

NBER WORKING PAPER SERIES

VOLATILITY MANAGED PORTFOLIOS

Alan Moreira
Tyler Muir

Working Paper 22208
<http://www.nber.org/papers/w22208>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2016

We thank Matthew Baron, Jonathan Berk, Olivier Boguth, John Campbell, John Cochrane, Kent Daniel, Peter DeMarzo, Wayne Ferson, Marcelo Fernandes, William Goetzmann, Mark Grinblatt, Ben Hebert, Steve Heston, Jon Ingersoll, Ravi Jagannathan, Bryan Kelly, Serhiy Kosak, Hanno Lustig, Justin Murfin, Stefan Nagel, David Ng, Lubos Pastor, Myron Scholes, Ivan Shaliastovich, Ken Singleton, Tuomo Vuoltenahoo, Jonathan Wallen, Lu Zhang and participants at Yale SOM, UCLA Anderson, Stanford GSB, Michigan Ross, Chicago Booth, Ohio State, Baruch College, Cornell, the NYU Five Star conference, the Colorado Winter Finance Conference, the Jackson Hole Winter Finance Conference, the ASU Sonoran Conference, the UBC winter conference, and Arrowstreet Capital for comments. We especially thank Nick Barberis for many useful discussions. We also thank Ken French for making data publicly available and Alexi Savov, Adrien Verdelhan and Lu Zhang for providing data. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Alan Moreira and Tyler Muir. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Volatility Managed Portfolios
Alan Moreira and Tyler Muir
NBER Working Paper No. 22208
April 2016
JEL No. G0,G12

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%. Contrary to conventional wisdom, we show that long horizon investors can benefit from volatility timing even when time-variation in volatility is completely driven by discount rate volatility.

Alan Moreira
Yale School of Management
Yale University
165 Whitney, Room 4520
New Haven, CT 06520
alan.moreira@yale.edu

Tyler Muir
Yale School of Management
165 Whitney Avenue, Room 4516
New Haven, CT 06511
and NBER
tyler.muir@yale.edu

1. Introduction

We construct portfolios that scale monthly returns by the inverse of their previous month's realized variance, decreasing risk exposure when variance was recently 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. We then interpret these results from both a portfolio choice and a general equilibrium perspective.

We motivate our analysis from the vantage point of a mean-variance investor, who adjusts their allocation 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 short horizons, 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. Annualized alphas with respect to the original factors are substantial, and Sharpe ratios increase by 50% to 100% of the original factor Sharpe ratios.

Figure 1 provides intuition for our results for the market portfolio. In line with our trading strategy, we group months by the previous month's realized volatility and plot average returns, volatility, and the mean-variance trade-off over the subsequent month. There is little relation between lagged volatility and average returns but there is a strong relationship between lagged volatility and current volatility. This means that the mean-variance trade-off weakens in periods of high volatility. From a portfolio choice perspective, this pattern implies that a short-horizon mean-variance investor should time volatility, i.e. take more risk when the mean-variance trade-off is attractive (volatility is low), and take less risk when the mean-variance trade-off is unattractive (volatility is high). From a general equilibrium perspective, this pattern presents a challenge to models focused on the dynamics of risk premia. The empirical pattern in Figure 1 implies that investor's willingness to take stock market risk would have to be higher in periods of high stock market volatility, which is counter to most theories. Sharpening the puzzle is the fact that volatility is typically high during recessions, financial crises, and in the aftermath of market crashes when theory generally suggests investors should, if anything, be

more risk averse relative to normal times.

Our volatility managed portfolios reduce risk taking during these bad times– times when the common advice is to increase or hold risk taking constant.¹ For example, in the aftermath of the sharp price declines in the fall of 2008, it was a widely held view that those that reduced positions in equities were missing a once-in-a-generation buying opportunity.² Yet our strategy 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 throughout several crisis episodes, including the Great Depression, the Great Recession, and 1987 stock market crash. More broadly, we show that our volatility managed portfolios take substantially less risk during recessions.

These facts may be surprising in light of evidence showing that expected returns are high in recessions (Fama and French, 1989) and in the aftermath of market crashes (Muir, 2013). 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 such as the price-to-earnings ratio and the yield spread between Baa and Aaa rated bonds. Using a vector autoregression (VAR) we show that in response to a variance shock, the conditional variance initially increases by far more than the expected return, making the risk-return trade-off initially unattractive. A mean-variance investor would decrease his or her risk exposure by 60% after a one 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. This difference in persistence reconciles the evidence on countercyclical expected returns with the profitability of our strategy.

We go through an extensive list of exercises to evaluate the robustness of our result. We show that our volatility managed strategy survives transaction costs, is different from

¹For example, in August 2015, a period of high volatility, Vanguard—a leading mutual fund company—gave advice consistent with this view :“What to do during market volatility? Perhaps nothing.” See <https://personal.vanguard.com/us/insights/article/market-volatility-082015>

²See for example Cochrane (2008) and Buffett (2008) for this view. However, in line with our main findings, Nagel et al. (2016) find that many households respond to volatility by selling stocks in 2008 and that this effect is larger for higher income households who may be more sophisticated traders.

strategies that explore low risk anomalies in the cross-section such as risk parity (Asness et al., 2012) and betting against beta (Frazzini and Pedersen, 2014), is less exposed to volatility shocks than the original factors (ruling out explanations based on the variance risk premium), and cannot be explained by downside market risk (Ang et al., 2006a; Lettau et al., 2014), disaster risk or jump risk. We also show that our strategy works just as well if implemented through options, which suggests that leverage constraints are unlikely to explain the high alphas of our volatility managed strategies. In the Appendix we show that our strategy works across 20 OECD stock market indices, that it can be further improved through the use of more sophisticated models of variance forecasting, that it does not generate fatter left tails than the original factors or create option-like payoffs, and that it outperforms not only using alpha and Sharpe ratios but also manipulation proof measures of performance (Goetzmann et al., 2007).

To provide more economic content to these findings we start by studying their portfolio choice implications. We first take the perspective of a short-horizon mean-variance investor. Our findings imply that they should reduce their exposure when volatility increases because future volatility will be high but expected returns won't be. We show that utility benefits of volatility timing are large, on the order of 50% to 90% of lifetime utility. These are substantially larger than those coming from expected return timing. However, the empirical evidence that there is some mean-reversion in returns, i.e. there are both permanent and mean-reverting shocks to stock prices, implies that volatility is not the right measure of risk for a long-horizon investor.³ Thus, investors with long horizons might not benefit from volatility timing. Both Cochrane (2008) and Buffett (2008) articulated this view in the fall of 2008 as volatility spiked to extreme levels. The idea is that an increase in the volatility of mean-reverting shocks implies that an investor has a greater chance of waking up poorer tomorrow, but expects to be just as rich in the long run. Thus, we dig deeper into the portfolio problem of a long-horizon investor in order to see how they should respond to volatility once mean-reversion is taken into account.

