

When is the price of risk high?*

Alan Moreira and Tyler Muir

September 9, 2015

Abstract

We show that the price of risk and quantity of risk are negatively correlated in the time-series for benchmark factors in equities and currencies. Managed portfolios that increase factor exposures when volatility is low and decrease exposure when volatility is high thus produce positive alphas and increase factor Sharpe ratios. We also find volatility timing to be more beneficial to a mean variance investor than expected return timing by a fairly wide margin. These portfolio timing strategies are simple to implement in real time and are contrary to conventional wisdom because volatility tends to be high at the beginning of recessions and crises when selling is often viewed as a mistake. The facts pose a challenge to economic theory because they imply that effective risk aversion would have to be low in bad times when volatility is high, and vice versa. Hence, these facts pose a challenge for existing structural explanations of the excess volatility puzzle.

*Yale School of Management. We thank Jon Ingersoll for comments. We also thank Ken French and Adrien Verdelhan for making data publicly available.

1. Introduction

We provide estimates of the covariance between the price of risk and the quantity of risk for asset returns. We find this covariance to be strongly negative for benchmark pricing factors in equities and currencies over long samples spanning from 1926-present. Because risk, as measured by variance, is easily forecastable, this means that managed portfolios that increase risk exposure when variance is low, and decrease exposure when variance is high can increase Sharpe ratios and offer higher risk adjusted returns. These managed factors thus expand the mean variance frontier and therefore have positive alpha with respect to the original factors. We calculate the benefits of these strategies to a mean variance investor and find that forecasting the conditional variance is more beneficial than forecasting conditional means of factor returns. We also show that these facts are a challenge to theory because leading asset pricing models typically have the price of risk remaining roughly constant or have the price of risk increase when volatility is high due to increased effective risk aversion.

To fix ideas, consider an asset pricing factor with expected excess return $E_t[R_{t+1}]$ and variance σ_t^2 . The no-arbitrage restriction implies the following notion of price of risk, b_t ,

$$E_t[R_{t+1}] = b_t \sigma_t^2.$$

The time-series behavior of the price of risk summarizes how compensation for risk evolves over time. This is an interesting quantity in several di-

mensions. From a portfolio choice perspective, this is the quantity that a mean-variance investor would need to know to properly adjust his portfolio. From a general equilibrium perspective, b_t tracks the effective risk-aversion in the economy.

In this paper we show how to empirically estimate the covariance between b_t and σ_t^2 . We find it to be negative for a wide range of factors including the market, size, value, momentum, profitability, and investment factors in equities as well as the currency carry trade. We also consider linear combinations of these portfolios that are mean variance efficient. The estimate for the covariance between price and quantity and risk is negative and statistically significant in almost every case.

Consider the implications of this fact for an investor trying to time the market in order to increase his overall Sharpe ratio. Standard portfolio choice theory roughly tells the investor to set factor risk exposure proportional to b_t where the proportion will depend on risk aversion (a standard example would be for portfolio demand to equal $\frac{1}{\gamma}b_t$). If b_t is negatively related to σ_t^2 , then a strategy that increases risk exposure when volatility is low, and reduces risk exposure when volatility is high, will increase the overall portfolio Sharpe ratio.

We find that indeed this is the case. Managed portfolios that scale factors by the inverse of their past variance generally produce positive risk adjusted returns. We show that the scaled factors taken together can increase Sharpe ratios by large amounts, with annualized Sharpe ratios increasing by 0.6.

These strategies are easy to implement in real time because volatility is fairly easy to forecast over the short term, in contrast to forecasts of future expected returns. Volatility is also strongly counter-cyclical for all factors, which leads these portfolios to increase risk-adjusted returns while uniformly avoiding taking risk in recessions or major crises.

These facts are a challenge to economic theory because theory generally states that investors will, if anything, require higher compensation for taking risk during bad times, such as recessions, when volatility is high. Put differently, if effective risk aversion in the economy weakly increases in bad times when volatility is high, the correlation between the price and quantity of risk should be weakly positive. Indeed, structural asset pricing models (i.e., Campbell and Cochrane (1999), He and Krishnamurthy (2012), Bansal and Yaron (2004), Wachter (2013), Barberis et al. (2001)) generally either have b_t being a constant, so that expected returns vary through time purely as compensation for risk changing through time, or have b_t increasing with risk. This means the covariance between the price and quantity of risk should theoretically be either near zero or positive. We show that neither case is consistent with the data.

These facts may be surprising because there is a lot of evidence that expected returns are high in recessions and therefore recessions are viewed as attractive periods for taking risk (French et al. (1987)). In order to better understand the business cycle behavior of the risk-return trade-off, we combine information about time-variation in both expected returns and variance,

using predictive variables like the price to earnings ratio. We find that in times when volatility spikes, such as the Great Depression, Crash of 1987, or Great Recession, expected returns do not rise immediately by enough to make the risk return tradeoff attractive. However, since volatility movements are much less persistent than movements in expected returns, our optimal portfolio strategy prescribes a large reduction in risk exposure when volatility spikes, followed by a gradual increase in the exposure as the initial volatility shock fades. This allows an investor to avoid the unfavorable risk return trade-off during high volatility times.

These results square our findings with a longer literature that focuses on market timing based on return predictability rather than volatility (for example, Barberis (2000) and Campbell and Thompson (2008)). Our results make the point that time-varying volatility is a critical but often ignored determinant of market timing strategies and has the advantage that, over relatively short horizons, volatility is easier to forecast in real time than expected returns. Indeed, simple computations show that a mean-variance investor is significantly better off paying attention to expected volatility than to expected returns. Using standard predictive variables, an agent is able to increase his overall expected return on his portfolio through return forecasting by around 35% (Campbell and Thompson (2008)). We show using volatility timing that an agent can increase his expected return by around 85%. Thus, a mean-variance agent will find it more beneficial to forecast volatility than to forecast future returns. We show that the amount of return predictability

needed to overcome this result and make return forecasting more favorable than volatility forecasting requires a monthly *out of sample* forecasting R^2 of over 1%, generally exceeding the predictive power of standard variables in the literature.

Our results also builds on several other strands of literature. The first is the long literature on volatility forecasting (e.g. Andersen and Bollerslev (1998)). The consensus of this literature is that it is possible to accurately forecast volatility over relatively short horizons. We consider alternative models that vary in sophistication, but our main results hold for even a crude model that assumes next months volatility is equal to realized volatility in the current month. Our main results are enhanced by, but do not rely on, more sophisticated volatility forecasts which increase precision. This is important because it shows that a rather naive investor can implement these strategies in real time.

The second strand of literature debated whether or not the relationship between risk and return is positive (Glosten et al. (1993), Lundblad (2007), Lettau and Ludvigson (2003), among many others). Typically this is done by regressing future realized returns on estimated volatility or variance. The results of a risk return tradeoff are surprisingly mixed. The coefficient in these regressions is typically found to be negative or close to zero but is occasionally found to be positive depending on the sample period, specification, and horizon used. The question in this paper is different. In this paper we show that not only the sign, but the strength of this relationship

has qualitative implications for portfolio choice and structural asset pricing models. Even if this relationship is positive, we show that it can still imply a negative relationship between volatility and the price of risk.

The third strand of literature is the cross sectional relationship between risk and return. Recent studies have documented a low risk anomaly in the cross section where stocks with low betas or low idiosyncratic volatility have high risk adjusted returns (Ang et al. (2006), Frazzini and Pedersen (2014)). Our results complement these studies but are quite distinct from them. In particular, our results are about the time-series behavior of risk and return for a broad set of factors. We use the volatility of priced factors rather than idiosyncratic volatility of individual stocks and we show that our results hold for a general set of factors rather than using only CAPM betas. Consistent with this intuition, we show that controlling for a betting against beta factor (BAB) does not eliminate the risk adjusted returns we find in our volatility managed portfolios.

This paper proceeds as follows. Section 2 discusses the theoretical relation between risk and return for a general no-arbitrage factor model. It shows how to estimate the covariance between the quantity and price of risk empirically. Section 3 shows our main empirical results including our estimates of the covariance between quantity and price of risk as well as results related to our volatility managed portfolios. Section 4 discusses implications for structural asset pricing models. Section 5 discusses portfolio choice implications and derives dynamic market timing rules. Section 6 compares the welfare benefit

to a mean-variance investor of forecasting the variance versus the conditional mean of stock returns. Section 7 concludes.

2. Relation between price and quantity of risk

In the absence of arbitrage a pricing kernel, or stochastic discount factor (SDF), exists that prices all assets. Without loss of generality, we can assume a single factor model for the SDF, such that $SDF = a_t - b_t f_{t+1}$ where f is an excess return though we also consider a multi-factor model in the next section.¹² Then, it is well known that the pricing equation for this to be a valid SDF results in

$$E_t[R_{i,t+1}] = b_t cov_t(f_{t+1}, R_{i,t+1}) \quad (1)$$

where $R_{i,t+1}$ is any excess return.³

We then apply Equation 1 to f_{t+1} itself, which is also an excess return.

