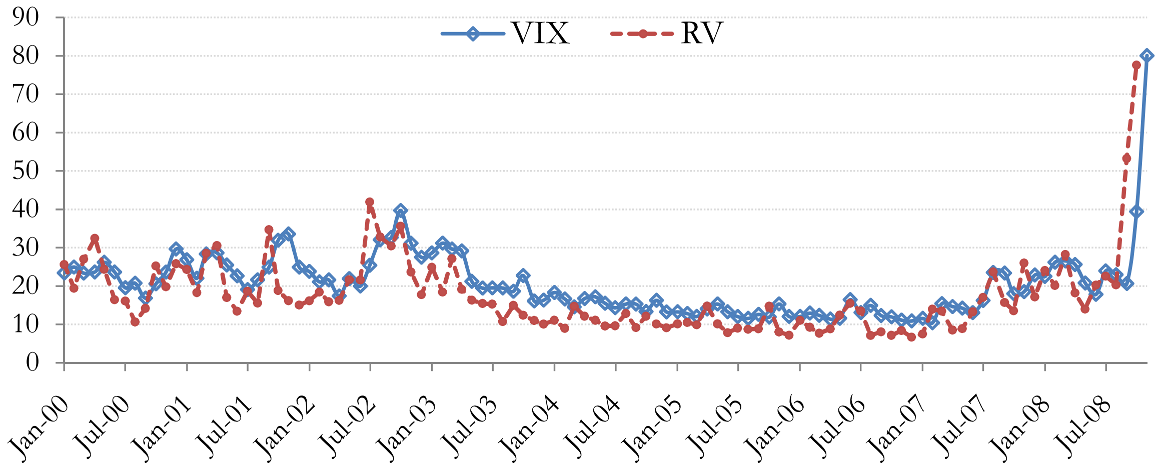

Figure 1: Unforecastable volatility innovations. We compare the monthly realized volatility, RV, of the S&P 500 index (calculated as the square root of the annualized sum of squared daily returns) with the option-implied risk-neutral expected S&P 500 volatility that month (captured by the VIX). The y-axis can be interpreted as annualized return standard deviation, ranging from less than 10% per year to close to 90% per year.

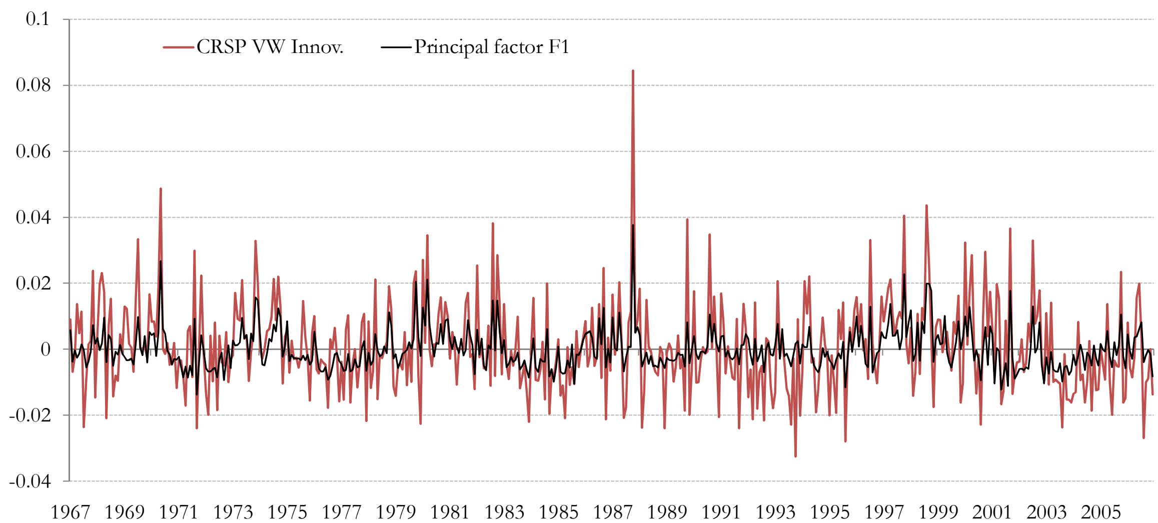

Figure 2: Time series plot of standardized principal component. The principal component extracted from the cross-section of individual stock log volatility innovations vs. the time series of standardized market log volatility innovations. The innovation series are based on ARMA(1,1) forecast errors of log realized volatility.