We find that mean-reversion can either increase or decrease the benefits of volatility timing for a long-horizon investor. A long-horizon investor should time volatility less ag-

³Mean-reverting shocks in stock returns are often called discount rate shocks.

gressively than a short-horizon investor only if changes in volatility are associated with greater-than-proportional changes in the amount of mean-reversion in returns, i.e. the share of return volatility due to mean-reverting shocks increases with volatility. Intuitively, an increase in the volatility of mean-reverting shocks makes stocks relatively safer in the long run compared to the short run. Following the same logic, if changes in volatility are associated with less-than-proportional increases in the amount of mean-reversion, the long-horizon investor should time volatility even more aggressively than the short-horizon investor.

Empirically, there is yet no evidence on how volatility and the amount of mean-reversion co-move. Furthermore, even in the case of extreme co-movement, when volatility variation is completely driven by variation in the volatility of mean-reverting shocks, we show that long-horizon investors still find it optimal to time volatility. The key for this result is that in the data mean reversion takes many years, making even mean-reverting shocks risky for realistic investment horizons.⁴ Overall, our analysis conclusively shows that long-horizon investors should also time volatility.

Lastly, we study the general equilibrium implications of our results. In equilibrium, it is useful to think of “effective risk aversion” as $\gamma_t = E_t[R_{t+1}]/Var_t[R_{t+1}]$. Our results imply that γ_t is negatively related to volatility. For example, in Figure 1 we find that effective risk aversion declines by nearly an order of magnitude as we go from low to high volatility quintiles, despite the fact that volatility is strongly related to the business cycle. Equilibrium asset pricing theories all feature the opposite prediction, namely that the correlation between effective risk aversion, γ_t , and variance is weakly positive. This is because in bad times when volatility increases, effective risk aversion in these models also increases, driving up the compensation for risk. This is a typical feature of standard rational, behavioral, and intermediary models of asset pricing alike. We argue that this correlation is important for these models. Ultimately, the goal of these theories is to generate a large and volatile equity premium, and the co-movement in the price and quantity of risk plays a key role in achieving this result.

⁴Sharpe ratios for stocks increase only slowly with investment horizon (Poterba and Summers, 1988), and valuation ratios that predict returns are highly persistent with auto-correlation close to one (Campbell and Shiller, 1988).

The general equilibrium results also highlight how our approach differs from other asset allocation papers which use volatility because our results can speak to the evolution of the aggregate risk return tradeoff. For example, Fleming et al. (2001) and Fleming et al. (2003) study daily asset allocation across stocks, bonds, and gold based on estimating the conditional covariance matrix which performs cross-sectional asset allocation. Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2015) study volatility timing related to momentum crashes.⁵ Instead, our approach focuses on the time-series of many aggregate priced factors allowing us to give economic content to the returns on the volatility managed strategies.

This paper proceeds as follows. Section 2 documents our main empirical results. Section 3 explains the profitability of our strategy in more detail and provides various robustness checks. Section 4 explores portfolio choice implications for short- and long-horizon investors. Section 5 discusses implications for structural asset-pricing models. Section 6 concludes.

2. Main Empirical Results

2.1 Data Description

We use both daily and monthly factors from Ken French's website 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 an investment factor that they use in their 5 factor model (Fama and French (2015), Novy-Marx (2013)). Mom represents the momentum factor which goes long past winners and short past losers. We also include daily and monthly data from Hou et al. (2014) which includes an investment factor, IA, and a return on equity factor, ROE. We also use data on currency returns from Lustig et al. (2011) provided by Adrien Verdelhan. 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

⁵Daniel et al. (2015) also look at a related strategy to ours for currencies.

for the high and low portfolios to construct our volatility measure. We refer to this factor as “Carry” or “FX” to save on notation and to emphasize that it is a carry factor formed in foreign exchange markets.

We compute realized volatility (RV) for a given month t for a given factor f by taking the square root of 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 forming portfolios for the next month $t + 1$. Our approach is simple and uses only return data. Figure 2 displays our monthly estimates for realized volatility for each factor.

2.2 Portfolio construction

We construct managed portfolios by scaling each factor by the inverse of its realized 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

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

The constant c controls average risk exposure of the strategy. For ease of interpretation, we choose c so that the managed factor has the same unconditional standard deviation as the non-managed factor.⁶

The idea is that if variance does not forecast returns, the risk-return trade-off will deteriorate when variance increases. In fact, this strategy is exactly what a mean-variance optimizing agent should do if 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 model. An appealing feature of this approach is that can be easily implemented by an investor in real time. Appendix A.1 considers the use of more sophisticated forecasting models.

⁶Importantly c has no effect on our strategy’s Sharpe ratio, thus the fact that we use the full sample to compute c does not impact our results.

2.3 Main results

Table 1 reports the results from running a regression of the volatility managed portfolios on the original factors. A positive intercept in this time-series regression implies that the volatility managed portfolios expand the unconditional mean-variance frontier and therefore that an investor could increase his or her Sharpe ratio relative to a buy-and-hold strategy. In a world where the trade-off between risk and return is constant, this time-series alpha should be approximately zero (see Section 3.1). Intuitively, this is because when the risk return trade off is strong there is no scope for volatility timing. Conversely, if there is no relationship between risk and return in the time-series, the volatility managed portfolios should earn large positive alphas. Intuitively, the managed portfolio takes advantage of the more attractive compensation for risk during low volatility times and avoids the poor risk-return trade-off during high volatility times.

We see positive, statistically significant constants (α 's) in most cases in Table 1. The managed market portfolio on its own deserves special attention because this strategy would have been easily available to the average investor in real time; moreover the results in this case directly relate to a long literature on market timing that we discuss later.⁷ The scaled market factor has an annualized alpha of 4.86% and a beta of only 0.6. While most alphas are strongly positive, the largest is for the momentum factor.⁸

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 (or "appraisal ratio") given by $\frac{\alpha}{\sigma_\epsilon}$, 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 Sharpe ratio given by the original non-scaled factor. For example, in Table 1, scaled momentum has an α of 12.5 and a root mean square error around 50 which means that its annualized appraisal ratio is $\sqrt{12} \frac{12.5}{50} = 0.875$. The scaled markets' annualized appraisal ratio is 0.34.⁹ Other

⁷The typical investor will likely find it difficult to trade the momentum factor, for example.

