

Volatility Managed Portfolios

Alan Moreira and Tyler Muir*

November 1, 2015

Abstract

Managed portfolios that take less risk when volatility is high produce large, positive alphas and increase factor Sharpe ratios by substantial amounts. We document this fact for the market, value, momentum, profitability, return on equity, and investment factors in equities, as well as the currency carry trade. Our portfolio timing strategies are simple to implement in real time and are contrary to conventional wisdom because volatility tends to be high after the onset of recessions and crises when selling is typically viewed as a mistake. Instead, our strategy earns high average returns while taking less risk in recessions. We study the portfolio choice implications of these results. We find volatility timing provides large utility gains to a mean variance investor, with increases in lifetime utility ranging from 50-90%. We then study the problem of a long horizon investor and show that, perhaps surprisingly, long horizon investors can benefit from volatility timing even when time-variation in volatility is completely driven by time-varying discount rate volatility.

*Yale School of Management. We thank Jon Ingersoll, Lu Zhang, and seminar participants at Yale SOM, Ohio State, and Baruch for comments. We especially thank Nick Barberis for many useful discussions. We also thank Ken French for making data publicly available and Adrien Verdelhan and Lu Zhang for providing data.

1. Introduction

We construct portfolios that scale monthly returns by the inverse of their expected variance, decreasing risk exposure when the returns variance is expected to be high and vice versa. We call these volatility managed portfolios. We document that this simple trading strategy earns large alphas across a wide range of asset pricing factors, suggesting that investors can benefit from volatility timing. Motivated by these results, we study the portfolio choice implications of time-varying volatility. We find that both short and long-term investors alike can benefit from volatility timing and that utility gains are substantial. Further, we show that the optimal portfolio can be approximated by a combination of a buy-and-hold portfolio and the volatility managed portfolio that we introduce in this paper.

We motivate our analysis from the vantage point of a simple mean variance investor, who adjusts his allocation in the risky asset according to the attractiveness of the mean variance trade-off, $E_t[R_{t+1}]/Var_t(R_{t+1})$. Because variance is highly forecastable at horizons up to one year, and variance forecasts are only weakly related to future returns at these horizons, our volatility managed portfolios produce significant risk adjusted returns for the market, value, momentum, profitability, return on equity, and investment factors in equities as well as for the currency carry trade. In addition we show the strategy survives transaction costs and works for most international indices as well. Annualized alphas with respect to the original factors are substantial, and Sharpe ratios increase by 50-100% of the original factor Sharpe ratios. Utility benefits of volatility timing for a mean variance investor are on the order of 50-90% of lifetime utility – substantially larger than those coming from expected return timing. Moreover, parameter instability for an agent that estimates volatility in real time is negligible, in contrast to strategies that try to time expected returns (Goyal and Welch (2008)).

Our volatility managed portfolios reduce risk taking after market crashes and volatility spikes, while the common advice is to increase or hold risk taking constant after mar-

ket downturns. Thus, on average, our volatility managed portfolios reduce risk exposure in recessions. For example, in the aftermath of the sharp price declines and huge spikes in volatility in the fall of 2008, it was a widely held view that market movements created a once in a generation buying opportunity, and that those that reduced positions in equities were making a large mistake. Yet the volatility managed portfolio cashed out almost completely and returned to the market only as the spike in volatility receded. We show that in fact our simple strategy turned out to work well historically and throughout several crisis episodes including the Great Depression, Great Recession, and 1987 stock market crash.

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 and BaaAaa spread. We run a VAR which includes both the conditional variance and conditional expected return and show that, in response to a variance shock, the conditional variance increases initially by far more than the expected return, making the risk-return trade-off initially unattractive. A mean variance investor would decrease his risk exposure by 60% after a 1 standard deviation shock to the market variance. However, since volatility movements are less persistent than movements in expected returns, our optimal portfolio strategy prescribes a gradual increase in the exposure as the initial volatility shock fades. On average it takes 18 months for portfolio exposure to return to normal and at horizons beyond this it is optimal to increase exposure further to capture the persistent increase in expected returns. The difference in persistence allows an investor to keep much of the expected return benefit while at the same time reducing his overall risk exposure.

Having understood these results from the perspective of a short horizon mean-variance investor, we then study the portfolio choice problem of a long horizon investor. Indeed,

an important aspect of our analysis is the role of investment horizon. The evidence that equity returns are somewhat predictable implies that investors with a long horizon should perceive equities as less risky. For example, suppose returns were extremely mean reverting so that a negative realization on any given day would be followed by a positive realization the following day of the same magnitude. Then an investor with a one day horizon would consider stocks risky, but an investor with a two day horizon would perceive stocks as perfectly safe. In contrast, if returns are independent across days, so that a bad realization today gives no information about returns tomorrow, then the investors horizon would not matter. This is the classic insight from Samuelson (1969) and Merton (1973). Empirically, returns have both features – there is some mean reversion in returns but also a large fraction of returns are permanent price changes that are independent of returns in the future (Campbell and Shiller, 1988). As a result of this empirical evidence, fluctuations in return volatility can come in two different types. They can be driven by shocks that are permanent or shocks that eventually mean-revert.

While a short horizon investor is indifferent and responds uniformly to changes in volatility, the type of volatility matters greatly to a long horizon investor. If the increase in volatility is due to increased volatility of the permanent component of returns, then a long horizon investor will respond more aggressively to volatility changes than a short horizon investor. Intuitively, stocks become relatively riskier as the share of permanent shocks increases. But if the increase in volatility is instead due to increased volatility of the transitory or mean reverting portion of returns, then a long horizon will respond less aggressively than a short horizon investor. The reason the long horizon investor perceives transitory shocks as less risky, is that he or she can wait until the price recovers. Thus, it matters how quickly this mean reversion takes place. Empirically, it appears that while returns do display mean reversion, this mean reversion plays out over many years — Sharpe ratios for stocks increase only slowly with investment horizon and valuation ratios that predict returns are highly persistent with auto-correlation close to one. Given the

time it takes for these transitory movements in prices to mean revert, our empirical estimates suggest that a long horizon investor will still care substantially about time varying volatility even if it is purely related to the mean reverting portion of returns. Thus, even long horizon investors will find some degree of volatility timing beneficial.

Motivated by these insights, we study how to implement the optimal portfolio for long-horizon investors. We find that a simple two fund theorem holds: all investors, regardless of horizon, will want to hold a linear combination of a passive buy and hold portfolio and our volatility managed portfolio. Each investor will choose *static* weights on these two funds. These weights will depend both on investment horizon and on whether volatility moves through time because of the permanent or transitory part of returns.¹ First, short horizon or mean variance agents will put no weight on the passive portfolio, and will instead put all their weight on our volatility managed portfolio, regardless of any other parameters. For long horizon investors, their weight on the volatility managed portfolio will be high when time-variation in volatility comes from the permanent part of returns, but will be lower when time-variation in volatility comes from the mean reverting or transitory portion of returns. For a calibration where investor horizon is 30 years, risk aversion is $\gamma = 10$, and *all* time variation in volatility is due to the mean-reverting portion of returns, the long horizon investor will load about half as much on our volatility managed portfolio as a short horizon investor would. This suggests even long horizon investors will generally find a fairly large amount of volatility timing beneficial. However, we again note that these results depend on the empirical estimates for how fast prices mean revert, taken from the persistence of price dividend ratios. To the extent that mean reversion is faster, long-horizon investors will load less on the volatility managed portfolio.

Finally, in studying investment horizon, we are better equipped to understand the reaction of agents during large market downturns. We use October 2008 as an example,

¹Weights will of course depend on other parameters such as risk aversion as well.

where volatility was around 60% or more, but prices had also fallen making expected returns likely higher. This was also a time period where the common wisdom was not to sell or panic, but instead to buy as prices had fallen dramatically (e.g., Buffett (2008)). Our framework suggests that if the higher volatility at the time was not due to the permanent part of returns but instead due completely to the transitory part of returns, then long horizon investors may indeed have less cause for concern. However, unless it was believed that this mean reversion in this particular episode would occur more quickly than usual, our calibration would still suggest that long horizon investors would want to sell and only return once volatility had subsided.

This paper proceeds as follows. Section 2 reviews additional related literature. Section 3 shows our main empirical results related to our volatility managed portfolios. Section 4 discusses portfolio choice implications and derives dynamic market timing rules. Section 4.2 compares the welfare benefit to a mean-variance investor of forecasting the variance versus the conditional mean of stock returns. Section 5 discusses implications for structural asset pricing models. Section 6 concludes.

2. Literature review

Our results build 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 the quality of our forecasts. This is important because it shows that a rather naive investor can implement these strategies in real time. Our volatility timing results are related to Fleming et al. (2001) and Fleming et al. (2003) who estimate the full variance co-variance matrix of returns across asset classes (stocks,

bonds, and gold) and use this to make asset allocation decisions across these asset classes at a daily frequency.

The second strand of literature debates 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. Even if this relationship is positive, volatility timing can still be beneficial if expected returns do not rise by *enough* compared to increases in volatility. Moreover, we take a portfolio strategy approach to this view by showing portfolios can be formed in real time that take advantage of the risk-return regressions and produce very large risk-adjusted returns.