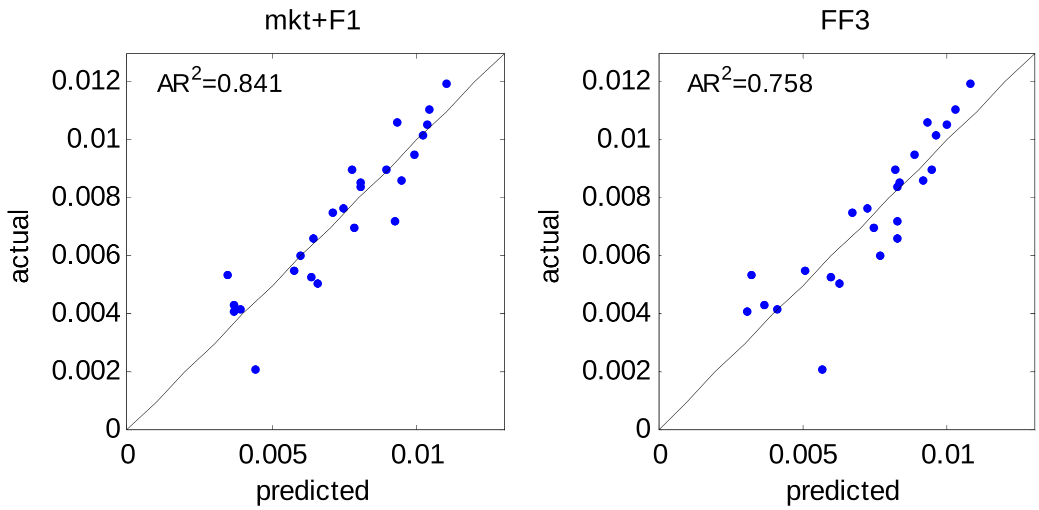

Figure 3: Realized and fitted excess returns of the 25 Fama-French size and book-to-market sorted portfolios. This figure shows average excess returns for the 25 size and book-to-market sorted portfolios against the fitted excess returns from the two-factor model and the Fama-French three-factor model.