⁸This is consistent with Barroso and Santa-Clara (2015) who find that strategies which avoid large momentum crashes by timing momentum volatility perform exceptionally well.

⁹We 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

notable appraisal ratios across factors are: HML (0.20), Profitability (0.41), Carry (0.44), ROE (0.80), and Investment (0.32).

Figure 3 plots the cumulative nominal returns to the volatility managed market factor compared to a buy-and-hold strategy from 1926-2015. We invest \$1 in 1926 and plot the cumulative returns to each strategy on a log scale. From this figure, we can see relatively steady gains from the volatility managed factor, which cumulates to around \$20,000 at the end of the sample vs. about \$4,000 for the buy-and-hold strategy. The lower panels of Figure 3 plot the drawdown and annual returns of the strategy relative to the market, which helps us understand when our strategy loses money relative to the buy-and-hold strategy. Our strategy takes relatively more risk when volatility is low (e.g., the 1960's) hence its losses are not surprisingly concentrated in these times. In contrast, large market losses tend to happen when volatility is high (e.g., the depression or recent financial crisis) and our strategy avoids these episodes. Because of this, the worst time periods for our strategy do not overlap much with the worst market crashes. This illustrates that our strategy works by shifting when it takes market risk and not by loading on extreme market realizations as profitable option strategies typically do.

From the vantage point of a sophisticated investor, a natural question that emerges from our findings is whether our volatility managed portfolios are capturing risk premia captured by well-known asset pricing factors. This question is relevant if the investor is already invested across multiple factors; in that case it is important to know if our volatility managed portfolios expand the unconditional mean-variance frontier relative to a particular investment opportunity set. We approach this question by considering alternative MVE portfolios formed using a different set of factors. Specifically, we compute the mean-variance efficient portfolio formed using static weights on a set of factors and then construct a volatility managed version of this portfolio using the realized variance of the portfolio in a given month. Thus, our volatility managed portfolios only shifts the conditional beta on the static MVE portfolio, but does not change the relative weights across individual factors. Given an investor's investment opportunity set, it is well known that the investor will want to choose the mean-variance efficient portfolio and then decide

still be a monthly number.

between this portfolio and the risk free asset. Therefore, this is precisely the portfolio the investor will want to volatility time. In Table 2, we show that the volatility timed mean-variance efficient portfolios have positive alpha with respect to the original MVE portfolio for all combinations of factors we consider including the Fama French three and five factors, or the Hou, Xue, and Zhang factors. This finding is robust to including the momentum factor as well. Appraisal ratios are all economically large and range from 0.33 to 0.91.

We also analyze these mean-variance portfolios across three 30-year sub-samples (1926-1955, 1956-1985, 1986-2015) in Panel B. The results generally show the earlier and later periods as having stronger, more significant alphas, with the results being weaker in the 1956-1985 period, though we note that point estimates are positive for all portfolios and for all subsamples. This should not be surprising as our results rely on a large degree of variation in volatility to work. For example, if volatility were constant over a particular period, our strategy would be identical to the buy-and-hold strategy and alphas would be zero. Thus, alphas tend to be larger over subsamples where volatility varies the most. Volatility varied far less in the 1956-1986 period, consistent with lower alphas during this time.

Overall, our volatility managed portfolios substantially expand the mean-variance frontier. This is true in a univariate sense, when one considers each factor in isolation, but also in a multi-factor sense because the volatility managed mean-variance efficient portfolios have substantial appraisal ratios.¹⁰

3. Understanding the profitability of volatility timing

In this section we investigate why our strategy works. Section 3.1 shows how the profitability of volatility timing is related to the risk-return trade-off. Section 3.2 investigates potential explanations for our strategy's high risk-adjusted returns. We first focus on risk-based explanations and then address explanations based on leverage constraints,

¹⁰Static MVE Sharpe ratios are likely to be overstated relative to their true population moments, since the weights are constructed in sample. Thus, the increase in Sharpe ratios we document are likely to be understated. We thank Tuomo Vuolteenaho for this point.

borrowing frictions, and transaction costs. We conclude this section by discussing several additional robustness exercises which are left to Appendix A.

We conclude that there is no obvious explanation for our findings. A reader interested in the portfolio choice implications of our main finding is free to skip to Section 4.

3.1 Volatility timing and the risk-return trade-off

We start by showing theoretically that the unconditional alpha of our strategy is tightly linked to the strength of the risk-return trade-off in the time-series.

We find it most useful to discuss our results in terms of conditional vs. unconditional models (see, e.g., Jagannathan and Wang (1996)). In particular, all conditional alphas are zero in our setting because our strategies are managed versions of the factors themselves. More specifically, our strategies are of the form $z_t R_{t+1}$ which, by definition, must have zero conditional alpha with respect to R_{t+1} . Thus, it follows that $E_t[z_t R_{t+1}] = \beta_t \mu_t$, where $\beta_t = z_t$ and $\mu_t = E_t[R_{t+1}]$. Taking unconditional expectations, letting β be the unconditional beta, and $\beta\mu$ be the expected return implied by an unconditional model, we have¹¹

$$\alpha = E[z_t R_{t+1}] - \beta\mu = (E[\beta_t] - \beta)\mu + cov(\beta_t, \mu_t) \quad (2)$$

This decomposes the unconditional alpha into a component that captures the ability of our strategy to time volatility, $E[\beta_t] - \beta$, and a component that captures the co-movement of our strategy with expected returns, $cov(\beta_t, \mu_t)$. If the strategy can time volatility and takes risk when conditional variance is low then the strategy's unconditional beta will be lower than its average beta and this first term will be positive.¹² Importantly, this volatility timing only generates alpha when volatility is weakly related to expected returns, so that the second term does not off-set the reduction in unconditional risk.

It is worth considering two cases. First, when there is no time-series relation between expected returns and variance so that $cov(\beta_t, \mu_t) \approx 0$, our strategy alpha becomes

¹¹Note that if z_t is uncorrelated with the moments of returns, the unconditional alpha is zero.

¹²See Boguth et al. (2011) for how this effect can cause a bias when the goal is to test a conditional factor model. In our setting the conditional model works perfectly and we show that the unconditional alpha of our strategy measures how beneficial is volatility timing.