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. A notable related set of papers is Fleming et al. (2001) and Fleming et al. (2003) who conduct asset allocation across assets at daily frequencies by estimating the conditional covariance matrix across

assets which mixes both cross-sectional and time-series approaches. We study volatility timing on our factors *individually*, employ many more factors, and use monthly or longer horizons to assess the benefits, making our results apply to average investors. By focusing on systematic risk factors, we are also able to say something about the price of risk over time. As an example, volatility timing on an individual stock will not tell us about risk compensation over time if the majority of the stock's volatility is idiosyncratic, as typically appears to be the case.

3. Empirical Results

3.1 Data Description

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.² 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

²We thank Adrien Verdelhan for help with the currency daily data.

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.

3.2 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

$$\frac{c}{RV_t^2} f_{t+1} \quad (1)$$

We choose the constant c 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 will 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 1 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 the two factors (SMB and CMA) that have stronger co-movement between volatility and future returns (see Table 12), 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 in-

vestor in real time and it relates to a long literature in market timing that we refer to later.³ 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 $\sqrt{SR_{old}^2 + \left(\frac{\alpha}{\sigma_\epsilon}\right)^2} - SR_{old}$ where SR_{old} is the maximum Sharpe ratio given by the original non-scaled factors. For example, in Table 1, 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 0.34.⁴

Table 2 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 3, 4, and 5 add additional controls to the regressions and compute alphas relative to larger multi-factor models. Notably, Table reftable: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 4, and 5 show that the scaled factors expand the mean variance frontier of

³The average investor will likely have trouble trading the momentum factor, for example.

⁴ 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.

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

Figure 4 plots the cumulative returns to the volatility managed market factor compared to a buy and hold strategy from 1926-2015. We invest \$1 in both in 1926 and plot the (real) cumulative returns to each on a log scale.

As a robustness check, we also find that our strategies survive transaction costs though this is beyond the goal of our paper. These results are given in Table 10. Specifically, we evaluate our volatility timing strategy for the market portfolio when including reasonable transactions costs. We consider alternative strategies that capture volatility timing but reduce trading activity, including using volatility instead of variance, using expected rather than realized variance, and only trading when variance is above the long run average. Each of these reduces trading and hence reduces transactions costs. We report the average absolute change in monthly weights, expected return, and alpha of each strategy. Then we report the alpha when including various trading cost assumptions. The 1bps cost comes from Fleming et al. (2003), the 10bps comes from Frazzini et al. (2015) when trading about 1% of daily volume, and the last column adds an additional 4bps to account for transaction costs increasing in high volatility episodes. Specifically, we use the slope coefficient of transactions costs on VIX from Frazzini et al. (2015) and evaluate this impact on a move in VIX from 20% to 40% which represents the 98th percentile of VIX. Finally, the last column backs out the implied trading costs in basis points needed to drive our alphas to zero in each of the cases. The results indicate that the strategy survives transactions costs, even in high volatility episodes where such costs likely rise (indeed we take the extreme case where VIX is at its 98th percentile). Alternative strategies that reduce trading costs are much less sensitive to these costs. Overall, we show that the annualized alpha of this strategy decreases somewhat for the market portfolio, but it still strongly

positive. We do not report results for all factors, since again this is not explicitly the goal of our paper, but we point out that realized volatility for the market varies by much more than for the other factors, implying more volatile weights and more trading. Hence, the trading costs for other factors is likely to be less. Note, however, we don't study trading costs of the original strategies (i.e., the costs of implementing the original momentum strategy as opposed to the volatility managed version).

We also consider strategies that impose leverage constraints. A simple strategy is one that only updates the portfolio when volatility is above its mean value. This portfolio both trades less frequently and also avoids the use of leverage from low volatility episodes where the risk weight would normally rise substantially. In unreported results, we find such constraints weaken the strategies considered, but not substantially so. There are still large alphas and substantial gains in Sharpe ratios.

As an additional robustness check, we demonstrate that our results hold when studying 20 OECD countries and focus on the broad stock market indices of each country. On average, the managed volatility version of the index has an annualized Sharpe ratio that is 0.15 higher than a passive buy and hold strategy, representing a substantial increase. The volatility managed index has a higher Sharpe ratio than the passive strategy in 80% of cases. These results are detailed in Figure 7 of our Appendix. Note that is a strong condition – a portfolio can still have positive alpha even when its Sharpe ratio is below the non-managed factor. The main text is devoted to understanding the better studied US factors.

3.3 Understanding the profitability of volatility timing

We first give intuition for why our volatility managed portfolios work in terms of generating positive risk adjusted returns. Then, we discuss how to reconcile these results with the return predictability literature. At first, it may sound contradictory that our volatility portfolios decrease risk exposure after large market downturns when volatility increases.

These are times when we think expected returns are rising. We show that the frequency of these two variables behave quite differently. Volatility tends to spike and recede quickly whereas expected returns are more persistent. This reconciles the two findings.

We start by showing that volatility for each factor doesn't predict the factor's future returns. We run monthly regressions of future 1 month returns on monthly realized volatility for each factor. Using log realized volatility or using realized variance, does not change the results.

Table 6 gives these results. We can see that coefficients across factors range from positive to negative but are, generally speaking, not significant. Therefore, we don't see any clear relationship between a factors' volatility and its future expected return. Mechanically, this is why the strategy we implement works. If a factors volatility does not predict an increase in returns, then an increase in volatility signifies a poor risk return tradeoff where reducing risk exposure is optimal. Our appendix maps this out more clearly by showing the exact conditions under which our strategy generates alpha. But, intuitively, the lack of risk return tradeoff in the data means volatility timing will be profitable.

Next, we try to understand our results in light of the return predictability literature. To better understand the co-movement between expected returns and conditional variance in the data, 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 to a volatility shock. We set risk aversion just above 2, where our choice is set so that the investor holds the market on average when using the unconditional value for variance and the equity premium (i.e., in the absence of movements in expected returns and vol, the portfolio weight is $w = 1$). 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 using a log normal model for volatility and including 3 lags for each factors realized volatility.

$$\ln RV_{t+1} = a + \sum_{j=1}^J \rho_j \ln RV_{t-j-1} + cx_t + \varepsilon_{t+1} \quad (2)$$

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

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

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.

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 portfolio choice 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 his 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 decreases 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. Thus, our volatility timing results are not in conflict with expected return timing results. Instead, after a large market crash such as October 2008, our strategy says to get out immediately to avoid an unfavorable risk return tradeoff, but by buying back in when the volatility shock subsides, one can capture much of the expected return increase.

We further push this intuition in Figure 5 where we show the behavior of our volatility timing strategy during four prominent market downturns: the Great Depression, the 2008 recession, the 1987 stock market crash, and the oil crisis and stock market downturn of 1974-1975. In all episodes one can see our strategy decreasing risk exposure after the initial decline in the market, and gradually increasing risk exposure as the episode moves on.

4. Portfolio choice implications

In this section, we focus on understanding the implications for portfolio choice of our findings. We start by computing the utility gain of volatility timing for a mean-variance investor, and find that these gains to be substantially larger than utility gains arising from standard expected return timing strategies. We then study the portfolio problem of a long-horizon investor.

4.1 The benefits of volatility timing for a mean-variance investor

We now use simple mean-variance preferences to give a sense of the benefits of timing fluctuations in conditional factor volatility. 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.

4.1.1 Expected return timing

Consider the following excess return process,

$$r_{t+1} = \mu + x_t + \sqrt{Z_t}e_{t+1} \quad (3)$$

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,

$$w = \frac{1}{\gamma} \frac{\mu}{E[Z_t] + \sigma_x^2}; \quad (4)$$

$$w(x_t) = \frac{1}{\gamma} \frac{\mu + x_t}{E[Z_t]}; \quad (5)$$

where Equations (4) and (4) reflect alternative conditioning sets. Equation (4) uses no conditional information, while Equation (4) 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:

$$E[wr_{t+1}] = \frac{1}{\gamma} \frac{\mu^2}{\sigma_x^2 + E[Z_t]} = \frac{1}{\gamma} S^2 \quad (6)$$

$$E[w(x_t)r_{t+1}] = \frac{1}{\gamma} \frac{\mu + \sigma_x^2}{E[Z_t]} = \frac{1}{\gamma} \frac{(S^2 + R_f^2)}{1 - R_f^2}, \quad (7)$$

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

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%.

4.1.2 Volatility timing

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 as in Campbell and Thompson (2008). We then study the general case which allows for both signals to operate simultaneously.

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

$$z_t = \bar{z} + y_t + \sigma_z \times u_t. \quad (8)$$

This produces optimal portfolio weights and expected returns given by

$$w(z_t) = \frac{1}{\gamma} \frac{\mu}{e^{2(\bar{z} + y_t + \sigma_z^2)}}; \quad (9)$$

$$E[w(z_t)r_{t+1}] = \frac{1}{\gamma} S^2 e^{\sigma_y^2}, \quad (10)$$

where this Equation (10) 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 gain 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 $R_z^2 = 38\%$, 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 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 substantially higher than is documented in the literature on return predictability. Moreover, even if some variables 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 a expected return timing strategy in terms of utility gain.