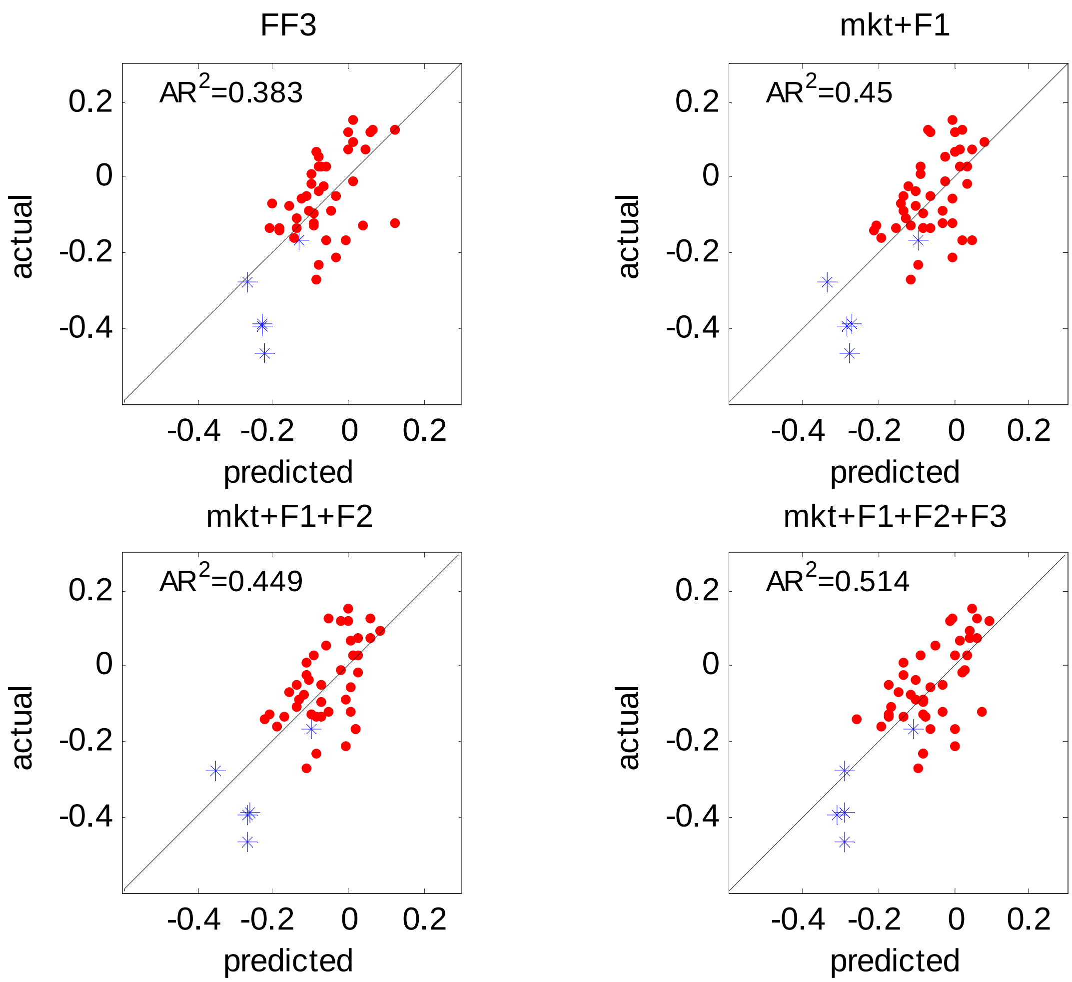

Figure 4: Average predicted and actual excess returns on variance swaps across stocks and indices. From 1997 to 2006, for each stock and index with more than 70 observations, we compute the time series averages of the actual excess returns on its variance swap. We then run a cross-sectional regression of these excess returns on their factor betas. We plot the factor-model-predicted excess returns against the actual excess returns on variance swaps separately for the Fama-French three-factor model ($FF3$) and two-factor ($MKT+F1$), three-factor ($MKT+F1+F2$) and four-factor ($MKT+F1+F2+F3$) models. The adjusted R-squares are also reported. Stock indices are represented using ($*$), and individual stocks are represented using ($\bullet$).