This results in

$$E_t[f_{t+1}] = b_t \sigma_{f,t}^2 \quad (2)$$

The above equation provides a natural decomposition of expected excess

¹This can be done, for example, by projecting the SDF onto the space of returns, see Cochrane (2009) Chapter 5.

²Special cases here would include a conditional CAPM (i.e., Jagannathan and Wang (1996)).

³For ease of notation, we work with the continuous time version of the above equation, which requires the approximation in discrete time that $(1 + r_{f,t}) \approx 1$ (specifically, the exact equation in discrete time is $E_t[R_{i,t+1}] = (1 + r_{f,t}) cov(-SDF_{t+1}, R_{i,t+1})$ though the equation above is commonly used as a close approximation). At shorter horizons, this approximation is relatively innocuous but we consider the results in the appendix from leaving the risk less rate in the equation. Generally speaking this common assumption simplifies the math without changing the main results.

returns for risk factors as determined by either movements in the price or quantity of risk.⁴ Note that the equation is distinct from studies that estimate conditional models and the price of risk in a purely cross-sectional manner (Jagannathan and Wang (1996)) because it uses only the *time-series* restrictions of the factors themselves rather than cross-sectional restrictions from a broader set of asset. Our goal is not to assess a model in terms of pricing errors for a cross-section of assets, but rather to explore the implied behavior of the price of risk b_t over time for a factor.

For ease of notation, we will denote $E_t[f_{t+1}] = \mu_t$ as the conditional expected factor return, $E[f_{t+1}] = \mu$ as the unconditional expected factor return, and $\sigma_{f,t}^2$ and $\sigma_{f,t}$ as the conditional variance and volatility of the factor. Our goal is to understand the relationship between movements in the price of risk, b_t , and the quantity of risk, $\sigma_{f,t}^2$, associated with these priced factors.

We have two results. The first result is that the coefficient in a predictive regression of future returns on the inverse of the factors' variance places sign restrictions on the covariance between the price and quantity of risk. Specifically, consider running the following regression:

$$f_{t+1} = a_0 + \beta \left(\frac{-1}{\sigma_{f,t}^2} \right) + \varepsilon_{t+1} \quad (3)$$

The following proposition shows that the coefficient, β in this regression is useful for estimating the covariance between price and quantity of risk.

⁴The decomposition is analogous to the Campbell and Shiller (1988) decomposition of price dividend ratios into either movements in cash flows or movements in risk premiums.

Proposition 1 (Covariance between price and quantity of risk). *The covariance between the price and quantity of risk is given by*

$$\text{cov}(b_t, \sigma_t^2) = \beta E[\sigma_t^2] \text{var}(1/\sigma_t^2) - \mu(E[\sigma_t^2]E[1/\sigma_t^2] - 1) \quad (4)$$

where β is defined as $\text{cov}\left(\frac{-1}{\sigma_{f,t}^2}, f_{t+1}\right) / \text{var}\left(\frac{1}{\sigma_{f,t}^2}\right)$ and can be estimated from the regression in (3).

Proof. See appendix. The intuition for the proof is straightforward. If changes in the quantity of risk do not strongly forecast returns, so that β is small, then the price of risk will have to negatively co-vary with the quantity of risk to justify this lack of predictability. \square

Examining this co-variance, it is worth pointing out that the second term is negative. This is because $E[\sigma_t^2]E[1/\sigma_t^2] > 1$ due to Jensen's inequality and $\mu > 0$ as we assume each factor has a positive premium (this is without loss of generality, since $-f$ could also serve as a factor). Then intuitively, the covariance can only be positive if β is large enough. This requires the conditional variance to have strong forecasting power for future returns. We return to estimating $\text{cov}(b_t, \sigma_t^2)$ in the next section.

An immediate corollary to Proposition 1 is that this covariance also puts a lower bound on how volatile the price of risk must be

Corollary 1.1 (Lower bound on volatility of the price of risk). *Given the estimate of the covariance between price and quantity of risk from Equation*

(4), we can place a lower bound on $\sigma(b_t)$

$$\sigma(b_t) \geq \frac{|\text{cov}(b_t, \sigma_t^2)|}{\sigma(\sigma_t^2)} \quad (5)$$

Proof. The covariance of the price of risk and quantity of risk can be expressed in terms of correlations as

$$\text{cov}(b_t, \sigma_t^2) = \text{corr}(b_t, \sigma_t^2)\sigma(b_t)\sigma(\sigma_t^2) \quad (6)$$

Taking absolute values of both sides, and using $|\text{corr}(b_t, \sigma_t^2)| \leq 1$ gives the result. \square

The Corollary states the obvious fact that a large non-zero covariance can come from two sources, namely the conditional variance can be very volatile or the price of risk can be very volatile. Because we can estimate the volatility of the conditional variance, we can determine how volatile the price of risk must be. Placing a lower bound on the volatility of the price of risk is useful in understanding what a model needs to match the data, much as the bounds in Hansen and Jagannathan (1991) place restrictions on the volatility of the SDF needed to match the data. It also helps us understand where the large negative covariance we see in the data must come from.

Next we show that the sign of the covariation between risk and price of risk gives us guidance on how to construct managed portfolios that generate risk-adjusted returns relative to original factors. Intuitively, if the price of risk is low during periods of high volatility, which is case if the covariance in Proposition 1 is negative, then a portfolio that has lower weights during

periods of higher volatility should earn higher alphas. Formally this can be seen in the following result

Corollary 1.2 (Managed portfolio alpha). *Let α be the intercept of the following regression*

$$\frac{1}{\sigma_{f,t}^2} f_{t+1} = \alpha + \beta f_{t+1} + \varepsilon_{t+1}, \quad (7)$$

then α is decreasing in the covariance between the price and quantity of risk $cov(b_t, \sigma_t^2)$.

Proof. Let $b = E[\mu_t]var(f_t)^{-1}$ be the unconditional price of risk associated with the factor, then the managed portfolio alpha can be written as

$$\alpha = E \left[cov_t \left(\frac{1}{\sigma_t^2} f_{t+1}, f_{t+1} \right) b_t \right] - cov \left(\frac{1}{\sigma_t^2} f_{t+1}, f_{t+1} \right) b \quad (8)$$

Using that $f_{t+1} = b_t \sigma_t^2 + \varepsilon_{t+1}$ and law of iterated expectations we obtain $\alpha = E[b_t] - b - cov(b_t, b_t \sigma_t^2) b$. We now use the following approximation,

$$b_t \sigma_t^2 \approx (b_t - E[b_t])E[\sigma_t^2] + (\sigma_t^2 - E[\sigma_t^2])E[b_t], \quad (9)$$

which after simple algebra yields the following equation for the managed portfolio alpha

$$\alpha = -cov(b_t, \sigma_t^2)E[b_t]b - (b - E[b_t] + bE[\sigma_t^2]var(b_t)). \quad (10)$$

As long $E[b_t]b > 0$, this implies that a more negative covariation between risk and price of risk translates into a larger managed portfolio alpha. \square

2.1 Generalizing the result to a multi-factor model

While we can always write a single factor model for the SDF, in practice we typically work with multi-factor models. We thus generalize to the case when there are K factors, so that $F' = [f_1, \dots, f_K]$. Then the SDF takes the form $SDF_{t+1} = a_t - B'_t F_{t+1}$ where $B'_t = [b_1, \dots, b_K]$ is a vector of weights. Without any loss of generality we consider factor models that with orthogonal risk factors. Empirically, this is often how factors are constructed (see, e.g., Fama and French (1996)). Specifically, when forming a “size” factor it is common to double sort on size and book to market to orthogonalize the factors. The empirical motivation is to separate the premium related to size to that related to, say, book to market and the procedure roughly approximates principle components that capture orthogonal risk factors in an APT sense. Since size and book to market both use market value of equity they are naturally correlated and this orthogonalization procedure helps to isolate the priced component of returns due to each factor distinctly. Therefore, the assumption of orthogonal factors is an approximation of how multi-factor models are formed empirically. Alternatively, one can always rotate the factor space so that factors are orthogonal.⁵

With this setup our main pricing equation for factor f_i is

$$E_t[f_{i,t+1}] = -cov_t(f_{i,t+1}, SDF_{t+1}) = -cov_t(f_{i,t+1}, a_t - B'_t F_{t+1}) \quad (11)$$

⁵e.g. Kozak et al. (2015)

So that

$$E_t[f_{i,t+1}] = B_t' cov_t(f_{i,t+1}, F_{t+1}) \quad (12)$$

Using the assumption that the factors are uncorrelated, this becomes

$$E_t[f_{i,t+1}] = b_{i,t} \sigma_{f,i,t}^2 \quad (13)$$

Therefore the orthogonality assumption generates the same equation, that a given factors conditional expected return is equal to its conditional quantity of risk times its conditional price of risk. The equation suggests that we can compute these relationships for various factors individually if the orthogonality assumption holds.