Importantly, volatility timing can be implemented not only for the aggregate market, but for the several additional factors studied in Section 3, 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 aggregate market portfolio do not necessarily apply to other portfolios such as value, momentum, or the currency carry trade. Thus, we would need to come up with additional return forecasting variables for each of the different factors. Volatility timing on the other hand is easy to replicate across factor because 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 Campbell and Thompson (2008) 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,

$$\begin{aligned} w(Z_t, x_t) &= \frac{1}{\gamma} \frac{\mu + x_t}{Z_t}; \\ E[w(Z_t, x_t)r_{t+1}] &= \frac{1}{\gamma} \left(\frac{S^2 + R_f^2}{1 - R_f^2} E[Z_t]E[Z_t^{-1}] + 2\mu cov(x_t, Z_t^{-1}) + cov(x_t^2, Z_t^{-1}) \right), \end{aligned} \quad (11)$$

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 $R_z^2 = 53\%$ for the variance signal, one would obtain a 236% increase in expected returns. But there is some risk-return trade-off in the data, thus we need to consider the other terms as well.

The second term we can construct directly from our estimates in Table 12, 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

$$\text{cov}(x_t^2, Z_t^{-1}) \geq -1\sigma(x_t^2)\sigma(Z_t^{-1}) = -1\sqrt{2}\sigma_x^2\sigma(Z_t^{-1}), \quad (12)$$

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

$$E[w(Z_t, x_t)r_{t+1}] \geq \frac{1}{\gamma} \left(\frac{S^2 + R_r^2}{1 - R_r^2} E[Z_t]E[Z_t^{-1}] - 2\beta\mu\sigma^2(Z_t^{-1}) - \sqrt{2}\sigma_x^2\sigma(Z_t^{-1}) \right). \quad (13)$$

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 trade-off 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 the combination of information provides additional gains.

4.2 Volatility timing for a long-horizon investor

We now study the problem of a long-horizon investor and investigate how much he or she can benefit from volatility timing. Our analysis is motivated by the idea that a large fraction of stock price movements are transitory, implying that stock market volatility might not be the relevant measure of risk for a long-horizon investor.

Formally, the idea is that changes in volatility might be mostly driven by *discount rate* volatility. A high discount rate volatility means that you might wake-up poorer tomorrow, but expected returns moved exactly so you expect to be just as rich fifty years from now. Thus, according to this view increases in volatility does not pose a risk to a long run investor because these discount rate shocks wash out in the long run. John Cochrane articulates this view nicely in a 2008 Wall Street Journal article:

And what about volatility?(...) if you were happy with a 50/50 portfolio with an expected return of 7% and 15% volatility , 50% volatility means you should hold only 4.5% of your portfolio in stocks! (...) expected returns would need to rise from 7% per year to 78% per year to justify a 50/50 allocation with 50% volatility. (...) The answer to this paradox is that the standard formula is wrong. (...) Stocks act a lot like long-term bonds – when prices decline and dividend yields rise, subsequent returns rise as well.(...)If bond prices go down more, bond yields and long-run returns will rise just enough that you face no long-run risk.(...)the same logic explains why you can ignore "short-run" volatility in stock markets.

Our goal here is to understand better this advice: how long-run an investment horizon needs to be for investors to safely ignore discount rate volatility? We find that only extremely long run investors can safely ignore discount rate volatility. For example, we show that even an investor with a 30 year horizon still cares about discount rate volatility and finds it optimal to reduce his portfolio weight when volatility goes up. The key for this result is that, in the data, shocks to expected returns seem to be very persistent, so an investor has to be indifferent with respect to fluctuations in the value of his or her wealth for many years in the future. To the extent that these expected returns were less persistent, long horizon investors would perceive increases in discount rate volatility as less risky.

Earlier work on portfolio choice has studied expected return variation, volatility variation, or volatility variation with a constant risk-return trade-off.⁵ In order to understand how much long-horizon investor can benefit from the empirical patterns we document in the data, we solve the dynamic portfolio problem in a stochastic environment that (i) allows for independent variation in expected returns and volatility; and (ii) allows for variation in the mix of cash-flow and discount-rate shocks. These two ingredients are

⁵Example of work that study volatility timing in a dynamic environment are Chacko and Viceira (2005) and Liu (2007).

novel and essential to interpret our findings.

We now discuss the assumed stochastic environment, followed by investors preferences, and a brief description of the portfolio optimization problem. Our interest is quantitative, so we focus on contrasting how the long-horizon investor responds differently to time-variation in volatility.

4.2.1 Investment opportunity set

We assume there are two assets. A risk-less bond that pays constant interest rate $r dt$, and risky asset S_t , with dynamics given by

$$\frac{dS_t}{S_t} = (r + x_t)dt + \sqrt{y_t}DdB_t + FdZ_t, \quad (14)$$

where S_t is the value of a portfolio fully invested in the asset and that reinvests all dividends, x_t is a scalar that drives the risky asset excess return, and y_t is a scalar that drives variation in return volatility. The shocks dB_t and dZ_t are independent three dimensional Brownian motions. An important innovation in our analysis is to allow the mix of volatility to vary, what happens as long as F and D are not proportional to each other.

The state variables evolve as follows,

$$dx_t = \kappa_x(\mu_x - x_t)dt + \sqrt{y_t}GdB_t + HdZ_t \quad (15)$$

$$dy_t = \kappa_y(\mu_y - y_t)dt + \sqrt{y_t}LdB_t. \quad (16)$$

The volatility process follows a CIR process, so it is bounded below by zero. We impose the appropriate conditions to guarantee that the zero boundary is reflexive ($2\kappa_y\mu_y > L'L$). Vectors D, F, G, H, L are three by one constant vectors, κ_x, κ_y are positive scalars that control the rate of mean-reversion of discount rate and volatility shocks, and μ_x, μ_y are the unconditional averages of expected returns and volatility.

The vectors G and H have the the first two rows equal to zero. So only the third shock of each Brownian moves discount rates. The fact that discount rate shocks only

have a transitory effect on asset prices implies $G[3] = -D[3]\kappa_x$ and $H[3] = -F[3]\kappa_x$. The vector L has the first entry equal to zero. The second entry captures pure volatility shocks that are contemporaneously unrelated to discount rate shocks, and the third entry captures the fact that volatility and expected returns might go up at the same time. The vectors D and F have in the first two entries permanent cash-flow shocks. In the first entry we have the exposure to shocks unrelated to volatility or expected returns, and in the second shocks related to volatility. In the third entry we have the discount rate shocks. Shocks that by construction mean-revert in the long run.

Two features of this environment are important for our analysis. First, this specification allows flexibility to fit the weak relationship between expected returns and volatility we see in the data. Expected returns and volatility can potentially co-move by making $|L'G| > 0$ but they don't need to. Second, the specification allow us flexibility to choose both the average mix of discount rate and cash-flow shocks, and how the distribution of these shocks evolve over time. In particular, we consider extreme cases where all *time-variation* in volatility is driven by discount rate volatility.

4.2.2 Investor preferences and optimization

Investors preferences are described by Epstein and Zin (1989) utility. We adopt the Duffie and Epstein (1992) continuous time implementation:

$$J = E_t \left[\int_t^\infty f(C_s, J_s) ds \right] \quad (17)$$

$$f(C, J) = \rho \frac{1 - \gamma}{1 - \psi^{-1}} J \times \left[\left(\frac{C}{((1 - \gamma)J)^{\frac{1}{1-\gamma}}} \right)^{1-\psi^{-1}} - 1 \right] \quad (18)$$

where ρ is rate of time preference, γ the coefficient of relative risk aversion, and ψ is the elasticity of intertemporal substitution. Power utility is the knife edge case where $\gamma = \psi^{-1}$. In the limit $\psi \rightarrow 1$, the aggregator $f(C, J)$ converges to

$$f(C, J) = \rho(1 - \gamma)J \times \left[\log(C) - \frac{\log((1 - \gamma)J)}{1 - \gamma} \right]. \quad (19)$$

Let W_t denote the investor wealth and w_t the allocation to the risky asset, then his budget constraint can be written as,

$$\frac{dW_t}{W_t} = [w_t x_t + r - \frac{C_t}{W_t}]dt + w_t \sqrt{y_t} D dB_t + w_t F dZ_t. \quad (20)$$

The investor maximizes 17 subject to his intertemporal budget constraint and the evolution of state variables x_t and y_t .

4.2.3 Solution

The optimization problem has three state variable. The investor wealth plus the two drivers (expected return x and volatility y) of the investment opportunity set. The Bellman equation for this problem is

$$\begin{aligned} 0 &= \sup_{w,C} f(C_t, J_t) + [w_t x_t W_t + r W_t - C_t] J_W + \frac{1}{2} w_t^2 W_t^2 J_{WW} (y_t D' D + F' F) \\ &+ \kappa_y (\mu_y - y_t) J_y + \frac{1}{2} Y_t J_{yy} L' L + \kappa_x (\mu_x - x_t) J_x + \frac{1}{2} J_{xx} (y_t G' G + H' H) \\ &+ y_t J_{xy} G' L + w_t W_t J_{xW} (G' D y_t + F' H) + w_t W_t J_{yW} L' D Y_t \end{aligned} \quad (21)$$

This problem can be simplified to two state variable by exploring homogeneity of the problem with respect to wealth. In particular, a function of the form $J(W, x, Y) = \frac{W^{1-\gamma}}{1-\gamma} e^{V(x,y)}$ satisfies the above equation. We use collocation methods to solve this problem numerically.