$\alpha = \mu (E[\beta_t] - \beta)$, which is increasing in our ability to forecast volatility. Second, when the relationship between variance and expected returns is strong¹³, so that $\mu_t = \gamma\sigma_t^2$, the unconditional alpha is approximately zero because reducing risk exposure when volatility is high also sacrifices high expected returns rather than just avoiding high volatility. That is, if the risk-return tradeoff is strong there is no gain to volatility timing. Appendix B shows this in more detail.

In summary equation (2) makes it transparent that our strategy produces alpha if volatility is predictable and the relationship between volatility and expected returns is weak, not necessarily zero. Empirically, however, we find that for the volatility managed market portfolio, $\mu (E[\beta_t] - \beta) = 4.8\%$, while the actual alpha of our strategy is 4.86%. This confirms that indeed the term $cov(\beta_t, \mu_t)$ is near zero, consistent with no risk-return tradeoff in the data.

3.1.1 The strength of the risk return trade-off

We have shown that our strategy alpha is quantitatively consistent with no time-series relation between expected returns and risk. We now study this relationship directly and show that indeed our alpha results are consistent with alternative ways of measuring the risk return trade-off.

We start by showing that realized volatility for each factor strongly predicts future volatility (Table 3), but doesn't predict the factor's future returns (Table 4). 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 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 one reason why our strategy works. If a factors realized volatility is persistent and does not predict a large increase in expected returns, then an increase in volatility signifies a worse risk return tradeoff. These results complement a longer literature on the risk return tradeoff, which generally finds weak

¹³As would be the case in a general equilibrium model with a representative agent with risk aversion γ that holds the market.

forecasting power of variance for future returns (Glosten et al. (1993), Whitelaw (1994), Lundblad (2007), Lettau and Ludvigson (2003)).¹⁴

To see that the relationship between risk and return is indeed weak, we approach this question with a less parametric approach. We sort the time-series of market returns into five buckets based on the current month's realized volatility. We then look at the behavior of realized returns over the following month. Figure 1 plots the average realized return, volatility of returns, and the risk return tradeoff across these five buckets. The basic pattern is clear: high volatility this month forecasts high return volatility next month but not higher average returns. This means Sharpe ratios are lower after high volatility months. Even more dramatically, the relevant quantity $E_t(R_{t+1})/Var_t(R_{t+1})$ is strongly decreasing in volatility, which again shows that a mean-variance investor would want to time volatility.

The question in this paper is also different from the standard risk return tradeoff literature. 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. Our paper takes a portfolio strategy approach to the risk return tradeoff by showing portfolios can be formed in real time that take advantage of the risk-return regressions and produce very large risk-adjusted returns. This not only takes advantage of the weak risk return tradeoff but also the large variation in volatility.¹⁵ The large alphas, dramatic increase in Sharpe ratios, and large implied utility benefits from our volatility managed portfolios help quantify the failure of the risk return tradeoff in economic terms.

¹⁴See also related work by Bollerslev et al. (2016) and Tang and Whitelaw (2011). Also, relatedly, Cheng (2015) finds that the variance risk premium appears to *fall* in magnitude as variance rises, meaning that the risk return tradeoff in VIX futures potentially appears negative rather than positive.

¹⁵For example, even if the risk return tradeoff were indeed violated, if volatility didn't vary dramatically there wouldn't be a large benefit to volatility timing.

3.1.2 The dynamics of the risk return trade-off

Next, we frame 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 of a volatility shock for a myopic mean-variance investor. We set risk aversion just above 2 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 volatility, the portfolio weight is $w = 1$). This provides a natural benchmark to compare to.

We first estimate 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 month's stock returns on the price to earnings ratio, default spread, and realized variance (adding additional lags of each does not change the results). Expected variance is formed using a log normal model for volatility and including three lags for the market realized volatility.

We then take the estimated conditional expected return and variance and run a VAR with three lags of each variable. We use an Impulse Response Function to trace out the effects of a variance shock where we choose the ordering of the variables so that the variance shock can affect contemporaneous expected returns as well.

The results are in Figure 4. 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.

It is well known that movements in both stock-market variance and expected returns

are counter-cyclical (French et al., 1987; Lustig and Verdelhan, 2012), that is, both risk and expected returns tend to be high in recessions and low in booms. 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 gets out of the market immediately to avoid an unfavorable risk return tradeoff, but captures much of the expected return increase by buying back in when the volatility shock subsides.

These results square our findings with the portfolio choice literature on return predictability. 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 re-lever her positions once the volatility shock has died out to capture the low frequency movements in expected returns in bad times.

3.2 Economic explanations for the profitability of volatility timing

3.2.1 Business cycle risk

In Figure 3, we can see that the volatility managed factor has a lower standard deviation through recession episodes like the Great Recession where volatility was high. Table 5 makes this point more clearly across our factors. Specifically, we run regressions of each of our volatility managed factors on the original factors but also add an interaction term that includes an NBER recession dummy. The coefficient on this term represents the conditional beta of our strategy on the original factor during recession periods relative to non-recession periods. The results in the table show that, across the board for all factors, our strategies take less risk during recessions and thus have lower betas during recessions. For example, the non-recession market beta of the volatility managed market factor is 0.83 but the recession beta coefficient is -0.51, making the beta of our volatility managed portfolio conditional on a recession equal to 0.32. Finally, by looking at Figure 2 which plots the time-series realized volatility of each factor, we can clearly see that volatil-

ity for all factors tends to rise in recessions. Thus, our strategies decrease risk exposure in NBER recessions. This makes it difficult for a business cycle risk story to explain our facts. However, we still review several specific risk based stories below.

3.2.2 Other risk based explanations

Variance risk premia: Because our strategy aggressively times volatility a reasonable concern is that our strategy's high Sharpe ratio is due to a large exposure to variance shocks which would require a high risk premium (Ang et al., 2006b; Carr and Wu, 2009). However, it turns out that our strategy is much less exposed to volatility shocks than the buy-and-hold strategy. This follows from the fact that volatility of volatility is higher when volatility is high. Because our strategy takes less risk when volatility is high, it also less sensitive to volatility shocks.

Downside risk: In unreported results, we find that the downside betas of our strategy following the methodology in Lettau et al. (2014) are always substantially lower than unconditional betas. For example, for the volatility managed market return, the downside beta we estimate is 0.11 and isn't significantly different from zero. Thus, alphas would be even larger if we evaluated them relative to the downside risk CAPM (Ang et al. (2006a) and Lettau et al. (2014)). Intuitively, periods of very low market returns are typically preceded by periods of high volatility when our strategy has a low risk exposure.