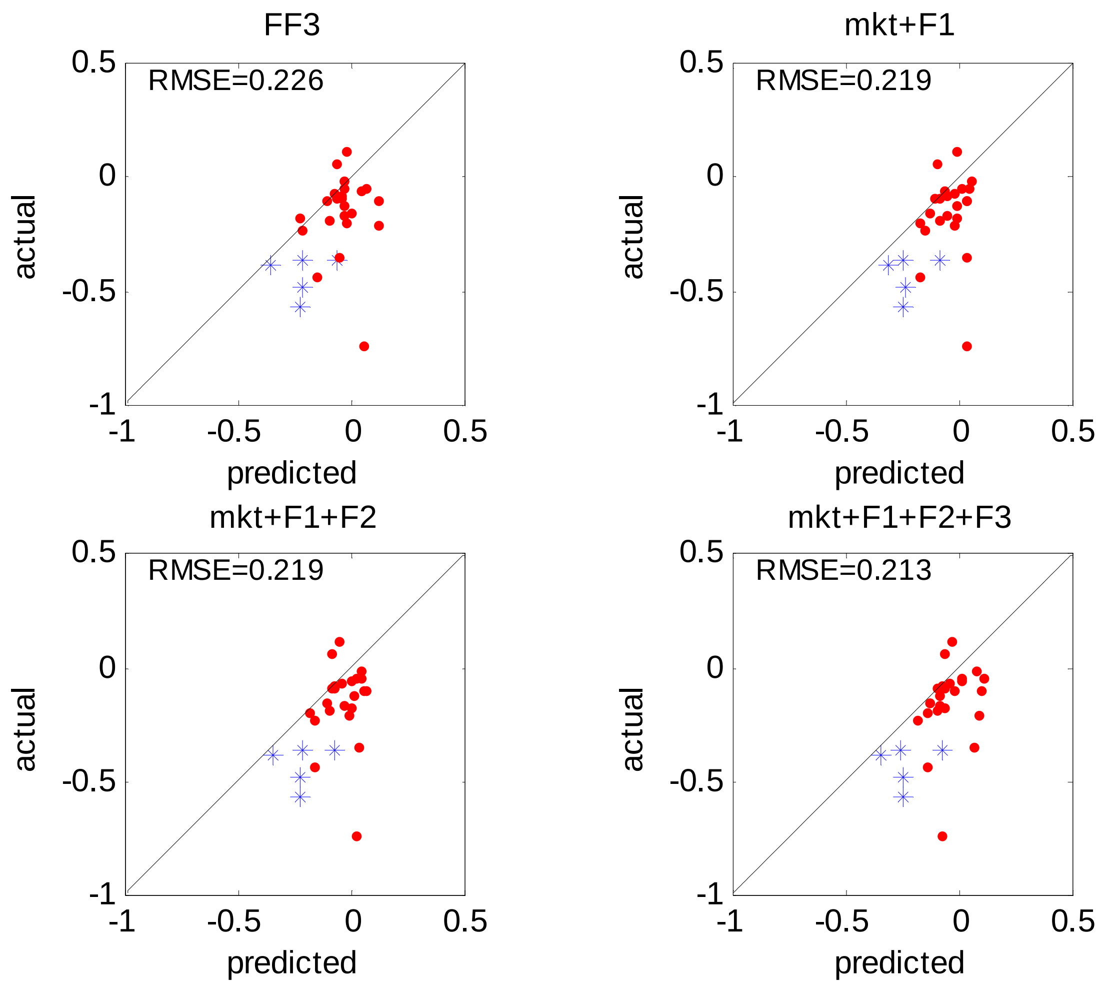

Figure 5: Average predicted and actual excess returns on variance swaps across stocks out-of-sample. We estimate factor loadings and risk premium during the in-sample period from 1997 through 2004 and then price the average variance swap returns during the out-of-sample period from 2005 to 2006. We keep stock and index with more than 50 observations during the in-sample period and more than 16 observations during the out-of-sample period. We use the corresponding average variance swap returns and factor loadings during the in-sample period to estimate factor risk premia. For each stock and index, we then compute a predicted excess return on the variance swap as an inner product between factor betas (computed in the rolling window ending in December 2004) and estimated factor risk premia suggested by the pricing equation (8). We then plot the predicted excess returns on variance swaps against the actual average excess returns during the out-of-sample period. We also report root mean square errors (RMSE). Stock indices are represented using ($*$), and individual stocks are represented using ($\bullet$).

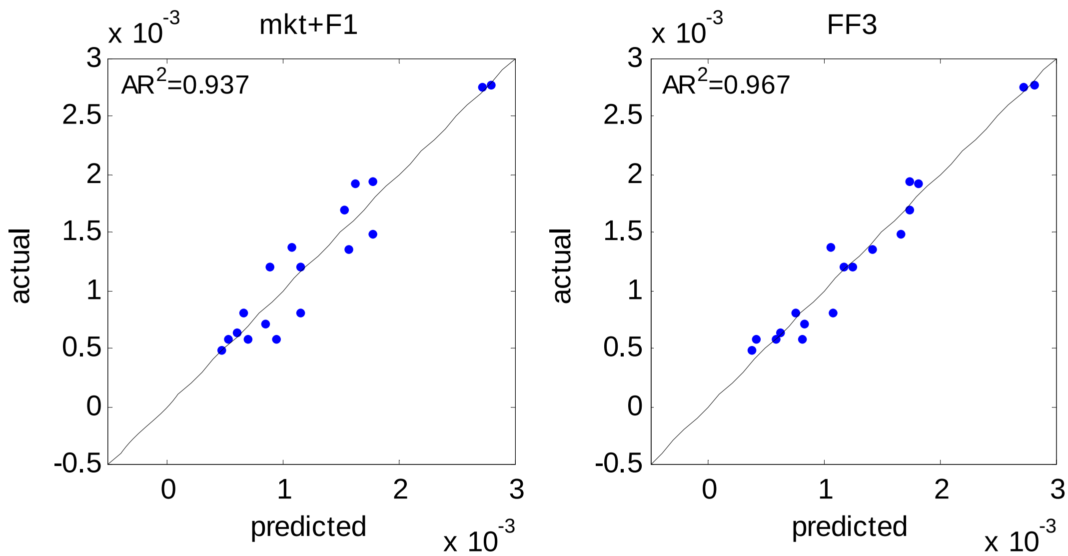

Figure 6: Realized and fitted excess returns of the 19 bond index portfolios. This figure shows average excess returns for the 19 bond index portfolios against the fitted excess returns from the two-factor model and the Fama-French three-factor model.