To get intuition about the results that follow it is useful to stare at the optimal portfolio policy,

$$w_t = \frac{x_t}{\gamma (D' D y_t + F' F)} + V_x \frac{G' D y_t + F' H}{\gamma (D' D y_t + F' F)} + V_y \frac{L' D y_t}{\gamma (D' D y_t + F' F)}. \quad (22)$$

The first term is the myopic demand, which reflects the portfolio choice of a short-horizon investor. The two additional terms are the Mertonian hedging demands. Past studies of volatility timing have focused on the third term, which is important to the extent that volatility innovations are strongly correlated with returns. Chacko and Viceira (2005)

shows that for parameters consistent with the data, this term turn out not to be quantitatively important. The return predictability literature has emphasized the second term. The fact that expected returns tend to increase after low returns, make investment in the risky asset a natural investment opportunity set hedge. This effects leads to a higher average position in the the risky asset (V_x is typically negative). Our analysis emphasizes how the strength of this hedging demand fluctuates with volatility.

Assuming for illustration purposes that $L'D = 0$, we can write

$$w_t = \frac{x_t}{\gamma (D'Dy_t + F'F)} - \kappa_x V_x \frac{DR_{share}(y_t)}{\gamma}, \quad (23)$$

where $DR_{share}(y_t) = \frac{G'Dy_t + F'H}{\kappa_x(D'Dy_t + F'F)}$ is the share of return volatility which is driven by discount rate shocks. If most variation in volatility is about discount rate volatility, than $DR_{share}(y_t)$ will be high in periods of high volatility. This increase in the discount rate share will tend to increase the hedging demand in periods of high volatility, counteracting the myopic demand, which calls for a reduction in position. This is the effect John Cochrane alludes to in the previous quote.

Investment horizon plays a role in equation (23) through V_x , the sensitivity of the investor value function with respect to the expected return state variable. A long horizon implies the investor can benefit more from the variation in expected returns, increasing hedging demands accordingly. Intuitively, the risky asset is a better hedge for a long-horizon investor because he or she can wait until the price recovers after a increase in discount rates.

Note that the same force that makes long-horizon investors less responsive to changes in volatility driven by discount rates, will make them more responsive to time-variation in volatility when the driver is cash flow volatility. In this case, $DR_{share}(y_t)$ will be low in periods of high volatility, further reinforcing the myopic demand. In the case the mix of shocks is constant ($F \propto D$ and $G \propto H$), the optimal policy only deviates from the myopic policy to the extent that V_x changes with y_t .

The analysis that follows is focused on understanding how large is V_x and how V_x

changes with the volatility state variable y_t .

4.2.4 Analysis

We study alternative calibrations of the return process that vary in the composition of volatility shocks. Our focus here is to show that in this fairly rich environment the optimal portfolio response of the long-horizon investor with respect to changes in volatility can be approximately described by a strategy that is a constant weight combination of the buy-and-hold portfolio and the volatility managed portfolio.

Table 9 reports the moments used in our calibration. Calibration (a) is the case in which fluctuations in volatility are driven by changes in the volatility of the permanent shock to prices (cash-flow volatility). In calibration (b) we study the case where all time-variation in volatility is the result of variation in the volatility of shocks that have a transitory effect on prices (discount rate volatility). Calibration (c) studies the case of a constant *mix* of transitory and permanent shocks, where changes in volatility do not change the relative importance of permanent and transitory shocks.

We calibrate the stochastic processes to be consistent with the following moments: average volatility, the R-square of a predictability regression of year-ahead returns on the price-dividend ratio, the auto-correlation coefficient of the logarithm of realized variance, the standard deviation of the logarithm of realized variance, and the auto-correlation of the expected return process.

We study CRRA preferences with risk aversion of 5 and 10, and Epstein and Zin utility with with risk aversion of 5 and 10, and IES of 0.5, 1 and 1.5. Following the analysis in Blanchard (1985) and Gârleanu and Panageas (2015), we use the preference parameter ρ as a proxy for investor horizon. We choose ρ so that the half-life⁶ of utility weights ranges from 5 year to 30 years.

⁶Specifically we map horizon into ρ as follows: For a given horizon T_j , look for ρ_j such that $\frac{\int_0^{T_j} e^{-\rho_j \times t} dt}{\int_0^{\infty} e^{-\rho_j \times t} dt} = \frac{1}{2}$

To focus on the effects of volatility variation we initially abstract from subtle hedging demand effects arising from any contemporaneous correlation between return and volatility shocks and set $D[2] = 0$, and assume zero correlation between volatility and expected return shocks $L[3] = 0$. Chacko and Viceira (2005) studies the effect of such hedging demands and show that the effects are small for plausible levels of risk-aversion. We will later revisit these assumptions.

We are interested in understanding how much a long-horizon investor deviate from the optimal mean-variance portfolio. To focus on the effect of volatility we fit the following linear relation on the optimal policy evaluated at the unconditional equity premia $w^*(\mu_x, y_t) = a + b \times \frac{1}{D'Dy_t + FF}$. The policies of a mean-variance investor fit this linear relation perfectly with $a = 0$ and $b = \frac{\mu_x}{\gamma}$. In short, the mean-variance investor puts zero weight on the buy-and-hold portfolio and weight 1 on the volatility managed portfolio. In Tables 7 (a-c) we report $b \times \frac{\gamma}{\mu_x}$, so the reported coefficients have the direct interpretation of how much weight the investor places on the volatility managed portfolio. A coefficient lower than one implies the long-horizon investor trades volatility less aggressively than the mean-variance investor. In the last column we give the minimum R-square for each horizon across specifications in a given row where the row denotes investor horizon.

Table ??(a) shows the case where all time-variation in volatility is driven by cash-flow volatility. The long-horizon investor responds more aggressively than the mean-variance investor (weight on the volatility managed portfolio is higher than one). Intuitively, when cash-flow volatility goes down, the proportion of discount-rate shocks to cash-flow shocks decreases, making stocks relatively safer for a long-horizon investor. Thus, the long-horizon investor responds by increasing his portfolio allocation more than proportionally with the decrease in variance. Importantly this does not speak to their average positions on the risky asset, but only speak to how positions change with volatility.

Table 7(b) shows the case where the average mix is consistent with the data, but all time-variation in volatility is discount rate volatility. For example, both John Cochrane

(2008)⁷ and Warren Buffet(2008)⁸ interpreted the massive increase in volatility in the fall of 2008 as mostly about discount rate volatility and based on that they argued that it was a good time to buy for a long-horizon investor. The linear policy is still an almost perfect description of the optimal policy, but now long-horizon investors respond less aggressively to changes in volatility. Consistent with the intuition in Buffet and Cochrane, long-horizon investors do trade less aggressively on changes in discount rate volatility, but inconsistent with their views a long-horizon investor still trades quite a bit. Coefficients are always quite high, implying that the dynamic investor invests quite a bit on the volatility managed portfolio. For example, during the fall of 2008 volatility spiked from 20% to 60%. This would induce the mean-variance investor to shrink his weight by 90%, while an investor with half-life of 30 years and coefficient of relative risk-aversion 5, would shrink his portfolio by 42%(0.47*90%). This aggressive response to discount rate volatility is at odds with the view expressed by Buffet and Cochrane. Long-horizon investors perceive discount-rate volatility as risky because measured discount-rate variation is extremely persistent in the data. Even though shocks eventually mean-revert, they take a long time.

Table 7(c) shows the case where the average mix of discount-rate and cash-flow volatility is constant. This is a natural benchmark given that we are not aware of any work that have shown how the mix of discount rate and volatility shocks evolve over-time in a predictable way. Again coefficients are basically equal to 1. The mean-variance policy is still a good description of how the long-horizon investor should *change* his portfolio. Obviously the presence of discount rate shocks will induce the long-horizon investor to perceive less risk and hold more stocks on average, but his response to changes in volatility is approximately equal to the mean-variance investor behavior.

In Tables 8 (d) and (e) we study the effect of introducing contemporaneous correlation between return realizations and volatility innovations. Chacko and Viceira (2005) stud-

⁷WSJ, Nov/2008 "Is now the time to buy sotcks?"

⁸NYT, Oct/2008 "Buy American. I am."

ies the hedging demands that result from this correlation. Here our focus is not in level effects on the investor allocation, but whether these correlations change how an investor should respond to volatility shocks. We start from calibration (c) and introduce a correlation between realized returns and volatility shocks. In Table 8 (d) is the case where the correlation between returns and volatility is entirely due to the discount rate component. That is volatility and expected returns shock are positively related in this case. Table 8 (e) analyzes the other extreme where none of the correlation between returns and volatility is due to the discount rate component. In this case the expected return process is again completely independent of the discount rate volatility process. In both cases we see that the introduction of this correlation did not meaningfully change how investors respond to changes in volatility. As is was the case for calibration (c) the long-horizon investor responds roughly in the same fashion as the mean-variance investors

4.2.5 Summary

Overall, this analysis shows that long horizon investors can substantially benefit from volatility timing, and their optimal policy can be described by a simple strategy that combines the buy-and-hold and the volatility managed portfolio.

The distinction between the *type* of volatility that is time-varying plays an important role in determining how much a long-horizon investor weights the volatility managed portfolio. Specifically we distinguish between volatility of cash flow shocks, shocks that have permanent effects on the level of prices, and volatility of discount rate shocks, shocks that have only transitory effects on the price level.