3. Empirical Results

3.1 Data Description and Volatility Forecast

We use both daily and monthly factors from Ken French on Mkt, SMB, HML, Mom, RMW, and CMA. The first three factors are the original Fama-French 3 factors (Fama and French (1996)), while the last two are a profitability and investment factor that they use in their 5 factor model (Novy-Marx (2013), Hou et al. (2014)). Mom represents the momentum factor which goes long past winners and short past losers. We also use data on currency returns from Adrien Verdelhan used in Lustig et al. (2011). We use the monthly high minus low carry factor formed on the interest rate differential, or forward discount, of various currencies. We have monthly data on returns and use daily data on exchange rate changes for the high and low portfolios

to construct our volatility measure.⁶ We refer to this factor as “Carry” or “FX” to save on notation to emphasize that it is a carry factor formed in foreign exchange markets. Finally, we form two mean-variance efficient equity portfolios which are the ex-post mean variance efficient combination of the equity factors using constant unconditional weights. The first uses the Fama-French 3 factors along with the momentum factor and begins in 1926, which the second adds RMW and CMA but because of things begins only in 1963 (we label these portfolios MVE and MVE2, respectively). The idea is that these portfolios summarize all the asset pricing implications of the individual factors. It is thus a natural benchmark to consider.

We compute realized variance (RV) for a given month t for a given factor f by taking the variance of the past daily returns in the month. This information is known at the end of month t and we use this as conditioning information in predicting returns and forming portfolios for the next month $t + 1$. Our approach is simple and uses only return data.

Next we form forecasts for volatilities, variances, and inverse variances by projecting realized log variance onto past information. We run this in logs for several reasons. First, volatility is approximately log normal (Andersen et al., 2001). Second, logs de-emphasize extreme realizations of variances or volatility. Third, the log forecast allows us to easily convert between volatility, variance, and inverse variance in a parsimonious way thus requiring

⁶We thank Adrien Verdelhan for help with the currency daily data.

only one model to form all of our forecasts. Specifically, we run

$$\ln RV_{t+1} = a + \sum_{j=1}^J b_j \ln RV_{t-j-1} + cx_t + \varepsilon_{t+1} \quad (14)$$

Where x_t are controls and RV_t represents realized variance in a given month. For each factor we include three monthly lags of its own realized variance so $J = 3$, and we include the log market variance as a control (note: this would be redundant for the market itself). We plan to explore additional controls in the future, though we note the regressions generally produce a reasonably high R-square. We can then form conditional expectations of future volatility or variance or inverse variance easily from our log specification. Taking conditional expectations forms our forecasts; specifically

$$\sigma_t^n = E_t[RV_{t+1}^n] = \exp(n(a + b \ln \sigma_t + cx_t)) \exp\left(\frac{n^2}{2} \sigma_\varepsilon^2\right)$$

We can then set $n = 1/2, 1, -1$ as our forecast of volatility, variance, or inverse variance. The last term in the above equation takes into account a Jensen's inequality effect.

Table 9 in the Appendix gives the results of these regressions. We only report the coefficient on the first lag. The Table confirms the well known result that volatility is highly forecastable at the monthly horizon and has a monthly persistent ranging from 0.5-0.8 across factors. This finding is important for our results because it means expected volatility next months is fairly easily predicted using volatility from recent months.

We also plot realized volatility for all factors used in Figure 1. Two things are worth noting. The first is that the volatility of all factors is strongly

counter-cyclical, rising in the beginning of recessions. The second is that factor volatilities co-move strongly, as they tend to rise together in persistent and predictable ways.

3.2 The Covariance Between Price and Quantity of Risk

We begin by running regressions of future 1 month returns on each of the individual factors on their past month's negative inverse variance $-1/\sigma_t^2$. This specification is theoretically motivated by Equation (4), as the estimated slope has direct implications for the correlation between the price of risk and variance. Table 1 gives these regressions. We can see immediately that, consistent with the results in the risk-return tradeoff literature, the coefficients are generally close to zero and statistically insignificant. A notable exception is SMB which is positive and significant. The fact that these estimates are noisy means volatility doesn't strongly predict future returns, consistent with a weak risk-return tradeoff.

We then use the β from the above regression to estimate Equation (4) directly, thus giving us a direct estimate of $cov(b_t, \sigma_t^2)$. We compute standard errors and confidence intervals for $cov(b_t, \sigma_t^2)$ using bootstrap. The results, presented in Table 2, show that the point estimate for the covariance is negative for all factors besides SMB. Again, this is to a large extent the result of the β in the forecasting regression being either negative or small in magnitude. In most cases the 95% confidence intervals do not include 0, with

the exception of SMB and CMA. Therefore, for most factors the data easily rejects a value of the covariance above 0, and imply negative point estimate in almost all cases.

Next, we use these covariances to place restrictions on the volatility of the price of risk, as shown in Corollary 1.1. Because we can estimate the volatility of the conditional variance, the covariance tells us a lower bound for the volatility of the price of risk. We give these results in Table 3. We find that the *monthly* volatility of the price of risk is economically quite large and comparable to the unconditional price of risk (reported in the last column).

3.3 Managed Volatility Portfolios

We construct managed portfolios by scaling each factor by the inverse of its variance. That is, each month we increase or decrease our risk exposure to the factors by looking at the realized variance over the past month. The managed portfolio is then $\frac{c}{RV_t^2} f_{t+1}$ for a constant c which we choose so that the managed factor has the same unconditional standard deviation as the non-managed factor. The idea is that if variance does not forecast returns, the risk-return trade-off deteriorate when variance increases. In fact, this is exactly what a mean-variance optimizing agent should do if she believes volatility doesn't forecast returns. In our main results, we keep the managed portfolios very simple by only scaling by past realized variance instead of the optimal expected variance computed using a forecasting equation. The reason is that this specification does not depend on the forecasting model

used and could be easily done by an investor in real time.

Table 4 reports the regression of running the managed portfolios on the original factors. We can see positive, statistically significant constants (α 's) in most cases. Consistent with Corollary 1.2 the two factors (SMB and CMA) that have stronger co-movement between volatility and future returns (see Table 1), are the ones with small managed portfolio alphas. Intuitively, alphas are positive because the managed portfolio takes advantage of the larger price of risk during low risk times and avoids the poor risk-return trade-off during high risk times. The managed market portfolio on its own likely deserves special attention because this strategy would have been easily available to the average investor in real time and it relates to a long literature in market timing that we refer to later.⁷ The scaled market factor has an annualized alpha of around 5% and a beta of only 0.6. While most alphas are strongly positive, the largest is momentum. This is consistent with Barroso and Santa-Clara (2015) who find that strategies which avoid large momentum crashes perform exceptionally well.

In all tables reporting α 's we also include the root mean squared error, which allows us to construct the managed factor excess Sharpe ratio, thus giving us a measure of how much dynamic trading expands the slope of the MVE frontier spanned by the original factors. More specifically, the Sharpe ratio will increase by precisely $\sqrt{SR_{old}^2 + \left(\frac{\alpha}{\sigma_\varepsilon}\right)^2} - SR_{old}$ where SR_{old} is the

⁷The average investor will likely have trouble trading the momentum factor, for example.

maximum Sharpe ratio given by the original non-scaled factors. For example, in Table 4, scaled momentum has an α of 12.5 and a root mean square error around 50 which means its annualized appraisal ratio is $\sqrt{12} \frac{12.5}{50} = 0.875$. The scaled markets annualized appraisal ratio is 0.34.⁸

Table 5 shows the results when, instead of scaling by past realized variance, we scale by the expected variance from our forecasting regressions. This offers more precision but comes at the cost of assuming an investor could forecast volatility using the forecasting relationship in real time. As expected, the increased precision generally increases significance of alphas and increases appraisal ratios. We favor using the realized variance approach because it does not require a first stage estimation and it also has a clear appeal from a more practical implementation perspective.

Tables 6, 7, and 8 add additional controls to the regressions and compute alphas relative to larger multi-factor models. Notably, Table retable:alpha2 shows that our results are not explained by the betting against beta factor (BAB). Thus our *time-series* volatility managed portfolios are distinct from the low beta anomaly documented in the cross section. Tables 7, and 8 show that the scaled factors expand the mean variance frontier of the existing factors because the appraisal ratio of HML, RMW, Mom are strongly positive when including all factors. The MVE portfolio's appraisal ratio here is 0.62 which is economically very large. We conclude that managing factors using

⁸ We need to multiply the monthly appraisal ratio by $\sqrt{12}$ to arrive at annual numbers. We multiplied all factor returns by 12 to annualize them but that also multiplies volatilities by 12, so the Sharpe ratio will still be a monthly number.

their past variance provides a powerful way to expand the mean variance frontier.