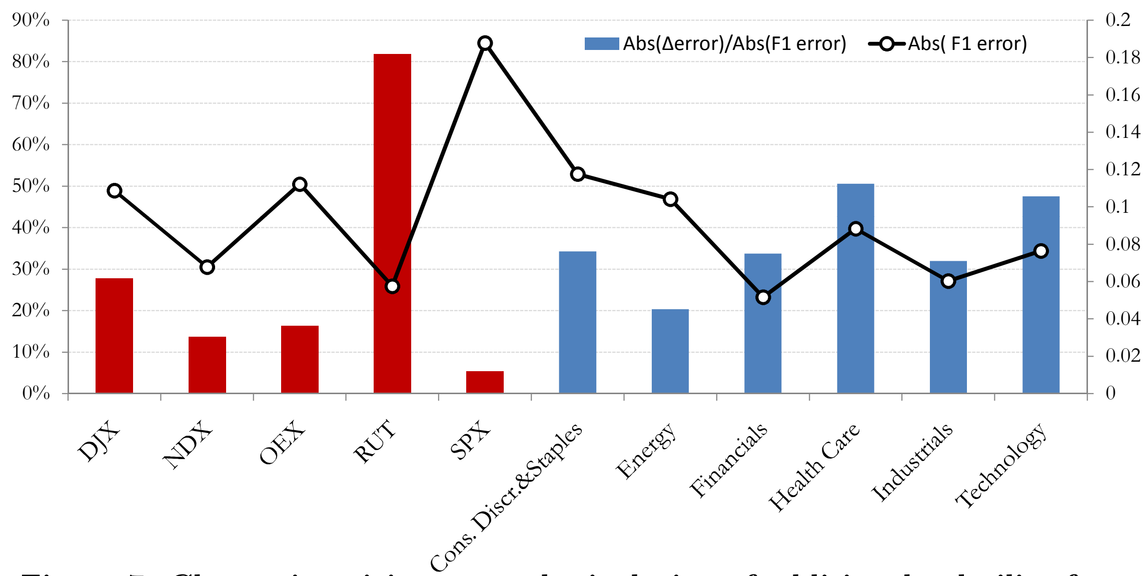

Figure 7: Change in pricing errors by inclusion of additional volatility factors. We group individual stock option by their two-digit Global Industry Classification System (GICS) classification and display the relative improvement in absolute pricing error from including two additional volatility factors (vertical bars corresponding to the left y-axis) along with the average absolute pricing error from the benchmark two-factor model (solid line corresponding to the right y-axis). An individual stock symbol is included if the synthetic swap contract can be constructed during at least 70 months. Index options are indicated by their tickers (first five bars).