This distinction interacts with the investor horizon in an interesting fashion. In particular, long-horizon investors are better able to deal with discount rate volatility. Cash-flow volatility on the other hand impacts all investors equally because there are permanent shocks to asset prices.

We find that only when all variation in volatility is driven by discount rate volatility,

the long horizon investor responds less aggressively to changes in volatility. But even in this case, the optimal weight on the volatility managed portfolio is positive and substantial for horizons up to 30 years.

5. Implications for 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 we are mindful that theories linking macro and asset prices are typically designed to think about lower frequencies than the studied in this paper. Here our goal is to briefly contrast the empirical pattern we document in the data with these models predictions.

Specifically, our results speak to these theories because their aim is to explain time-variation in risk-premia through a combination of time-variation in risk and the price of risk. Our empirical work allow us to study whether the joint dynamics these models rely on is consistent with the data, at least at the frequencies that we study.

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,

$$\max_i \frac{E_t[R_{i,t+1}^e]}{\sigma_t^2(R_{i,t+1}^e)} \approx f(s_t) \quad (24)$$

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.⁹ In both these economies, negative

⁹In CC positive shocks to consumption increase the distance of the agent to it's habit level, reducing the

shocks to consumption makes marginal utility 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 paper with this feature is Barberis et al. (2001) who have effective risk aversion occurring after realized losses.

While it is undisputed that the data features a lot of excess volatility, these models seem to 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.

Wachter (2013) and Bansal and Yaron (2004) have a different flavor as they rely on time variation in cash-flow risk to generate fluctuations in risk premia. 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 about persistent movements in fundamental volatility. This framework 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 very close to 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 counter-factual given the empirical dynamics of aggregate consumption.

Overall, our results seem to be consistent with a model where the bulk of time-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.

variation in volatility is discount rate volatility along the lines of (Campbell and Cochrane, 1999) and (He and Krishnamurthy, 2012), but where discount rate volatility is not as tightly related to the level of risk-premia as in these models.

6. Conclusion

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. 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.
- 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., 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.
- Blanchard, O. J. (1985). Debt, deficits, and finite horizons. *The Journal of Political Economy*, 93(2):223–247.
- Buffett, W. E. (2008). Buy American. I am. *The New York Times*, October 16th, 2008.
- 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.
- Duffie, D. and Epstein, L. G. (1992). Stochastic differential utility. *Econometrica: Journal of the Econometric Society*, pages 353–394.
- Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica: Journal of the Econometric Society*, pages 937–969.
- Fama, E. F. and French, K. R. (1996). Multifactor explanations of asset pricing anomalies. *Journal of Finance*, 51(1):55–84.

- Fleming, J., Kirby, C., and Ostdiek, B. (2001). The economic value of volatility timing. *The Journal of Finance*, 56(1):329–352.
- Fleming, J., Kirby, C., and Ostdiek, B. (2003). The economic value of volatility timing using ?realized? volatility. *Journal of Financial Economics*, 67(3):473–509.
- Frazzini, A., Israel, R., and Moskowitz, T. (2015). Trading costs of asset pricing anomalies. *working paper*.
- 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.
- Gârleanu, N. and Panageas, S. (2015). Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy*, 123(3):670–685.
- 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.
- Goyal, A. and Welch, I. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508.
- 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.

- 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.
- Liu, J. (2007). Portfolio selection in stochastic environments. *Review of Financial Studies*, 20(1):1–39.
- 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.
- Merton, R. C. (1973). An intertemporal asset pricing model. *Econometrica*, 41(5):867–887.
- Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1):1–28.
- Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *The review of economics and statistics*, pages 239–246.
- Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, forthcoming.

7. Tables and Figures

Table 1: 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.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	SMkt	SSMB	SHML	SMom	SRMW	SCMA	SMVE	SFX	SROE	SIA
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)		
ROE									0.63 (0.07)	
IA										0.68 (0.05)
α	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)	5.48 (0.97)	1.55 (0.67)
N	1,065	1,065	1,065	1,060	621	621	1,060	360	575	575
R2	0.37	0.38	0.32	0.22	0.38	0.46	0.33	0.51	0.40	0.47
rmse	51.39	30.44	34.92	50.37	20.16	17.55	25.34	21.78	23.69	16.58

Table 2: 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 3: 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 4: 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.

	(1)	(2)	(3)	(4)	(5)
VARIABLES	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 5: 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

Table 6: We regress returns at time $t + 1$ for each factor on the realized volatility of the factor at time t . The regressions give us a sense for the risk return relation across factors. 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.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Mkt	SMB	HML	Mom	RMW	CMA	MVE	FX	ROE	IA
$\sigma_{Mkt,t}$	-0.04 (0.44)									
$\sigma_{SMB,t}$		0.38 (0.38)								
$\sigma_{HML,t}$			0.70 (0.58)							
$\sigma_{Mom,t}$				-1.10 (0.59)						
$\sigma_{RMW,t}$					0.87 (0.83)					
$\sigma_{CMA,t}$						1.19 (0.62)				
$\sigma_{MVE,t}$							-0.02 (0.35)			
$\sigma_{FX,t}$								-0.10 (0.67)		
$\sigma_{ROE,t}$									-0.90 (0.76)	
$\sigma_{IA,t}$										0.87 (0.75)
N	1,065	1,065	1,065	1,060	621	621	1,060	360	575	575
R^2	0.00	0.00	0.01	0.02	0.01	0.02	0.00	0.00	0.01	0.01

Table 7: Panels (a)-(c) show linear approximations to the optimal portfolio policy for calibrations (a)-(c). In calibration (a) variation in volatility is driven by cash-flow volatility, in calibration (b) it is driven by cash-flow volatility, and calibration (c) is the case where the volatility mix is constant. Specifically, the panels report $b \times \frac{\gamma}{\mu_x}$, where b is the slope of a regression of the optimal policy on the inverse variance. The last column reports the minimum R-squares across these specifications. The calibrations are in Table 9.

(a) Discount rate volatility									
Horizon	Power utility (γ)		Epstein and Zin (γ, ψ)						R^2
	(5)	(10)	(5,0.5)	(10,0.5)	(5,1)	(10,1)	(5,1.5)	(10,1.5)	
30	0.68	0.48	0.60	0.41	0.44	0.33	0.35	0.28	0.99
20	0.68	0.48	0.61	0.43	0.47	0.37	0.38	0.33	0.99
10	0.69	0.49	0.64	0.47	0.55	0.45	0.47	0.43	0.99
5	0.71	0.51	0.68	0.54	0.63	0.55	0.59	0.55	0.97
0	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.00
(b) Cash flow volatility									
Horizon	Power utility (γ)		Epstein and Zin (γ, ψ)						R^2
	(5)	(10)	(5,0.5)	(10,0.5)	(5,1)	(10,1)	(5,1.5)	(10,1.5)	
30	1.69	2.20	1.85	2.27	2.14	2.36	2.34	2.44	1.00
20	1.67	2.18	1.79	2.18	2.01	2.20	2.19	2.24	1.00
10	1.60	2.12	1.63	1.95	1.73	1.86	1.83	1.82	1.00
5	1.51	2.01	1.45	1.67	1.43	1.50	1.44	1.41	1.00
0	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.00
(c) Constant volatility mix									
Horizon	Power utility (γ)		Epstein and Zin (γ, ψ)						R^2
	(5)	(10)	(5,0.5)	(10,0.5)	(5,1)	(10,1)	(5,1.5)	(10,1.5)	
30	0.99	0.98	0.98	0.99	1.00	1.01	1.01	1.02	1.00
20	0.99	0.98	0.98	0.99	1.00	1.00	1.01	1.02	1.00
10	0.99	0.98	0.98	0.98	0.99	0.99	1.01	1.00	1.00
5	0.99	0.98	0.98	0.98	0.98	0.98	0.99	0.98	1.00
0	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.00

Table 8: Panels (d) and (e) show linear approximations to the optimal portfolio policy in the cases of constant mix of discount rate and cash-flow volatility. Panel (d) introduces a positive correlation between expected return and volatility shocks and Panel(e) introduces a correlation between cash-flow and volatility shocks. Specifically, the panels report $b \times \frac{\gamma}{\mu_x}$, where b is the slope of a regression of the optimal policy on the inverse variance process. The last column reports the minimum R-squares across these specifications. The calibrations are in Table 9.

(d)Volatility shocks positively correlated to discount rate shock									
Horizon	Power utility (γ)		Epstein and Zin (γ, ψ)						R^2
	(5)	(10)	(5,0.5)	(10,0.5)	(5,1)	(10,1)	(5,1.5)	(10,1.5)	
30	0.88	0.73	0.80	1.00	0.96	0.95	1.01	0.98	0.56
20	0.88	0.73	0.80	1.00	0.95	0.95	1.01	0.98	0.56
10	0.88	0.74	0.82	1.02	0.94	0.94	1.01	0.98	0.60
5	0.89	0.75	0.84	1.04	0.94	0.93	1.01	0.99	0.70
0	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.00
(e)Volatility shocks negatively correlated to cash-flow shock									
Horizon	Power utility (γ)		Epstein and Zin (γ, ψ)						R^2
	(5)	(10)	(5,0.5)	(10,0.5)	(5,1)	(10,1)	(5,1.5)	(10,1.5)	
30	0.92	0.93	0.93	0.66	0.95	0.95	0.93	0.92	0.90
20	0.92	0.93	0.93	0.67	0.95	0.95	0.93	0.94	0.90
10	0.92	0.93	0.92	0.69	0.95	0.95	0.94	0.95	0.91
5	0.92	0.92	0.91	0.72	0.93	0.95	0.93	0.96	0.92
0	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.00

Table 9: This Table show targeted moments in the data and in the eight alternative model analyzed in Tables 7 and 8. Model based quantities are found by simulating 100 different 100 years histories. Reported is the median moment realization across these 100 histories.