4. Economic Theory

This section sketches the implications of our empirical results for theories of time-varying risk-premia. We have in mind theories based on habit formation (Campbell and Cochrane, 1999), long-run risk (Bansal and Yaron, 2004), intermediaries (He and Krishnamurthy, 2012), and rare disasters (Wachter, 2013).

First, we want to acknowledge that these models were not designed to think about the cross-section of asset prices, and these theoretical explanations have not had much empirical success in accounting for several of the cross-sectional factors that we study. For example, none of these theories are successful in accounting for the momentum factor in the cross-sectional of stock returns.

However, we feel our point is broader. These theories rely on a combination of time-varying risk and time-varying risk-premia to match the amount of excess volatility and time-variation documented in the U.S. data. Our empirical work allow us to study whether the joint dynamics these models rely on is consistent with the data.

In He and Krishnamurthy (2012) (HK) and Campbell and Cochrane (1999) (CC) time-variation in risk and the price of risk is endogenous and a function of past shocks. In these models,

$$\max_i \frac{E_t[R_{i,t+1}^e]}{\sigma_t^2(R_{i,t+1}^e)} \approx f(s_t) \quad (15)$$

where $f'(s_t) < 0$, and s_t is a state variable that is an increasing function of past shocks to consumption. In both these models past consumption shocks shape the sensitivity of the marginal utility to future consumption shocks. In CC positive shocks to consumption increase the distance of the agent to its habit level, reducing the effective risk-aversion. In HK it increases the share of wealth held by financial intermediaries, also having the effect of reducing the effective risk-aversion in the economy. So in both these economies, negative shocks to consumption makes marginal utility more volatile, resulting in an endogenous increase in asset price volatility. So any asset with a positive risk-premia will also feature an endogenous increase in its return volatility in periods where the price of risk is high. Another notable paper with this feature is Barberis et al. (2001) who have effective risk aversion occurring after realized losses.

This property is a fundamental piece of how these models explain the excessive volatile puzzle: returns are more volatile than fundamentals because of fluctuations in discount rates are larger during bad times. Our empirical work shows that the data does not provide strong support for this mechanism. While it is undisputed that the data features a lot of excess volatility, these models generate volatility at the wrong times. In the data periods of higher than average volatility are associated with periods of lower than average price of risk. Including an additional state variable in these models might

be useful way to break this tight link between volatility and risk-premia.

Frameworks that rely on time-varying risk such as Wachter (2013) and Bansal and Yaron (2004) have less stark predictions for the joint dynamics of risk and the price of risk.

In Wachter (2013), it is the probability of a rare disaster, which simultaneously drives stock market variance and risk-premia. In her calibration, the co-variance between price of risk and variance is positive, so it is also qualitatively inconsistent with what we measure in the data. While we cannot rule out that there are parameter combinations able to deliver a negative relation, these combinations must feature less disaster risk, making it harder for the model to fit other features of the data.

In Bansal and Yaron (2004), it is time-varying fundamental volatility which drives variation in risk-premia and stock market risk. It is the only framework that produces a negative relationship between risk and the price of risk. Intuitively, because shocks to volatility are priced and do not scale up with volatility (volatility itself has constant volatility), increases in volatility have the effect of reducing the price of risk. In practice for the leading calibrations (Bansal and Yaron (2004), Bansal et al. (2009)) the relationship is flat. The model can produce a more pronounced negative relation between risk and the price of risk by substantially increasing the conditional volatility of the consumption volatility process, but this would be strongly counterfactual given the empirical dynamics of aggregate consumption.

5. Portfolio Choice Implications

The portfolio choice literature has focused on studying how portfolios that explores the predictability of returns in the data can increase portfolios expected returns and Sharpe ratios. This literature has concluded that the uncertainty around the return forecasting regressions estimated in the data makes market timing strategies harder to implement in practice (Barberis (2000)). Our paper emphasizes that return volatility is a order of magnitude more predictable than expected returns and the relationship between factor risk price and factor risk we document in the data is also stable to analysis across different sub-samples, contrasting with the instability of return forecasting regressions. The evidence in this paper implies that portfolio advice that exploits variance predictability (Chacko and Viceira (2005)) is empirically relevant across a wide range of risk factors. The data strongly favors strategies that take less risk in periods of high risk.

We estimate a VAR for expected returns and variances of the market portfolio. We then trace out the portfolio choice implications for a myopic mean variance investor. We set risk aversion just above 2, where our choice is set so that the investor holds the market on average when using unconditional value for variance and the equity premium (i.e., in the absence of movements in expected returns and vol, $\alpha = 1$). This gives a natural benchmark to compare to.

The VAR first estimates the conditional mean and conditional variance of

the market return using monthly data on realized variance, monthly market returns, the monthly (log) price to earnings ratio, and the BaaAaa default spread. The expected return is formed by using the fitted value from a regression of next months stock returns on the price to earnings ratio, default spread, and realized variance (adding additional lags of each does not change results). Expected variance is formed as described in Section 3. We then take the estimated conditional expected return and variance and run a VAR with 3 lags of each variable. We consider the effect of a variance shock where we choose the ordering of the variables so that the variance shock can affect contemporaneous expected returns as well. These results are meant to be somewhat stylized in order to understand our claims about the price of risk when expected returns also vary and to understand the intuition for how portfolio choice should optimally respond to a high variance shock.

The results are given in Figure 3. We see that a variance shock raises future variance sharply and immediately. Expected returns, however, do not move much on impact but rise slowly as time goes on. The impulse response for the variance dies out fairly quickly, consistent with variance being strongly mean reverting. Given the increase in variance but only slow increase in expected return, the lower panel shows that it is optimal for the investor to reduce her portfolio exposure from 1 to 0.6 on impact because of an unfavorable risk return tradeoff. This is because expected returns have not risen fast enough relative to volatility. The portfolio share is consistently below 1 for roughly 18 months after the shock. At this point, variance de-

creases enough and expected returns rise enough that an allocation above 1 is desirable. This increase in risk exposure fades very slowly over the next several years. These results square our findings with the portfolio choice literature. They say that in the face of volatility spikes expected returns do not react immediately and at the same frequency. This suggests reducing risk exposures by substantial amounts at first. However, the investor should then take advantage of favorable increases in expected returns once volatility has return to reasonable levels.

It is well known that both movements in stock-market variance and expected returns are counter-cyclical (French et al. (1987), Lustig and Verdelhan (2012)). Here, we show that the much lower persistence of volatility shocks implies the risk-return trade-off initially deteriorates but gradually improves as volatility recedes through the recession.

6. The Net Benefits of Volatility Timing

We now use simple mean-variance preferences to give a rough sense of the benefits of timing fluctuations in conditional factor volatility and compare this to forecasting conditional mean returns. For simplicity we focus on a single factor. For our numerical comparisons we will use the market as the factor though we emphasize that our volatility timing works well for other factors as well. In particular our goal here is to compare directly the benefits of volatility timing and expected return timing. Specifically, we extend the analysis in Campbell and Thompson (2008) to allow for time-variation in

volatility.

Consider the following process for excess return on the factor,

$$r_{t+1} = \mu + x_t + \sqrt{Z_t}e_{t+1} \quad (16)$$

where $E[x_t] = E[e_t] = 0$ and x_t, e_{t+1} are conditionally independent. For a mean-variance investor his portfolio choice and welfare can be written as follows,

$$w = \frac{1}{\gamma} \frac{\mu}{E[Z_t] + \sigma_x^2}; \quad (17)$$

$$w(x_t) = \frac{1}{\gamma} \frac{\mu + x_t}{E[Z_t]}; \quad (18)$$

where the different lines reflect alternative conditioning sets. The first row uses no conditional information, while the second row uses information about the conditional mean but ignores fluctuations in volatility. We now compare expected returns across the two conditioning sets multiplying factor returns by the portfolio weight before taking unconditional expectations. That is, we compute $E[w_t r_{t+1}]$ in each case. For the fixed weight and the expected return strategy this implies,

$$E[wr_{t+1}] = \frac{1}{\gamma} \frac{\mu^2}{\sigma_x^2 + E[Z_t]} = \frac{1}{\gamma} S^2 \quad (19)$$

$$E[w(x_t)r_{t+1}] = \frac{1}{\gamma} \frac{\mu + \sigma_x^2}{E[Z_t]} = \frac{1}{\gamma} \frac{(S^2 + R_r^2)}{1 - R_r^2}, \quad (20)$$

where S is the unconditional Sharpe ratio of the factor $S = \frac{\mu}{\sqrt{E[Z_t] + \sigma_x^2}}$ and R_r^2 is the share of the return variation captured by the forecasting signal x . The proportional increase in expected returns (and utility) is $\frac{1+S^2}{S^2} \frac{R_r^2}{1-R_r^2}$.