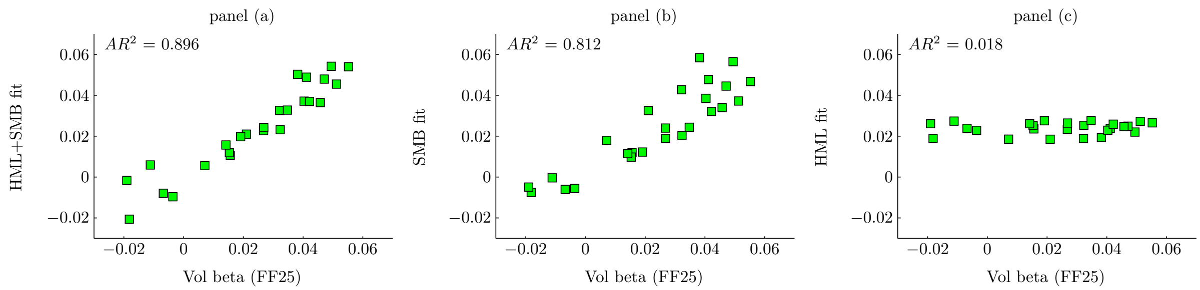

Figure 8: Spanning of the volatility beta by the Fama-French SMB and HML betas. We regress the volatility betas estimated from stocks (FF25 portfolios) on the corresponding Fama-French SMB and HML betas (including a constant). Panels (a)-(c) show the results based on the 25 book to market and size sorted portfolios. The volatility beta (x-axis) is plotted against the fitted value using the HML betas (panel c) or SMB betas (panel b) or both (panel a) on the y-axis.

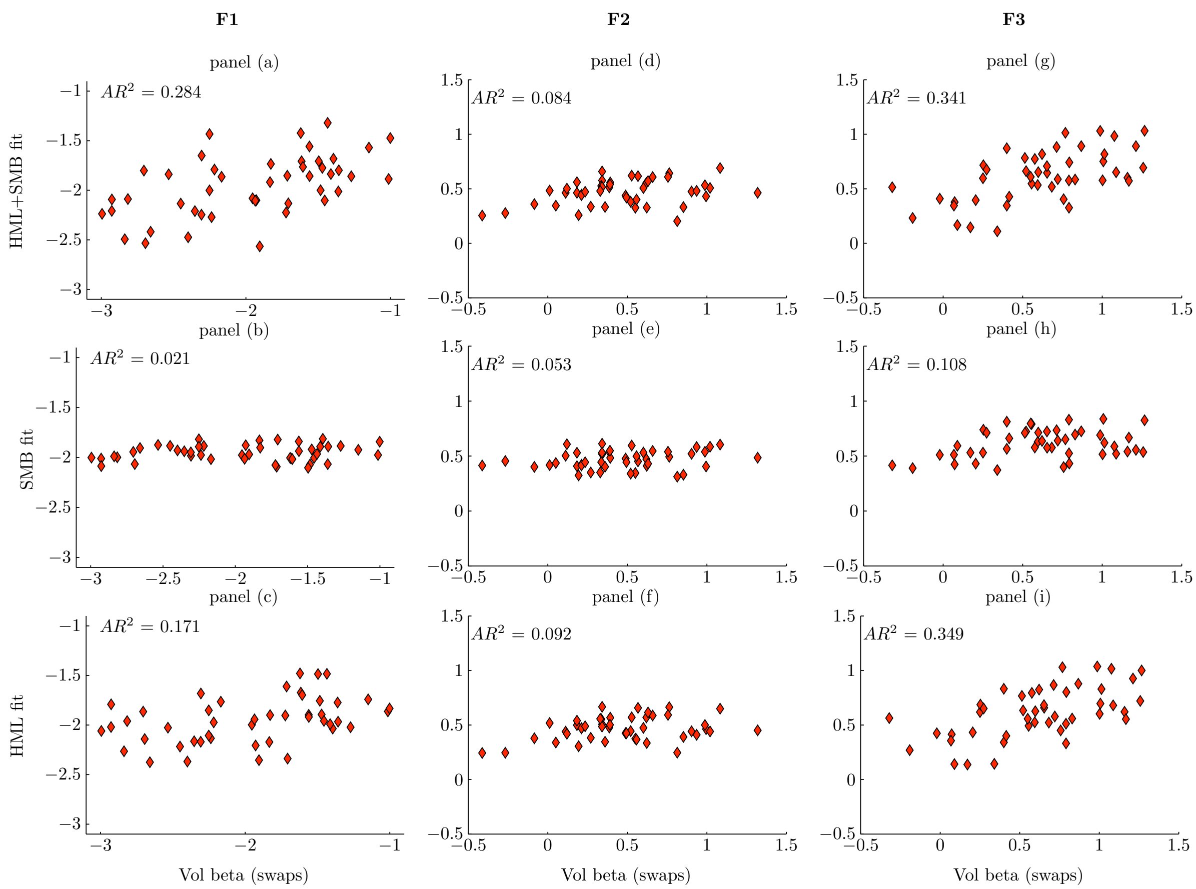

Figure 9: Spanning of the volatility beta by the Fama-French SMB and HML betas. For each of the three principal volatility factors (columns denoted F1, F2, F3), we regress the volatility betas estimated from options (synthetic variance swap contracts) on the corresponding Fama French SMB and HML betas (including a constant). Panels (a)-(c) show the results for the first volatility factor, F1, Panels (d)-(f) the results for the second volatility factor, F2, and Panels (g)-(i) the results for the third volatility factor, F3. In each case, the volatility beta (x-axis) is plotted against the fitted value using the HML betas (third row) or SMB betas (second row) or both (first row) on the y-axis.