Moment	Data	(a)	(b)	(c)	(d)	(e)
$\sqrt{E[Rvar(R_{t+1})]} * 12$	0.160	0.168	0.166	0.165	0.163	0.164
$corr(pd_t, R_{t \rightarrow t+12})$	0.060	0.000	0.056	0.072	0.067	0.062
$corr(\log(Rvar(R_{t+1})), \log(Rvar(R_t)))$	0.500	0.519	0.469	0.506	0.448	0.504
$corr(R_{t \rightarrow t+12}, 1/Rvar(R_t))$	0.000	0.016	0.016	0.016	0.013	0.015
$stdev(\log(Rvar(R_{t+1})))$	1.050	1.039	0.796	0.920	0.751	0.924
$corr(R_t, Rvar(R_t) - Rvar(R_{t-1}))$	-0.240	-0.000	-0.000	-0.001	-0.200	-0.177
$corr(pd_{t+1}, pd_t)^2$	0.980	1.000	0.975	0.975	0.975	0.975

Table 10: We evaluate our volatility timing strategy for the market portfolio when including transactions costs. We consider alternative strategies that capture volatility timing but reduce trading activity, including using volatility instead of variance, using expected rather than realized variance, and only trading when variance is above the long run average. We report the average absolute change in monthly weights, expected return, and alpha of each strategy. Then we report the alpha when including various trading costs. The 1bps cost comes from Fleming et al. (2003), the 10bps comes from Frazzini et al. (2015) when trading about 1% of daily volume, and the last column adds an additional 4bps to account for transaction costs increasing in high volatility episodes. Specifically, we use the slope coefficient of transactions costs on VIX from Frazzini et al. (2015) and evaluate this impact on a move in VIX from 20% to 40% which represents the 98th percentile of VIX. Finally, the last column backs out the implied trading costs in basis points needed to drive our alphas to zero in each of the cases.

w	Description	$ \Delta w $	$E[R]$	α	α After Trading Costs			
					1bps	10bps	14bps	Break Even
$\frac{1}{RV^2}$	Realized Variance	0.73	9.47%	4.86%	4.77%	3.98%	3.63%	56bps
$\frac{1}{RV}$	Realized Vol	0.38	9.84%	3.85%	3.80%	3.39%	3.21%	84bps
$\frac{1}{E_t[RV^2]}$	Expected Variance	0.37	9.47%	3.30%	3.26%	2.86%	2.68%	74bps
$\frac{1}{\max(E[RV^2], RV^2)}$	RV Above Mean	0.10	9.10%	2.20%	2.19%	2.08%	2.03%	183bps

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.

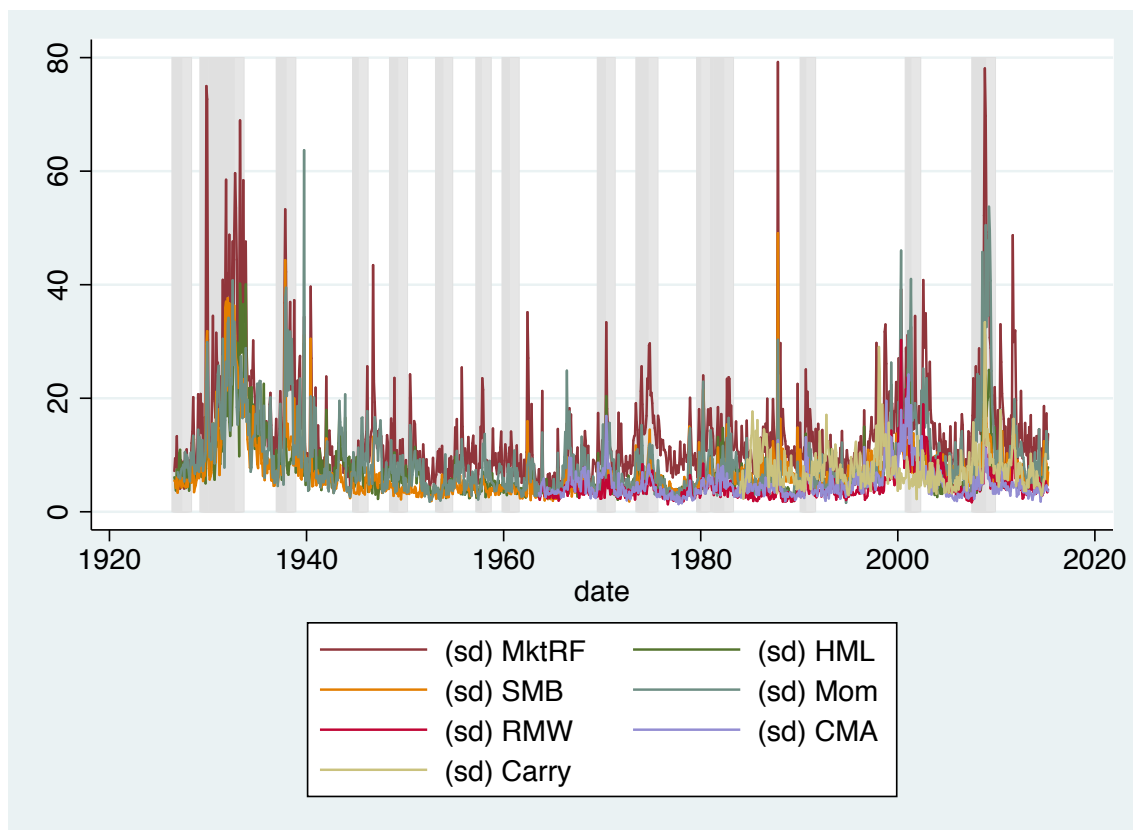


Figure 2: The figure plots the full distribution of scaled factors vs non-scaled factors estimated using kernel density. In particular, for each panel we plot the distribution of f_{t+1} (solid line) along with the distribution of $z_t f_{t+1}$ (dashed line).

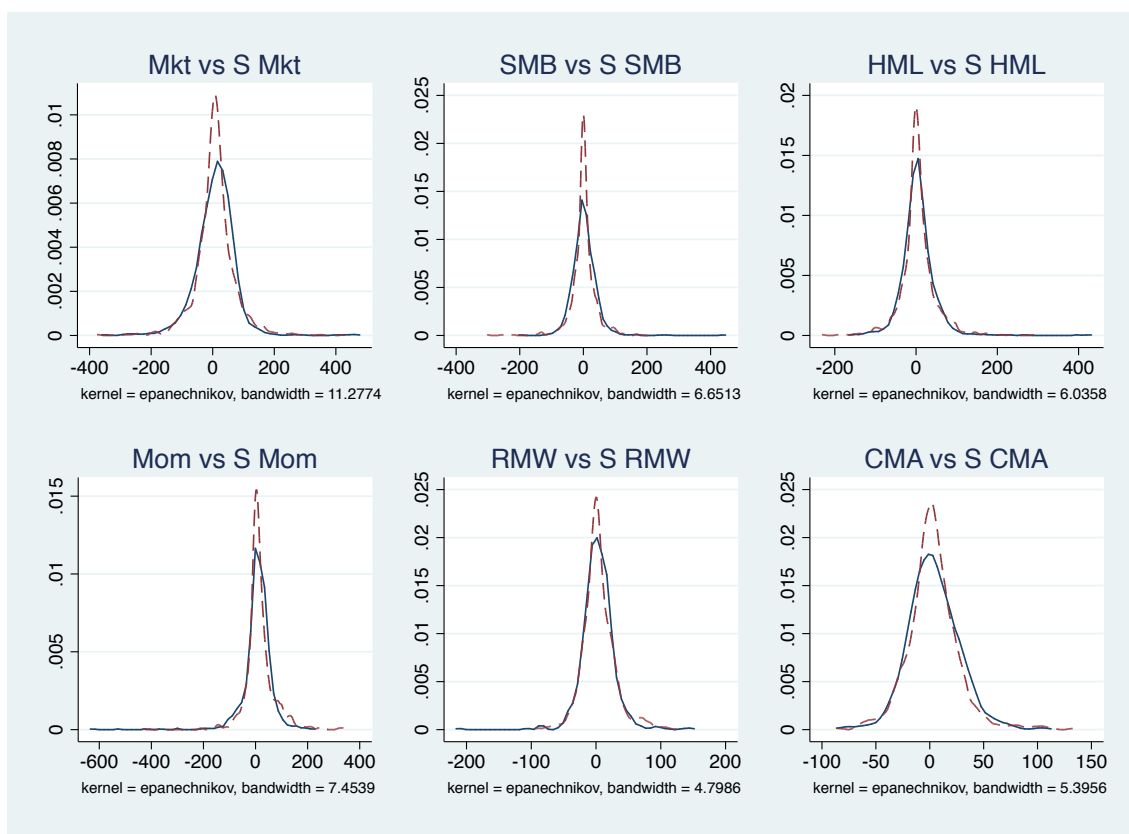


Figure 3: The figure plots the impulse response of the expected return and variance of the market portfolio for a shock to the variance. The x-axis is months. The last panel gives the portfolio choice implications for the differential movements in expected returns and variance to a variance shock at 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.