This essentially assumes that there is no risk-return trade-off in the time-series. With such an assumption Campbell and Thompson (2008) show that a mean-variance investors can experience a proportional increase in expected returns and utility of roughly 35% by using conditional variables know to predict returns such as the price-earnings ratio. For example if an investor had risk-aversion that implied an expected excess return on his portfolio of 5%, the dynamic strategy implies average excess returns of 6.75%.

To evaluate the value added by volatility timing we extend these computations by first adding only a volatility signal. This exercise also assumes away any risk-return co-movement that might exist in the data. We next study the general case which allows for both signals to operate simultaneously and possibly against each other.

For the simple case we assume the variance process is log-normal $Z_t = e^{z_t}$, with

$$z_t = \bar{z} + y_t + \sigma_z \times u_t. \quad (21)$$

This produces optimal portfolio weights and expected returns given by

$$w(z_t) = \frac{1}{\gamma} \frac{\mu}{e^{2(\bar{z} + y_t + \sigma_z^2)}}; \quad (22)$$

$$E[w(z_t)r_{t+1}] = \frac{1}{\gamma} S^2 e^{\sigma_y^2}, \quad (23)$$

where this last equation can be written as function of variance forecasting R-squares, $E[w(z_t)r_{t+1}] = \frac{1}{\gamma} S^2 e^{R_z^2 \times Var(z_t)}$. The proportional expected return again of such strategy is simply $e^{R_z^2 \times Var(z_t)}$. The total monthly variance of

log realized variance is 1.06 in the full sample (1926-2015), and a very naive model that simply uses lagged variance as the forecast of future variance achieves a $R_z^2 = 0.38$, implying a proportional expected return increase of 50%. A slightly less naive model that takes into account mean-reversion and uses the lagged realized variance to form a OLS forecast (i.e., an AR(1) model for variance) achieves $R_z^2 = 53\%$. This amount of predictability implies a proportional increase in expected returns of 75%. A sophisticated model that uses additional lags of realized variance can reach even higher values. Using more recent option market data one can construct forecasts that reach as much as 60% R-square, implying a expected return increase of 90%. These estimates do not depend on whether the R-square is measured in or out of sample as this relationship is incredible stable over time. These effects are economically large and rely on taking more risk when volatility is low, periods when leverage constraints are less likely to bind.

It is worth noting that all of these methods provide larger increases in expected returns than forecasts based on the conditional mean. A simple calculation shows that the forecasting power for the market portfolio would need to have an out of sample R-square above 1% per month to outperform our volatility timing method, which is quite high given the literature on return predictability. Moreover, even if some variables are found that are able to predict conditional expected returns above this threshold, it is not clear if an investor would have knowledge and access to these variables in real time. In contrast, data on volatility is much simpler and more available.

Even a naive investor who simply assumes volatility next month is equal to realized volatility last month will outperform the expected return forecast.

Importantly, these calculations can be done not just for the market return, but for *any* factor, with similar degrees of success. The degree of predictability we find for the conditional variance of different factors is fairly similar, with simple AR(1) models generally producing R-square values around 50-60% at the monthly horizon. In contrast, the same variables that help forecast mean returns on the market portfolio do not necessarily apply to other portfolios such as value, momentum, or the currency carry trade. Thus, while we would need to come up with additional return forecasting variables for each of the different factors, using lags of the factors own variance is a very reliable and stable way to estimate conditional expected variance across factors.

Nevertheless one needs to be cautious with these magnitudes. The same caveat that applied to CT applies here as well. Specifically, it is possible that co-movement between x_t and z_t erases much of these gains.

We now include both sources of time-variation in the investment opportunity set and allow for arbitrary co-variation between these investment signals. In this case we have,

$$\begin{aligned} w(Z_t, x_t) &= \frac{1}{\gamma} \frac{\mu + x_t}{Z_t}; \\ E[w(Z_t, x_t)r_{t+1}] &= \frac{1}{\gamma} \left(\frac{S^2 + R_r^2}{1 - R_r^2} E[Z_t]E[Z_t^{-1}] + 2\mu cov(x_t, Z_t^{-1}) + cov(x_t^2, Z_t^{-1}) \right), \end{aligned} \tag{24}$$

In the first term we have the total effect if both signals were completely

unrelated – that is, if there was no risk-return trade-off at all in the data. Under this assumption and using the $R_x^2 = 0.43\%$ from the CT study for the expected return signal and the more conservative $R_z^2 = 0.53\%$ for the variance signal, one would obtain a 236% increase in expected returns. But there is some risk-return trade-off in the data, so we need to consider the other terms as well.

The second term we can construct directly from our estimates in Table 1, using that $cov(x_t, Z_t^{-1}) = -\beta Var(Z_t^{-1})$. The third term is trickier but likely very small. One possibility is to explicitly construct expected return forecasts, square them, and compute the co-variance with realized variance. Here we take a more conservative approach and only characterize a lower bound

$$cov(x_t^2, Z_t^{-1}) \geq -1\sigma(x_t^2)\sigma(Z_t^{-1}) = -1\sqrt{2}\sigma_x^2\sigma(Z_t^{-1}), \quad (25)$$

where we assume that x is normally distributed. Substituting back in equation we obtain,

$$E[w(Z_t, x_t)r_{t+1}] \geq \frac{1}{\gamma} \left(\frac{S^2 + R_r^2}{1 - R_r^2} E[Z_t]E[Z_t^{-1}] - 2\beta\mu\sigma^2(Z_t^{-1}) - \sqrt{2}\sigma_x^2\sigma(Z_t^{-1}) \right). \quad (26)$$

Plugging numbers for σ_x consistent with a monthly R_r^2 of 0.43%, and $\sigma(Z_t^{-1})$ and β as implied by the variance model that uses only a lag of realized variance ($R_z^2 = 53\%$), we obtain an estimate of -0.16 for the second and last terms. This implies a minimum increase in expected return of 220% relative

to the baseline case of no timing. The main reason this number remains large is because the estimated risk return tradeoff in the data is fairly weak. Thus, while the conditional mean and conditional variance are not independent, they are not close to perfectly correlated either, meaning that a combination of information based off of each provides substantial gains.

7. Conclusion

This paper estimates the covariance between the price of risk and quantity of risk for a wide range of benchmark asset pricing factors across equities and currencies. In all cases the point estimate of this covariance is negative, and in almost all cases the confidence interval is below zero. The estimates are economically large and have strong implications for economic theory and portfolio choice. In portfolio choice terms, volatility managed portfolios offer superior risk adjusted returns and are easy to implement in real time. These portfolios lower risk exposure when volatility is high and increase risk exposure when volatility is low. Contrary to standard intuition, our portfolio choice rule would tell investors to sell during crises like the Great Depression or 2008 when volatility spiked dramatically so that investors behave in a “panicked” manner that we often think of as irrational. The facts are a puzzle from the perspective of economic theory because, if anything, theory generally posits that the price of risk will tend to rise in bad times when volatility is high. We demonstrate that structural models of asset pricing, which are successful in many other dimensions, generate a covariance between

the price and quantity of risk that is weakly positive rather than strongly negative as in the data. Finally, we conduct welfare implications for a mean-variance investor who times the market by observing the conditional mean and conditional volatility of stock returns. We find such an investor is better off paying attention to conditional volatility than the conditional mean by a fairly wide margin, suggesting that volatility is a key element of market timing.

References

- Andersen, T. G. and Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International economic review*, pages 885–905.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2001). Modeling and forecasting realized volatility. Technical report, National Bureau of Economic Research.
- Ang, A., Hodrick, R. J., Xing, Y., and Zhang, X. (2006). The cross-section of volatility and expected returns. *Journal of Finance*, 61(1):259–299.
- Bansal, R., Kiku, D., and Yaron, A. (2009). An empirical evaluation of the long-run risks model for asset prices. Technical report, National Bureau of Economic Research.

- Bansal, R. and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509.
- Barberis, N. (2000). Investing for the long run when returns are predictable. *The Journal of Finance*, 55(1):pp. 225–264.
- Barberis, N., Huang, M., and Santos, T. (2001). Prospect theory and asset prices. *The Quarterly Journal of Economics*, 116(1):1–53.
- Barroso, P. and Santa-Clara, P. (2015). Momentum has its moments. *Journal of Financial Economics*, 116(1):111–120.
- Campbell, J. Y. and Cochrane, J. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251.
- Campbell, J. Y. and Shiller, R. J. (1988). The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1(3):195–228.
- Campbell, J. Y. and Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, 21(4):1509–1531.
- Chacko, G. and Viceira, L. M. (2005). Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets. *Review of Financial Studies*, 18(4):1369–1402.