Disaster risk: For disaster risk to explain our findings, our volatility managed portfolio would have to be more exposed to disaster risk than the static portfolio. Because empirically, macro-economic disasters unfold over many periods (Nakamura et al., 2010) and feature above average financial market volatility (Manela and Moreira, 2013), the volatility timing strategy tends to perform better during disaster events than the static counterpart. This is further supported by the fact that our strategy takes less risk in the Great Depression and recent financial crisis (see Figure 5), the two largest consumption declines in our sample.

Jump risk: Jump risk is the exposure to sudden market crashes. To the extent that crashes after low volatility periods happen frequently, our strategy should exhibit much fatter tails than the static strategy, yet we do not see this when analyzing the uncon-

ditional distribution of the volatility managed portfolios. Overall, crashes during low volatility times are just not frequent enough (relative to high volatility times) to make our volatility managed portfolio more exposed to jump risk than the static buy-and-hold. If anything, jumps seem to be much more likely when volatility is high (Bollerslev and Todorov, 2011), suggesting that our strategy is less exposed to jump risk than the buy-and-hold portfolio.

3.2.3 Contrasting with cross-sectional low-risk anomalies

It is useful to contrast our analysis with strategies that explore a weak risk return trade-off in the *cross-section* of stocks.

The first strategy, popular among practitioners, is risk parity. Risk parity is mostly about cross-sectional allocation. Specifically, risk parity ignores information about expected returns and co-variances and allocates to different asset classes or factors in a way that makes the total volatility contribution of each asset the same (see for example Asness et al. (2012)). This implies that, if the volatility of one factor spikes relative to other factors, the strategy will rebalance from the high volatility factor to the low volatility factor. In contrast, when we time combinations of factors, as in Table 2, we keep the relative weights of all factors constant and only increase or decrease overall risk exposure based on total volatility. Thus, our volatility timing is conceptually quite different from risk parity. To assess this difference empirically, we follow Asness et al. (2012) and construct a risk parity factor that we then use in Table 6 as a control in our time-series regression. The alphas are basically unchanged. We thus find that controlling for the risk parity portfolios constructed following Asness et al. (2012) has no effect on our results, suggesting that we are picking up a different empirical phenomenon.

The second strategy is the betting against beta factor (BAB) of Frazzini and Pedersen (2014). They show that a strategy that goes long low beta stocks and shorts high beta stocks can earn large alphas relative to the CAPM and the Fama-French three factor model that includes a Momentum factor. Conceptually, our strategy is quite different. While the high risk-adjusted return of the BAB factor reflects the fact that differences in average returns are not explained by differences in CAPM betas in the cross-section, our strategy is

based on the fact that across time periods, differences in average returns are not explained by differences in stock market variance. Our strategy is measuring different phenomena in the data. In the last column of Table 6 we show further that a volatility managed version of the BAB portfolio also earns large alphas relative to the buy-and-hold BAB portfolio. Therefore, one can volatility time the cross-sectional anomaly. In addition to this, we also find that our alphas are not impacted if we add the BAB factor as a control. These details are relegated to the Appendix. Thus, our *time-series* volatility managed portfolios are distinct from the low beta anomaly documented in the cross-section.

In addition our results are different from Ang et al. (2006b) who study the link between *idiosyncratic* risk and the cross-section of returns. By focusing on systematic risk factors, we are able to say something about the evolution of the aggregate price of risk over time. Volatility timing on an individual stock will not tell us about risk compensation over time because the majority of the stock's volatility is idiosyncratic.

3.2.4 Leverage constraints

Black (1972), Jensen et al. (1972) and more recently Frazzini and Pedersen (2014) show that leverage constraints can distort the risk-return trade-off in the cross-section. Empirically, high beta assets have an abnormally low average return relative to that predicted by the CAPM. The idea is that the embedded leverage of high beta assets make them attractive to investors that are leverage constrained.

Applying this idea to the time-series, one could argue that low volatility periods are analogous to low beta assets, and as such have expected returns that are too high relative to the baseline model. Intuitively, leverage constrained investors cannot leverage their position during low volatility periods explaining why average returns are too high in these periods.

On a conceptual level, this explanation has potential. However we see two challenges for it to explain our findings. First, in order for leverage constraints to explain our findings, the constraint would need to be tighter during periods of low volatility. This would be inconsistent with both theory and empirical evidence that it is easier to take on leverage in periods of low volatility (Brunnermeier and Pedersen (2009), Adrian et al. (2014)).

Next, we consider the role of leverage constraints empirically. We start by considering a simple strategy that only updates the portfolio when volatility is above its mean value. This portfolio avoids the use of leverage in low volatility episodes where the risk weight would normally rise substantially (see next section where we discuss these results in the context of transaction costs). We also construct a volatility managed strategy for an investor whose portfolio weight on the market is on average 50% and who can't use leverage. We find that both strategies have large alphas and substantial gains in Sharpe ratios.

For investors whose risk-aversion is low enough, our baseline strategy uses leverage. For an investor that wants to be 100% in stocks on average, the volatility timing mechanically requires leverage about half of the time, when volatility is below its average. To address the issue that leverage might be costly, we implement our strategy using options in the S&P 500. Specifically we use the option portfolios from Constantinides et al. (2013). We focus on in-the-money call options with maturities of 60 and 90 days and whose market beta is around 7. Whenever the strategy prescribes leverage, we can instead use the option portfolios to achieve our desired risk exposure. In Table 7, we compare the strategy implemented with options with the one implemented with leverage. The alphas are very similar showing that our results are not due to leverage constraints.

In light of recent work by Frazzini and Pedersen (2012), the fact that our strategy can be implemented through options should not be surprising. Frazzini and Pedersen (2012) show that, for option strategies on the S&P 500 index with embedded leverage up to 10, there is no difference in average returns relative to strategies that leverage the cash index. This implies that our strategy can easily be implemented using options for relatively high levels of leverage.

In the end, the potential for the leverage story to explain our findings depends on how leveraged an investor wants to be on average. As desired average leverage increases beyond 10, the benefits of timing the weak risk-return trade-off in the time-series is increasingly offset by the weak risk return trade-off in the cross-section. However, the *average* investor is an investor who holds the market, and this investor can substantially benefit from our strategy without explicitly needing to use leverage.