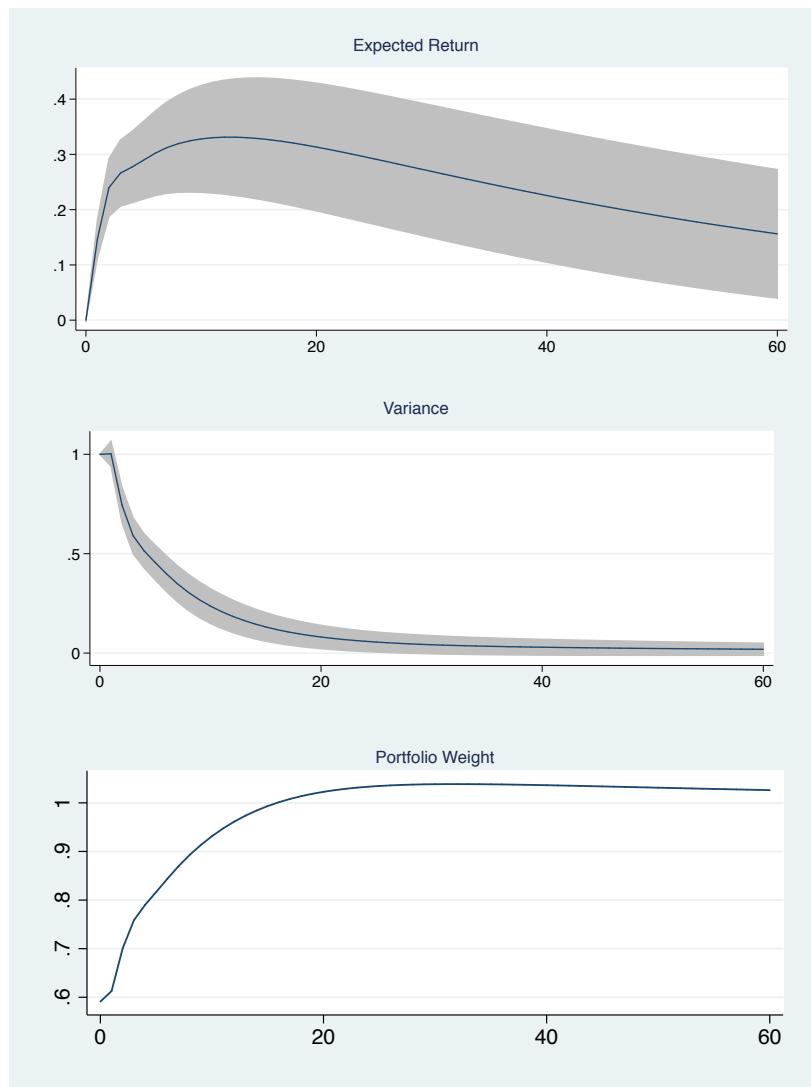


Figure 4: The figure plots the cumulative returns to a buy and hold strategy vs a volatility timing strategy for the market portfolio from 1926-2015. The y-axis is on a log scale and both strategies have the same unconditional monthly standard deviation.

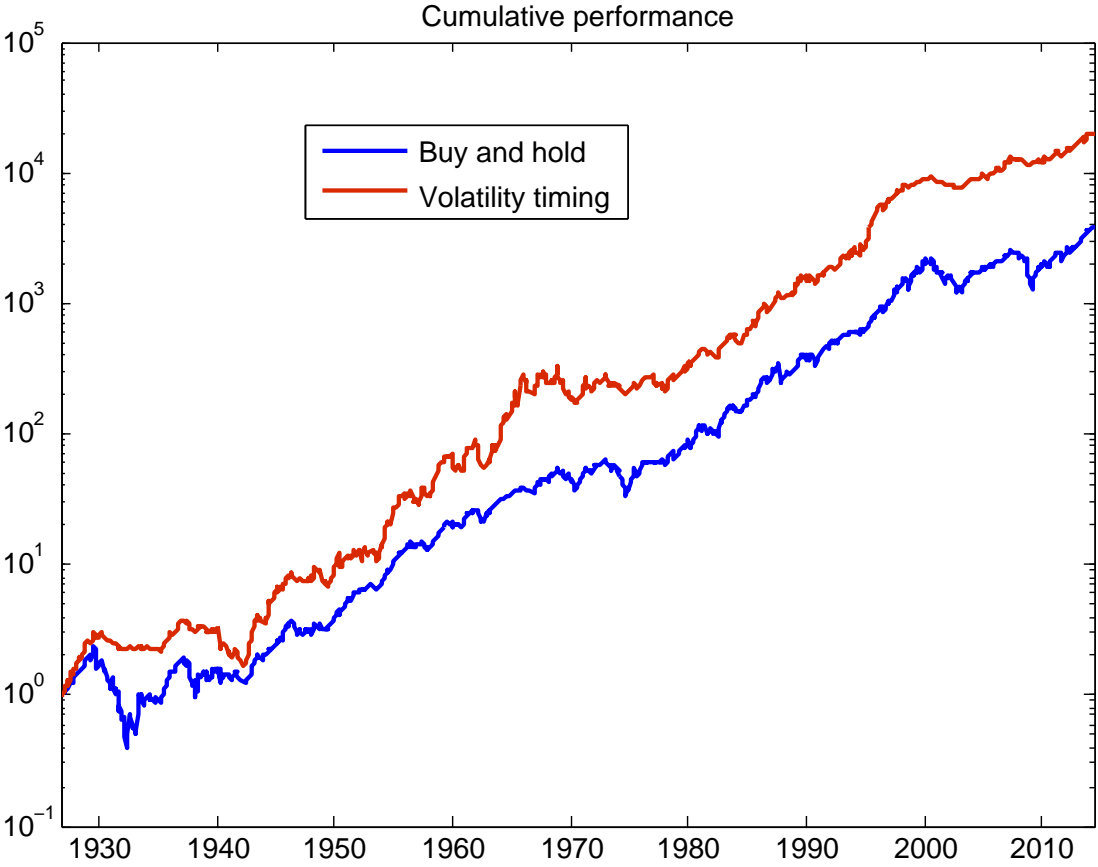


Figure 5: The figure plots the performance of our volatility timing strategy compared to a buy and hold strategy for the market return during times of market turmoil.

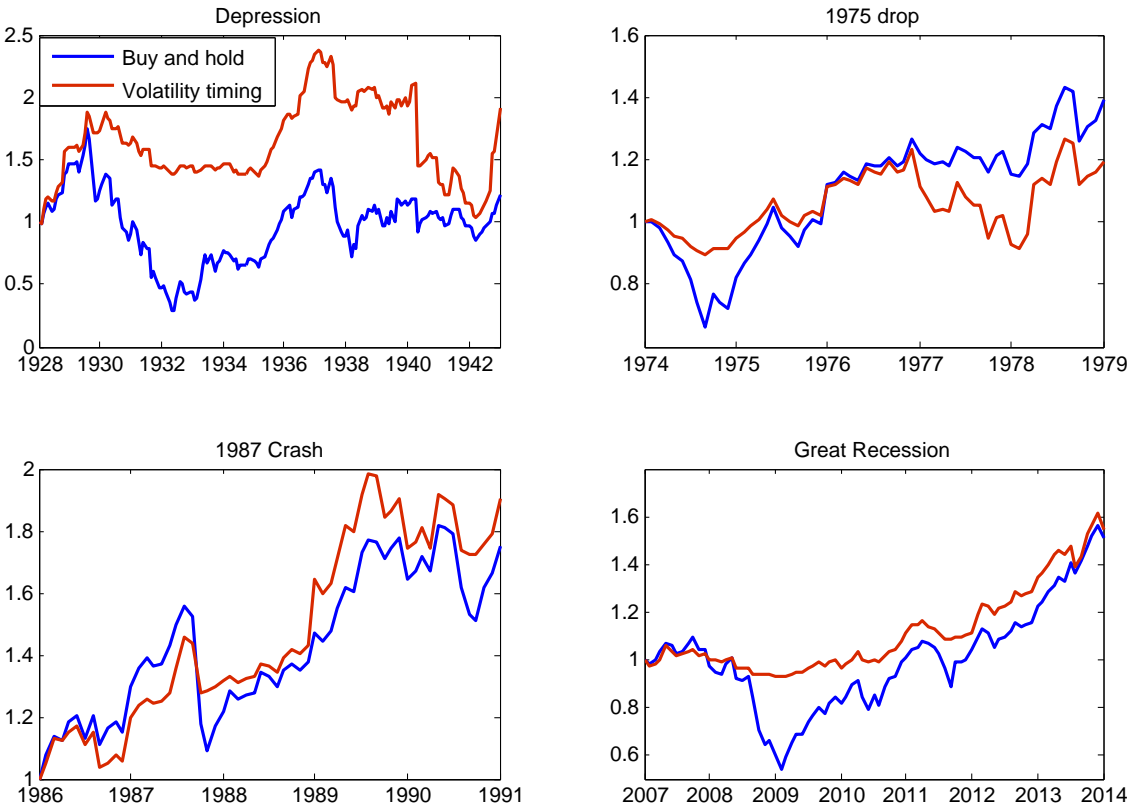
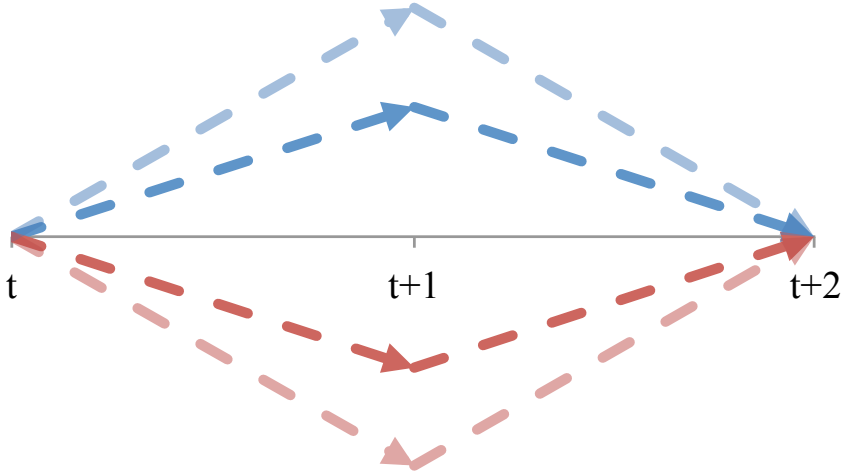