- Cochrane, J. H. (2009). *Asset Pricing:(Revised Edition)*. Princeton university press.
- Fama, E. F. and French, K. R. (1996). Multifactor explanations of asset pricing anomalies. *Journal of Finance*, 51(1):55–84.
- Frazzini, A. and Pedersen, L. H. (2014). Betting against beta. *Journal of Financial Economics*, 111(1):1 – 25.
- French, K. R., Schwert, G. W., and Stambaugh, R. F. (1987). Expected stock returns and volatility. *Journal of financial Economics*, 19(1):3–29.
- Glosten, L. R., Jagannathan, R., and Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The journal of finance*, 48(5):1779–1801.
- Hansen, L. P. and Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):pp. 225–262.
- He, Z. and Krishnamurthy, A. (2012). Intermediary asset pricing. *The American Economic Review*, forthcoming.
- Hou, K., Xue, C., and Zhang, L. (2014). Digesting anomalies: An investment approach. *Review of Financial Studies*, forthcoming.
- Jagannathan, R. and Wang, Z. (1996). The conditional capm and the cross-section of expected returns. *Journal of Finance*, 51(1):3–53.

- Kozak, S., Nagel, S., and Santosh, S. (2015). Interpreting factor models. *working paper*.
- Lettau, M. and Ludvigson, S. (2003). Measuring and modelling variation in the risk–return tradeoff, prepared for the handbook of financial econometrics, edited by y. *Ait-Sahalia and Lars–Peter Hansen*.
- Lundblad, C. (2007). The risk return tradeoff in the long run: 1836–2003. *Journal of Financial Economics*, 85(1):123–150.
- Lustig, H., Roussanov, N., and Verdelhan, A. (2011). Common risk factors in currency markets. *Review of Financial Studies*, page hhr068.
- Lustig, H. and Verdelhan, A. (2012). Business cycle variation in the risk–return trade-off. *Journal of Monetary Economics*, 59, Supplement(0):S35 – S49.
- Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1):1–28.
- Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, forthcoming.

8. Tables and Figures

Table 1: We regress each factors future return at $t + 1$ on its negative inverse variance at time t . The data is monthly and the sample is 1926-2015, except for the factors RMW and CMA which start in 1963, and the FX Carry factor which starts in 1983. Standard errors are in parentheses and adjust for heteroscedasticity.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Mkt	SMB	HML	Mom	RMW	CMA	MVE	FX
$-1/\sigma_{Mkt,t}^2$	0.15 (1.22)							
$-1/\sigma_{SMB,t}^2$		0.45 (0.26)						
$-1/\sigma_{HML,t}^2$			0.39 (0.31)					
$-1/\sigma_{Mom,t}^2$				-0.51 (0.36)				
$-1/\sigma_{RMW,t}^2$					-0.07 (0.16)			
$-1/\sigma_{CMA,t}^2$						0.26 (0.15)		
$-1/\sigma_{MVE,t}^2$							-0.08 (0.15)	
$-1/\sigma_{fx,t}^2$								-0.54 (0.61)
Constant	7.99 (3.07)	4.56 (2.00)	6.42 (2.43)	6.31 (2.81)	2.36 (2.23)	5.67 (1.77)	7.49 (1.43)	6.27 (2.72)
Observations	1,065	1,065	1,065	1,060	621	621	1,060	360
R-squared	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 2: We estimate the covariance between the price and quantity of risk for each factor and for the mean variance efficient portfolio (MVE). The data is monthly and the sample is 1926-2015, except for the factors RMW and CMA which start in 1963, and the FX Carry factor which starts in 1983. MVE2 computes the tangency portfolio for the later sample including RMW and CMA. Standard errors are in parentheses along with 95% confidence intervals (CI). Both are computed by bootstrap using Equation (4) in the text. The covariances are given in percentages.

	$cov(b_t, \sigma_t^2)$		
	Mean	Std Err	95% CI
MktRF	-0.50	0.17	[-0.84, -0.17]
SMB	0.03	0.12	[-0.21, 0.28]
HML	-0.27	0.13	[-0.52, -0.01]
Mom	-1.51	0.19	[-1.89, -1.12]
MVE	-0.62	0.09	[-0.80, -0.45]
RMW	-0.20	0.08	[-0.36, -0.05]
CMA	-0.06	0.07	[-0.19, 0.07]
MVE2	-0.29	0.06	[-0.41, -0.17]
Carry	-0.20	0.07	[-0.33, -0.06]

Table 3: We estimate the lower bound for the monthly volatility of the price of risk for each factor, $\sigma(b_t)$. The bound uses our estimate of the covariance between the price and quantity of risk, which can be expressed as the correlation between price and quantity of risk multiplied by the volatility of the price of risk and volatility of quantity of risk. We use an estimate for the monthly volatility of the quantity of risk and the fact that correlations have absolute value less than or equal to 1 to estimate this lower bound. The last column compares this to the unconditional price of risk for each factor, $E[R]/\sigma^2(R)$, to give a sense of how volatile the monthly price of risk is in relation to its unconditional value.

Lower bound on volatility of price of risk, $\sigma(b_t)$

	Lower bound for $\sigma(b_t)$	95% CI	Unconditional price of risk
MktRF	1.77	[0.59, 2.86]	2.24
SMB	0.33	[0.04, 2.97]	2.11
HML	2.71	[0.35, 5.25]	3.13
Mom	9.23	[7.11, 11.38]	2.98
MVE	8.70	[6.42, 11.12]	9.83
RMW	5.61	[1.14, 10.74]	5.48
CMA	1.93	[0.13, 6.72]	7.94
MVE2	12.79	[8.31, 17.47]	14.82
Carry	5.84	[2.06, 9.81]	9.72

Table 4: We run time-series regressions of each managed factor on the non-managed factor. The managed factor scales by the factors inverse realized variance in the preceding month. The data is monthly and the sample is 1926-2015, except for the factors RMW and CMA which start in 1963, and the FX Carry factor which starts in 1983. Standard errors are in parentheses and adjust for heteroscedasticity. All factors are annualized in percent per year by multiplying monthly factors by 12.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	S Mkt	S SMB	S HML	S Mom	S RMW	S CMA	S MVE	S FX
MktRF	0.61 (0.05)							
SMB		0.62 (0.08)						
HML			0.57 (0.07)					
Mom				0.47 (0.07)				
RMW					0.62 (0.08)			
CMA						0.68 (0.05)		
MVE							0.58 (0.03)	
Carry								0.71 (0.08)
Constant	4.86 (1.56)	-0.58 (0.91)	1.97 (1.02)	12.51 (1.71)	2.44 (0.83)	0.38 (0.67)	4.12 (0.77)	2.78 (1.49)
Observations	1,065	1,065	1,065	1,060	621	621	1,060	360
R-squared	0.37	0.38	0.32	0.22	0.38	0.46	0.33	0.51
rmse	51.39	30.44	34.92	50.37	20.16	17.55	25.34	21.78

Table 5: We run time-series regressions of each managed factor on the non-managed factor. Here our managed portfolios make use of the full forecasting regression for log variances rather than simply scaling by lagged realized variances. The data is monthly and the sample is 1926-2015, except for the factors RMW and CMA which start in 1963, and the FX Carry factor which starts in 1983. Standard errors are in parentheses and adjust for heteroscedasticity. All factors are annualized in percent per year by multiplying monthly factors by 12.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	S Mkt	S SMB	S HML	S Mom	S RMW	S CMA	S MVE	S FX
MktRF	0.73 (0.06)							
SMB		0.71 (0.09)						
HML			0.65 (0.08)					
Mom				0.59 (0.08)				
RMW					0.70 (0.08)			
CMA						0.78 (0.05)		
MVE							0.74 (0.03)	
Carry								0.89 (0.05)
Constant	3.85 (1.36)	-0.60 (0.78)	2.09 (0.92)	12.54 (1.67)	1.95 (0.75)	0.41 (0.57)	3.83 (0.67)	1.77 (0.90)
Observations	1,063	1,063	1,063	1,059	619	619	1,059	358
R-squared	0.53	0.51	0.43	0.35	0.49	0.61	0.54	0.81
rmse	44.33	27.02	32.06	46.01	18.31	14.96	20.97	13.66

Table 6: We run time-series regressions of each managed factor on the non-managed factor plus the betting against beta (BAB) factor. from Frazzini and Pedersen. The data is monthly and the sample is 1929-2012 based on availability of the BAB factor. Standard errors are in parentheses and adjust for heteroscedasticity. All factors are annualized in percent per year by multiplying monthly factors by 12.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	S Mkt	S SMB	S HML	S Mom	S RMW	S CMA	S MVE
MktRF	0.60 (0.05)						
BAB	0.09 (0.06)	0.01 (0.05)	0.02 (0.05)	-0.07 (0.04)	-0.13 (0.02)	-0.06 (0.02)	0.04 (0.02)
SMB		0.61 (0.09)					
HML			0.56 (0.07)				
Mom				0.47 (0.06)			
RMW					0.65 (0.08)		
CMA						0.69 (0.04)	
MVE							0.57 (0.04)
Constant	3.83 (1.80)	-0.77 (1.10)	2.05 (1.15)	13.52 (1.86)	3.97 (0.89)	0.94 (0.71)	4.10 (0.85)
Observations	996	996	996	996	584	584	996
R-squared	0.37	0.37	0.31	0.21	0.40	0.46	0.33
rmse	52.03	31.36	35.92	51.73	19.95	17.69	26.01

Table 7: We run time-series regressions of each managed factor on the 4 Fama-French Carhart factors. The data is monthly and the sample is 1926-2015. Standard errors are in parentheses and adjust for heteroscedasticity. All factors are annualized in percent per year by multiplying monthly factors by 12.