3.2.5 Transaction costs

We show that our strategies survive transaction costs. These results are given in Table 8. Specifically, we evaluate our volatility timing strategy for the market portfolio when including empirically realistic transaction costs. We consider various strategies that capture volatility timing but reduce trading activity, including using standard deviation instead of variance, using expected rather than realized variance, and only trading when variance is above the long run average (a strategy that does not use any leverage). Each of these reduces trading and hence reduces transaction costs. We report the average absolute change in monthly weights, expected return, and alpha of each strategy before transaction costs. Then we report the alpha when including various transaction cost assumptions. The 1bp cost comes from Fleming et al. (2003); the 10bps comes from Frazzini et al. (2015) which assumes the investor is 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 in a regression of transaction costs on VIX from Frazzini et al. (2015) to evaluate the impact of 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 the volatility managed strategy decreases somewhat for the market portfolio, but is still very large. 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 lower. Note, however, that we do not study trading costs of the original factors.

These results on transaction costs and the results dealing with leverage constraints together suggest that our strategy can be realistically implemented in real time.

3.3 Additional robustness checks

We conduct a number of additional robustness checks of our main result but leave the details to Appendix A. We show that our strategy works across 20 OECD stock market indices, that it can be further improved through the use of more sophisticated models of variance forecasting, that it does not generate fatter left tails than the original factors or create option-like payoffs, and that it outperforms not only using alpha and Sharpe ratios but also manipulation proof measures of performance (Goetzmann et al., 2007).

4. Portfolio choice implications

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

We focus on the portfolio problem of an investor that trades a risky portfolio and a risk-free bond, and calibrate the parameters in our analysis to be consistent with the market portfolio. Our results for the mean-variance investor in Section 4.1 extend directly to the other portfolios considered in Section 2; specifically, they carry over for the different ex-post MVE portfolios we study in Table 2. Our results for a long-horizon investor in Section 4.2 rely on the vast literature on return predictability which focuses on the aggregate market.

4.1 Volatility timing for a mean-variance investor

We now use simple mean-variance preferences to measure the benefits of volatility timing. The mean variance investor wants to set his risk exposure as $E_t[R_{t+1}]/Var_t(R_{t+1})$. A natural question is whether this investor is better off learning about the conditional mean or conditional variance. This comparison provides a useful benchmark to evaluate the benefits of volatility timing as the benefits of expected return timing are well studied

in the long literature on return predictability and portfolio choice (e.g., Barberis (2000)). Specifically, we extend the analysis in Campbell and Thompson (2008) to allow for time-variation in volatility, and use it to compare the benefits of timing volatility and expected returns for the market portfolio.¹⁶

4.1.1 Expected return timing

Consider the following excess return process,

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

where $E[x_t] = E[e_t] = 0$ and x_t, e_{t+1} are conditionally independent. We begin by studying an investor who follows an unconditional strategy and an expected return timing strategy. Later, we study an investor who times volatility rather than expected returns. For a mean-variance investor his portfolio choice can be written as follows,

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

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

where the portfolio policy in Equation (4) uses no conditional information, and the portfolio policy in Equation (5) uses information about the conditional mean but ignores fluctuations in volatility. We now compare expected returns across the two conditioning sets, which also measures the investor's expected utility. We multiply factor returns by the portfolio weight before taking unconditional expectations:

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

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

¹⁶See also Breen et al. (1989) and Busse (1999) on volatility timing.

where S is the unconditional Sharpe ratio of the factor $S = \frac{\mu}{\sqrt{E[Y_t] + \sigma_x^2}}$ and R_r^2 is the share of the return variation captured by the forecasting signal x (e.g., it is the R-square in a predictive return regression). The proportional increase in expected returns (and utility) is $\frac{1+S^2}{S^2} \frac{R_r^2}{1-R_r^2}$, which is increasing in the degree of return predictability.

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 investor can experience a proportional increase in expected returns and utility of roughly 35% by using conditional variables known 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. Following the analysis in Campbell and Thompson (2008) and consistent with our empirical findings this exercise assumes there is no risk-return trade-off in the time-series.¹⁷

We assume the variance process is log-normal $Y_t = e^{y_t}$, with $y_t = \hat{y}_t + u_t$, where \hat{y}_t is the forecastable component of stock market variance.

This produces optimal portfolio weights and expected returns given by

$$w(y_t) = \frac{1}{\gamma} \frac{\mu}{e^{2(\hat{y}_t + \sigma_u^2)}}; \quad (8)$$

$$E[w(y_t)r_{t+1}] = \frac{1}{\gamma} S^2 e^{R_y^2 \text{Var}(\hat{y}_t)}, \quad (9)$$

where equation (9) can be written as a function of the R-squared of the forecasting regression of future realized variance on the signal \hat{y} , $E[w(y_t)r_{t+1}] = \frac{1}{\gamma} S^2 e^{R_y^2 \text{Var}(y_t)}$. The proportional expected return gain of such a strategy is simply $e^{R_y^2 \text{Var}(y_t)}$. The total monthly variance of log realized variance is 1.06 in the full sample (1926-2015), and a simple model that uses lagged variance as the forecast of future variance achieves a $R_y^2 = 38\%$, implying a proportional expected return increase of 50%. A slightly less naive model that takes

¹⁷We relax this assumption in Appendix C.

into account mean-reversion and uses the lagged realized variance to form an OLS forecast (i.e., an AR(1) model for variance) achieves $R_y^2 = 53\%$. 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-squared, implying an expected return increase of 90%. These estimates do not depend on whether the R-squared is measured in or out of sample as this relationship is stable over time.

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-squared 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 of and access to these variables in real time. In contrast, data on volatility is much 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 for the several additional factors studied in Section 2, 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-squared 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 the other factors. 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 factors because lags of the factors' own variance are a reliable and stable way to estimate conditional expected variance across factors.

Because our focus is on volatility timing, in the main text we abstract from strategies that time both expected returns and volatility. In Appendix C we study the general case where investors time volatility and expected returns simultaneously. We show gains as high as 200%. The main reason this number is so large is because the estimated risk return

trade-off in the data is very weak, meaning that the combination of information provides large timing gains.

4.2 Volatility timing for a long-horizon investor

We now study the problem of a long-horizon investor and investigate how much they should adjust their portfolio to changes in volatility. This analysis is critically important to interpret our findings because a fraction of returns mean-revert (Campbell and Shiller, 1988; Poterba and Summers, 1988), i.e. low returns today are at least partially compensated by an increase in expected returns going forward. This empirical fact implies that short-term stock market volatility is not the the relevant measure of risk for a long-horizon investor.

Following the terminology in Merton (1973), the optimal portfolio demand for a long-horizon investor can be written as the sum of a myopic demand, the demand of a short-horizon investor, and a hedging demand, which emerges when returns are correlated with changes in the investment opportunity set. Specifically, the hedging demand will lead a long-horizon investor to choose more volatile portfolios than a short-horizon investor when returns feature some mean-reversion. Intuitively, a long-horizon investor likes mean-reversion because the associated increase in expected returns works as a partial hedge for the bad return realization, making stocks less volatile in the long run.