Figure 6: The figure plots a hypothetical description of a mean reverting volatility shock. For illustration, we consider a shock to returns that is fully mean reverting the next period. Thus, a negative return next period is perfectly offset by a positive return of equal magnitude of next period, and vice versa. Regardless of the volatility, an investor with a two period horizon does not view stocks as risky. However, an investor with a one period horizon will view the increase in volatility as an increase in risk.



A. Derivations

A.1 Deriving Alpha

We show conditions under which managed portfolios that scale factors by $1/\sigma_t^2$ generate risk-adjusted returns relative to original factors. Intuitively, if the price of risk, b_t , is low during periods of high volatility, 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

Proposition 1 (Managed portfolio alpha). *Let α be the intercept of the following regression*

$$\frac{1}{\sigma_{f,t}^2} f_{t+1} = \alpha + \beta f_{t+1} + \varepsilon_{t+1}, \quad (25)$$

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

$$\alpha = E \left[cov_t \left(\frac{1}{\sigma_t^2} f_{t+1}, f_{t+1} \right) b_t \right] - cov \left(\frac{1}{\sigma_t^2} f_{t+1}, f_{t+1} \right) b \quad (26)$$

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,

$$b_t \sigma_t^2 \approx (b_t - E[b_t])E[\sigma_t^2] + (\sigma_t^2 - E[\sigma_t^2])E[b_t], \quad (27)$$

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

$$\alpha = -cov(b_t, \sigma_t^2)E[b_t]b - \left(b - E[b_t] + bE[\sigma_t^2]var(b_t) \right). \quad (28)$$

□

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. In later sections we demon-

strate that indeed the covariance is negative. This stems from the simple fact that volatility does not strongly forecast returns. The next section derives this covariance term more carefully.

A.2 Derivation of Equation Covariance

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

$$\mu_t = b_t \sigma_t^2 \quad (29)$$

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

$$\beta = \frac{cov(f_{t+1}, \frac{-1}{\sigma_t^2})}{var(\frac{1}{\sigma_t^2})} \quad (30)$$

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

$$\beta = \frac{cov(b_t \sigma_t^2, \frac{1}{\sigma_t^2})}{var(\frac{1}{\sigma_t^2})} \quad (31)$$

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

$$cov(xy, z) = cov(x, y(z - E[z])) + E[x]cov(y, z) \quad (32)$$

This means

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

But then

$$cov(\sigma_t^2, 1/\sigma_t^2) = 1 - E[\sigma_t^2]E[1/\sigma_t^2] \quad (34)$$

So that

$$\beta = \left[\text{cov}(b_t, \sigma_t^2) / E[\sigma_t^2] + \frac{\mu}{E[\sigma_t^2]} (E[\sigma_t^2] E[1/\sigma_t^2] - 1) \right] / \text{var}(1/\sigma_t^2) \quad (35)$$

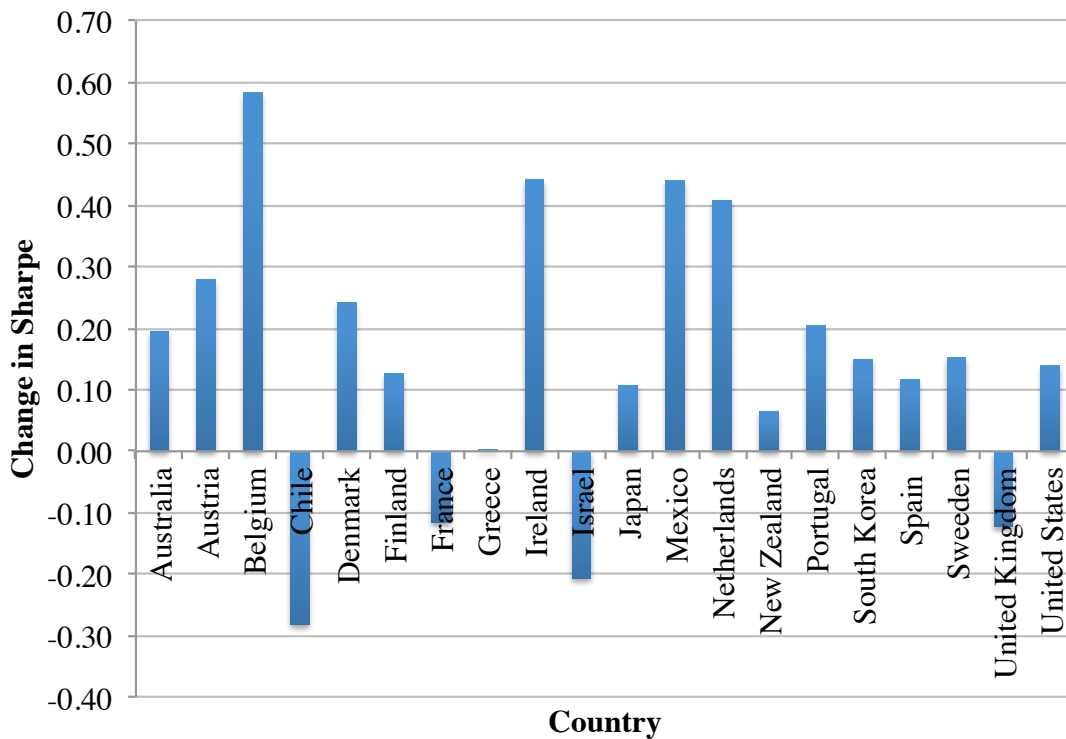
Rearranging,

$$\frac{1}{E[\sigma_t^2] \text{var}(1/\sigma_t^2)} \text{cov}(b_t, \sigma_t^2) = \beta - \frac{\mu}{E[\sigma_t^2]} \frac{(E[\sigma_t^2] E[1/\sigma_t^2] - 1)}{\text{var}(1/\sigma_t^2)} \quad (36)$$

We can thus estimate this co-variance. First, we could run the regression of returns on negative inverse variance then compute all other quantities. We find the co-variance to be strongly negative for almost all factors. This means the alpha shown in the previous section is typically going to be positive.

B. Additional Figures

Figure 7: The figure plots the change in Sharpe ratio for managed vs non-managed portfolios across 20 OECD countries. The change is computed as the Sharpe ratio of the volatility managed index minus the Sharpe ratio of the buy and hold country index. All indices are from Global Financial Data. For many series, the index only contains daily price data and not dividend data, thus our results are not intended to accurately capture the level of Sharpe ratios but should still capture their difference well to the extent that most of the fluctuations in monthly volatility is driven by daily price changes. All indices are converted to USD and are taking over the US risk-free rate from Ken French. The average change in Sharpe ratio is 0.15 and the value is positive in 80% of cases.



C. Additional Tables

Table 11: 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

Table 12: 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 13: 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 (??) 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 14: We repeat the previous tables estimates at both a *quarterly* as well as *annual* horizon, rather than a monthly horizon. We estimate the covariance between the price and quantity of risk for each factor and for the mean variance efficient portfolio (MVE). 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 (??) in the text. The covariances are given in percentages.

	$cov(b_t, \sigma_t^2)$			
	Quarterly		Annually	
	Mean	Std Err	Mean	Std Err
MktRF	-0.42	0.17	-0.44	0.16
SMB	-0.03	0.14	-0.04	0.14
HML	-0.26	0.16	-0.10	0.13
Mom	-1.37	0.23	-1.13	0.27
MVE	-0.55	0.09	-0.38	0.08
RMW	-0.15	0.08	-0.12	0.08
CMA	-0.09	0.07	0.03	0.05
MVE2	-0.45	0.10	-0.28	0.06
Carry	-0.11	0.07	0.01	0.04