VARIABLES	(1)	(2)	(3)	(4)	(5)
	S Mkt	S SMB	S HML	S Mom	S MVE
MktRF	0.70 (0.05)	-0.02 (0.01)	-0.10 (0.02)	0.16 (0.03)	0.23 (0.02)
HML	-0.03 (0.05)	-0.02 (0.04)	0.63 (0.05)	0.09 (0.05)	0.08 (0.02)
SMB	-0.05 (0.06)	0.63 (0.08)	-0.00 (0.05)	-0.10 (0.04)	-0.15 (0.02)
Mom	0.25 (0.03)	0.01 (0.03)	0.06 (0.04)	0.54 (0.05)	0.30 (0.02)
Constant	2.43 (1.60)	-0.42 (0.94)	1.96 (1.06)	10.52 (1.60)	4.47 (0.77)
Observations	1,060	1,060	1,060	1,060	1,060
R-squared	0.42	0.38	0.35	0.25	0.35
rmse	49.56	30.50	34.21	49.41	25.13

Table 8: We run time-series regressions of each managed factor on the 6 Fama-French Carhart factors. The data is monthly and the sample is 1963-2015. Standard errors are in parentheses and adjust for heteroscedasticity. All factors are annualized in percent per year by multiplying monthly factors by 12.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	S Mkt	S SMB	S HML	S Mom	S RMW	S CMA	S MVE	S MVE2
MktRF	0.79 (0.05)	0.03 (0.03)	-0.06 (0.03)	0.12 (0.04)	0.02 (0.02)	0.02 (0.01)	0.26 (0.03)	0.23 (0.02)
HML	0.11 (0.09)	0.09 (0.06)	1.03 (0.08)	0.15 (0.09)	-0.21 (0.04)	0.03 (0.03)	0.16 (0.04)	0.05 (0.03)
SMB	0.02 (0.05)	0.75 (0.05)	-0.05 (0.04)	-0.12 (0.07)	-0.02 (0.03)	-0.03 (0.02)	-0.15 (0.03)	-0.09 (0.02)
Mom	0.15 (0.03)	-0.01 (0.03)	0.05 (0.03)	0.64 (0.08)	-0.00 (0.02)	-0.02 (0.02)	0.32 (0.03)	0.23 (0.02)
RMW	0.15 (0.06)	0.23 (0.07)	-0.56 (0.08)	-0.04 (0.08)	0.64 (0.06)	-0.18 (0.04)	0.01 (0.04)	0.04 (0.03)
CMA	0.04 (0.12)	0.00 (0.07)	-0.28 (0.10)	-0.25 (0.11)	-0.00 (0.06)	0.63 (0.05)	-0.04 (0.06)	0.14 (0.04)
Constant	0.18 (1.87)	-1.68 (1.25)	4.16 (1.44)	12.91 (2.17)	3.21 (0.81)	1.07 (0.72)	4.00 (1.02)	3.03 (0.77)
Observations	622	622	622	622	621	621	622	621
R-squared	0.47	0.49	0.51	0.31	0.46	0.50	0.40	0.43
rmse	42.70	26.82	32.82	48.10	18.85	17.01	23.26	16.96

Figure 1: This figure plots the time-series of the monthly volatility of each individual factor. We emphasize the common co-movement in volatility across factors. Light shaded bars indicate recessions and show a clear business cycle pattern in volatility.

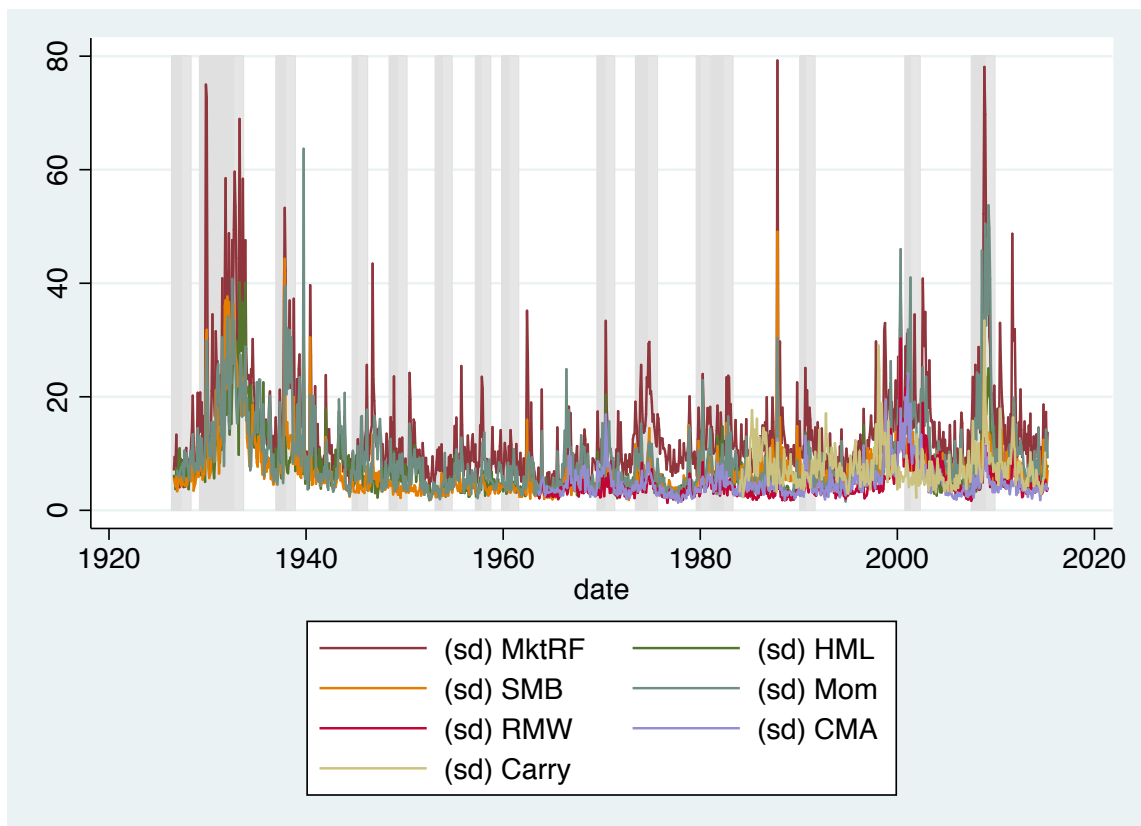


Figure 2: The figure plots the full distribution of scaled factors vs non-scaled factors estimated using kernel density. In particular, for each panel we plot the distribution of f_{t+1} (solid line) along with the distribution of $z_t f_{t+1}$ (dashed line).