This hedging effect means that the implications of our empirical findings for a long-horizon investor are not immediate. An increase in volatility might generate an increase in the hedging demand that is enough to completely offset the reduction in exposure due to the myopic demand– i.e. it might be that long-horizon investors should just ignore time-variation in volatility, in line with the advice of Cochrane (2008) and Buffett (2008).¹⁸ We therefore need to dig deeper into the long-horizon investor portfolio problem in order

¹⁸Cochrane (2008) articulates the argument nicely: “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 – (...)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.”

to see how the investor will behave once the hedging demand is taken into account.

We solve the portfolio problem of an infinitely lived investor in a stochastic environment that allows for (i) variation in volatility; (ii) mean-reversion in returns (i.e. time-variation in expected returns); (iii) weak time-series relation between expected returns and volatility and (iv) variation in the amount of mean-reversion in returns, i.e. variation in the share of return volatility due to mean-reverting shocks. Together, these ingredients are novel and essential to interpret our findings.¹⁹ Ingredients (i) and (iii) allow us to fit our main findings. Ingredient (ii) allow us to investigate whether the presence of mean-reversion changes the portfolio choice implication for a long-horizon investor.

Ingredient (iv) allows us to consider the possibility that an increase in volatility is associated with a greater-than-proportional increase in the amount of mean-reversion in stock returns. This turns out to be essential to capture the intuition behind Cochrane (2008) and Buffett (2008) argument. As we will show, the investment horizon will only impact how an investor should *respond* to changes in volatility if the amount of mean-reversion in stock returns does not vary proportionally with volatility, i.e. if the share return volatility due to mean-reverting shocks co-moves with volatility .

We now discuss the portfolio choice framework and analyze the solution with a focus on the optimal investment response to changes in volatility. Details are in Appendix D.

4.2.1 Investment opportunity set and preferences

We assume there is a riskless bond that pays a constant interest rate rdt , and a risky asset S_t , with dynamics given by

$$\frac{dS_t}{S_t} = (r + x_t)dt + \sqrt{y_t}DdB_t + FdZ_t, \quad (10)$$

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 expected excess return, and y_t is a scalar that drives return volatility. The shocks dB_t and dZ_t are independent three dimensional

¹⁹Earlier work on portfolio choice has studied expected return variation, volatility variation, or volatility variation with a constant risk-return trade-off. Examples of work that study volatility timing in a dynamic environment are Chacko and Viceira (2005) and Liu (2007).

Brownian motions, where each row captures shocks to realized returns, expected returns, and volatility. We need two different Brownian motions dB_t and dZ_t to allow for non-proportional variation in the volatility of mean-reverting shocks. Formally the risky asset variance is $\sigma^2(y) \equiv yD'D + F'F$ and the state variables evolve as follows,

$$dx_t = \kappa_x(\mu_x - x_t)dt + \sqrt{y_t}GdB_t + HdZ_t \quad (11)$$

$$dy_t = \kappa_y(\mu_y - y_t)dt + \sqrt{y_t}LdB_t. \quad (12)$$

The vector of return exposures D and F are chosen so that there are two different types of shocks impacting returns that we label permanent and mean-reverting. Permanent shocks are shocks to returns that are uncorrelated to shocks to expected returns. On the other hand, mean-reverting shocks are shocks to returns that are correlated with changes in expected returns in a way that the long run expected value of the asset S_t does not change. That is, these shocks lower returns today but raise expected returns by enough so that their long run effects are neutral. Permanent and mean-reverting shocks are also often referred in the literature as cash-flow and discount-rate shocks (Campbell and Vuolteenaho, 2004).

Investors have CRRA preferences, $E_t \left[\int_t^\infty e^{-\rho s} \frac{C_s^{1-\gamma}}{1-\gamma} ds \right]$, where ρ is the rate of time preference and γ the coefficient of relative risk aversion. In Appendix D we also consider more general Epstein and Zin (1989) preferences. The investor maximizes utility subject to his intertemporal budget constraint (Equation 13 below) and the evolution of state variables.

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

$$\frac{dW_t}{W_t} = \left[w_t x_t + r - \frac{C_t}{W_t} \right] dt + w_t \sqrt{y_t} D dB_t + w_t F dZ_t. \quad (13)$$

4.2.2 Solution

The optimization problem has three state variables: the investor's wealth plus the two drivers (expected return x and volatility y) of the investment opportunity set.

To get intuition about the results that follow, let the value function of the agent be given by $J(W, x, y) = \frac{W^{1-\gamma}}{1-\gamma} e^{V(x,y)}$. Then, in the case where realized returns are unrelated to innovations in volatility, $L'D = 0$, the optimal portfolio policy can be written as

$$w_t = \underbrace{\frac{x_t}{\gamma\sigma^2(y_t)}}_{\text{myopic demand}} - \underbrace{\frac{\kappa_x V_x}{\gamma}\theta(y_t)}_{\text{hedging demand}}, \quad (14)$$

where $\theta(y) = \frac{G'Dy_t + F'H}{\kappa_x\sigma^2(y)}$ is the share of return volatility driven by mean-reverting shocks.

The first term is the myopic demand, which reflects the portfolio choice of a short-horizon investor. The second term is the Mertonian hedging demand.²⁰ The fact that expected returns increase after low return realizations makes investment in the risky asset a natural investment opportunity set hedge. Thus, the hedging demand leads the long-horizon investor to have a higher average position in the risky asset (V_x is typically negative).

Our analysis emphasizes how the strength of this hedging demand depends on the composition of changes in volatility. If the volatility of mean-reverting shocks varies proportionally to total volatility, $\theta(y_t)$ is constant and this hedging demand does not fluctuate with volatility, thus the long-horizon investor responds to volatility changes just like the short-horizon investor.²¹ Thus, mean-reversion will always increase the average position of the long-horizon investor relative to the short-horizon investor, but the presence of mean-reversion alone does not imply that a long-horizon investor should ignore variation in volatility.

From Equation 14 we see that only if the volatility of mean-reverting shocks increases more than proportionally with volatility ($\theta'(y_t) > 0$), the hedging demand will call for an increase in risk exposure. This effect counteracts the myopic demand, which calls for a reduction in risk exposure. Thus, the advice that a long-horizon investor should

²⁰Earlier work on portfolio choice featuring time-varying volatility has focused on the hedging demand that arises when $L'D \neq 0$, which can be 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 turns out not to be quantitatively important. Here we assume $L'D = 0$ just for illustration purposes, though this has little effect on results as we show in Appendix D.