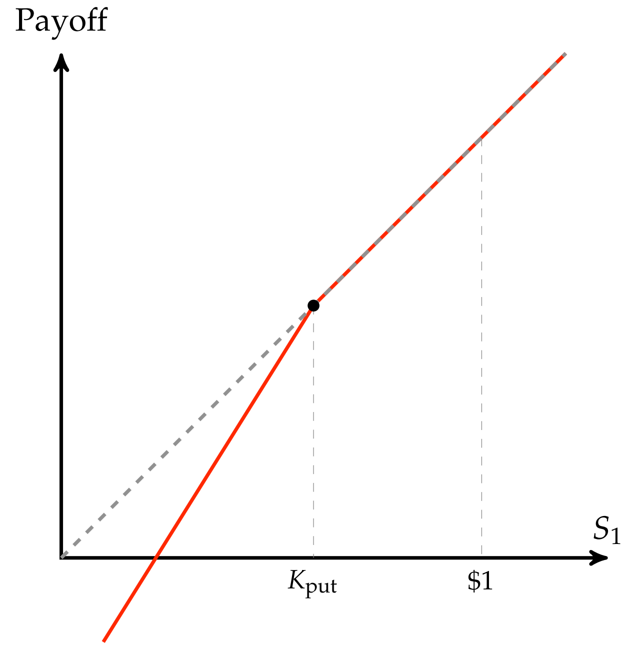

Figure 10: A Non-linear payoff. The one-period payoff to investing $S_0 = \$1$ in the stock index and shorting $\theta = \frac{1}{2}$ out-of-the-money puts (solid line) compared to the payoff to a $1 investment in the stock index (dashed line).

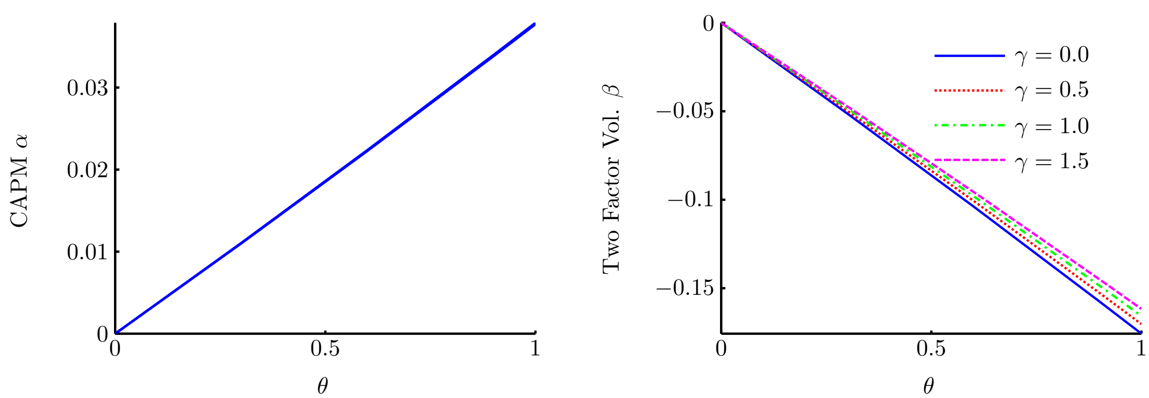

Figure 11: CAPM Alpha - Volatility Beta. The relation between CAPM α and the factor loading on realized volatility in an augmented CAPM. The stock price is assumed to follow the CEV model: $dS_t = \mu S dt + \sigma S^{\gamma/2} dB_t$, where $\mu = 8\%/year, \sigma = 20\%/year$. $\theta \in (0;1)$ is the number of puts shorted (as in Figure 10).