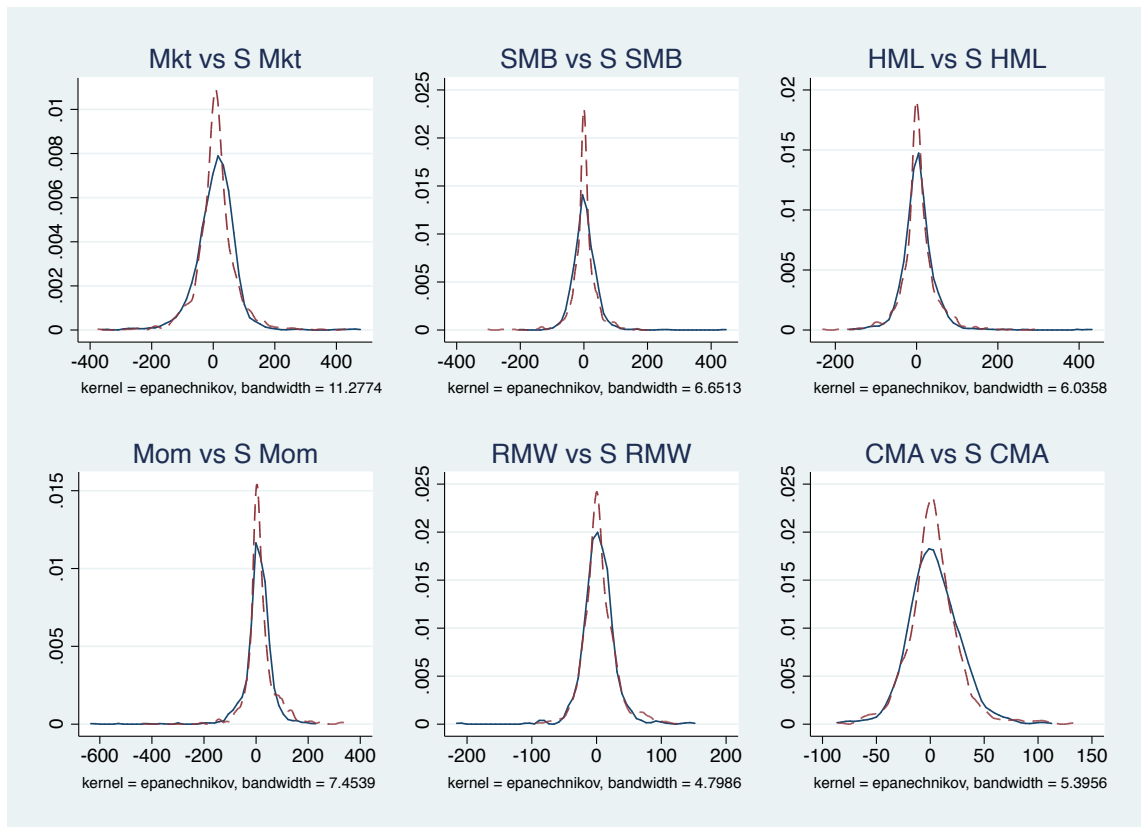
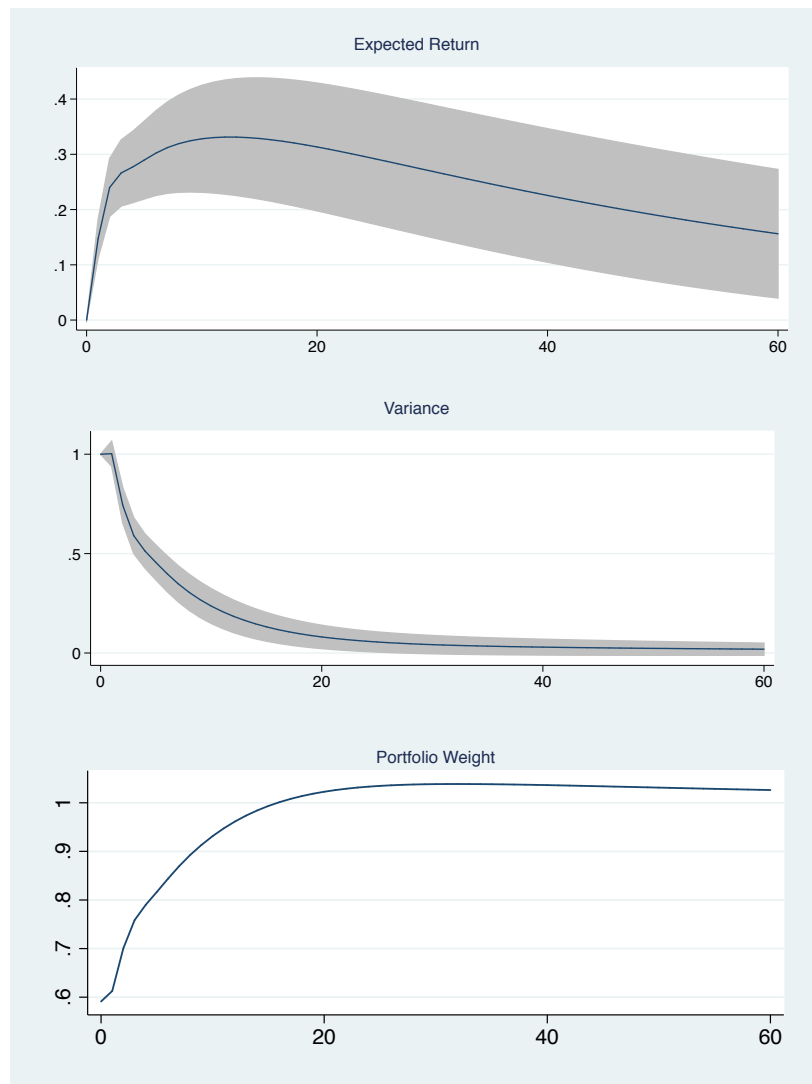


Figure 3: The figure plots the impulse response of the expected return and variance of the market portfolio for a shock to the variance. The x-axis is months. The last panel gives the portfolio choice implications for the differential movements in expected returns and variance to a variance shock at different horizons and are computed for a mean variance optimizing agent who holds the market portfolio on average. Expected returns are formed using a forecasting regression of future 1 month returns on Shiller's CAPE measure, the BaaAaa default spread, and forecasted volatility. Variance is the expected variance formed from our forecasting model. We compute impulse responses using a VAR with 3 monthly lags of each variable.



A. Derivations

A.1 Proof of Proposition 1

Let f_{t+1} be the excess return on the factor. Let μ_t denote the conditional expected return on the factor, $E_t[f_{t+1}]$, and σ_t^2 denote the conditional variance, $var_t[f_{t+1}]$. It is useful to keep in mind the equation

$$\mu_t = b_t \sigma_t^2 \quad (27)$$

The coefficient in the regression of future realized excess return f_{t+1} on past inverse volatility is

$$\beta = \frac{cov(f_{t+1}, \frac{-1}{\sigma_t^2})}{var(\frac{1}{\sigma_t^2})} \quad (28)$$

We can replace f_{t+1} with $f_{t+1} = \mu_t + \epsilon_{t+1}$ in the above equation. Note that the error term ϵ_{t+1} will have zero covariance with any time t variable. Next replace $\mu_t = b_t \sigma_t^2$ so that we have

$$\beta = \frac{cov(b_t \sigma_t^2, \frac{1}{\sigma_t^2})}{var(\frac{1}{\sigma_t^2})} \quad (29)$$

Next, we make use of the following identity in order to simplify the numerator above:

$$cov(xy, z) = cov(x, y(z - E[z])) + E[x]cov(y, z) \quad (30)$$

This means

$$\beta = \left[cov(b_t, \sigma_t^2 (E[1/\sigma_t^2] + cov(\sigma_t^2, 1/\sigma_t^2)/E[\sigma_t^2])) - \frac{\mu}{E[\sigma_t^2]} cov(\sigma_t^2, 1/\sigma_t^2) \right] / var(1/\sigma_t^2) \quad (31)$$

But then

$$\text{cov}(\sigma_t^2, 1/\sigma_t^2) = 1 - E[\sigma_t^2]E[1/\sigma_t^2] \quad (32)$$

So that

$$\beta = \left[\text{cov}(b_t, \sigma_t^2)/E[\sigma_t^2] + \frac{\mu}{E[\sigma_t^2]}(E[\sigma_t^2]E[1/\sigma_t^2] - 1) \right] / \text{var}(1/\sigma_t^2) \quad (33)$$

Rearranging,

$$\frac{1}{E[\sigma_t^2]\text{var}(1/\sigma_t^2)}\text{cov}(b_t, \sigma_t^2) = \beta - \frac{\mu}{E[\sigma_t^2]} \frac{(E[\sigma_t^2]E[1/\sigma_t^2] - 1)}{\text{var}(1/\sigma_t^2)} \quad (34)$$

which is the expression found in the main text.

B. Additional Tables

Table 9: We show that volatility is persistent for each individual factor by running a regression of current volatility on lagged volatility. Volatility is computed monthly by using daily observations for the given month. The data is monthly and the sample is 1926-2015, except for the factors RMW and CMA which start in 1963, and the FX Carry factor which starts in 1983. Standard errors are in parentheses and adjust for heteroscedasticity.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$\ln \sigma_{Mkt}$	$\ln \sigma_{SMB}$	$\ln \sigma_{HML}$	$\ln \sigma_{Mom}$	$\ln \sigma_{RMW}$	$\ln \sigma_{CMA}$	$\ln \sigma_{MVE}$	$\ln \sigma_{FX}$
$\ln \sigma_{Mkt,t-1}$	0.72 (0.02)							
$\ln \sigma_{SMB,t-1}$		0.78 (0.02)						
$\ln \sigma_{HML,t-1}$			0.78 (0.02)					
$\ln \sigma_{Mom,t-1}$				0.76 (0.02)				
$\ln \sigma_{RMW,t-1}$					0.78 (0.03)			
$\ln \sigma_{CMA,t-1}$						0.74 (0.03)		
$\ln \sigma_{MVE,t-1}$							0.69 (0.02)	
$\ln \sigma_{FX,t-1}$								0.52 (0.05)
Constant	0.69 (0.05)	0.40 (0.04)	0.40 (0.03)	0.47 (0.04)	0.30 (0.04)	0.39 (0.04)	0.55 (0.04)	0.93 (0.10)
Observations	1,064	1,064	1,064	1,060	620	620	1,060	372
R-squared	0.52	0.61	0.60	0.58	0.61	0.54	0.48	0.27