²¹Formally, changes in volatility could impact hedging demands through V_x , but this effect is small.

ignore variation in volatility can be right only if the increase in the hedging demand is sufficiently large. Intuitively, if an increase in volatility is associated with a greater-than-proportional increase in the amount of mean-reversion in stock returns, then, since a long-horizon investor likes mean-reversion, he will want to increase exposure, counteracting the myopic demand. Figure 6 provides a graphical illustration of this mechanism in the stark case where there are only mean-reverting shocks.

The opposite effect can also occur – if increases in volatility are associated with a reduction in the share of mean-reverting shocks ($\theta'(y) < 0$) then a long-horizon investor will want to volatility time more than a short horizon investor.

Investment horizon plays a role in shaping the strength of the hedging demand 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.

We solve for V numerically and study how the portfolio choice implied by equation (14) varies with volatility.

4.2.3 Calibration

Our focus is to evaluate how much a long-run investor will choose to time volatility for an empirically realistic process for returns and volatility. Our strategy is to match all moments that can be directly measured in the data. Because we are not aware of any work that measures how the amount of mean-reversion in returns co-moves with volatility, we consider alternative calibrations that embed different assumptions about this co-movement.

Table 9 reports the moments used in our calibration. In calibrations (a), (b), and (c), increases in volatility are associated with either greater-than-proportional (a), less-than-proportional (b), or proportional (c) increases in the amount of mean-reversion. Calibration (c) is a natural benchmark commonly *assumed* in papers that try to empirically identify mean-reverting and permanent shocks in stock returns (e.g. Campbell and Vuolteenaho, 2004).

We calibrate the stochastic processes to be consistent with the following moments: av-

erage volatility, the R-squared 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 price-dividend ratio (expected return state variable x).²²

The persistence and volatility of volatility determine how much predictable variation in volatility there is. The R-squared of the return predictability regression controls the average amount of mean-reversion. The auto-correlation of the expected return state variable controls how long mean-reversion takes. This auto-correlation plays a key role in determining how long an investment horizon an investor must have in order to be completely hedged with respect to shocks that mean-revert.

4.2.4 Analysis

In the main text we discuss results for an investor with CRRA preferences (with coefficient of relative risk aversion of 10).²³ Following the analysis in Blanchard (1985) and Gârleanu and Panageas (2015), we use the preference parameter ρ to vary the investor horizon. We choose ρ so that the half-life of utility weights ranges from 5 year to 30 years.²⁴

To focus on the effect of volatility we fit the following linear relation on the optimal policy evaluated at the unconditional equity premium²⁵

$$w^*(y) \approx a + b \times \frac{\mu_x}{\gamma\sigma^2(y)}. \quad (15)$$

The policies of a short-horizon investor fit this linear relation exactly with $a = 0$ and $b = 1$. In short, the short-horizon investor puts zero weight on the buy-and-hold portfolio and weight 1 on the volatility managed portfolio. In Figure 7 we report b as a function of investment horizon. The reported coefficients have the direct interpretation of how much

²²In the main text we focus on the effects of volatility variation and abstract from subtle hedging demand effects arising from any contemporaneous correlation between return and volatility shocks and assume zero correlation between volatility and expected return shocks $L[3] = 0$. In Appendix D.4 we study this case and show that it does not impact our results.

²³See Appendix D.4 for Epstein and Zin preferences and several values of risk aversion.

²⁴Specifically, for a given horizon T_j , we solve for ρ_j such that $\frac{\int_0^{T_j} e^{-\rho_j \times t} dt}{\int_0^\infty e^{-\rho_j \times t} dt} = \frac{1}{2}$

²⁵In Appendix D.4 we show that his linear approximation fits the optimal policy extremely well.

weight the investor places on the volatility managed portfolio. A coefficient lower than one implies the long-horizon investor times volatility less aggressively than the short-horizon investor.

The flat line shows case (c) where the amount of mean-reversion increases proportionally with volatility. Coefficients are equal to 1 and the long-horizon investor responds to volatility changes exactly as a short-horizon investor. Thus, the presence of mean-reversion in returns alone does not imply that a long-horizon investor responds differently to volatility changes than a short-horizon investor. Mean-reversion increases the average position a long-horizon investor has in stocks, but because in case (c) the increase in volatility is proportional, the risky asset becomes equally riskier as volatility increases. In other words, the hedging demand is unconditionally large but doesn't vary as a function of volatility.

For the intuition in Figure 6 to be qualitatively right, the amount of mean-reversion must increase more than proportionally with volatility (case (a)). In this case, long-horizon investors respond less than short-horizon investors to increases in volatility because the risky asset becomes relatively safer as a greater share of the volatility is due to less risky mean-reverting shocks. While qualitatively the intuition in Figure 6 is right, quantitatively the effect due to the hedging demand is not strong enough to fully counteract the myopic demand even in this extreme case. The loading on the volatility managed portfolio is still quite high. For example, during the fall of 2008 volatility spiked from 20% to 60%. This would induce the mean-variance investor to shrink his risk exposure by 90%, while an investor with a utility half-life of 30 years and coefficient of relative risk-aversion 10 would shrink his risk exposure by 45% ($0.5 \cdot 90\%$). Even long-horizon investors perceive mean-reverting shocks as somewhat risky because in the data mean-reversion in returns is slow. Thus, an investor needs a very long horizon in order to fully benefit from mean-reversion.

When the amount of mean-reversion increases less than proportionally with volatility, the long-horizon investor responds more than the short-horizon to changes in volatility. The risky asset becomes even riskier for a long-horizon investor as a smaller share of volatility is due to less risky mean-reverting shocks. The same mechanism that makes

long-horizon investors less responsive to changes in volatility driven by mean-reverting shocks makes them more responsive to changes in volatility driven by permanent shocks. Thus, the magnitude of the response of the long-horizon investor is highly dependent on the co-movement between volatility and the share of volatility due to mean-reverting shocks. But, quantitatively, it is clear from Figure 7 that for reasonable investment horizons, an investor always find it optimal to time volatility.

Both Cochrane (2008) and Buffett (2008) argued that investors with a long horizon should ignore the massive increase in volatility in the fall of 2008, and if anything the large drop in equity prices provided investors with a unique buying opportunity. Their intuition is qualitatively correct if the increase in volatility in the fall of 2008 was mostly about mean-reverting shocks. Quantitatively their intuition does not quite work. The change in the hedging demand is not large enough to offset the large reduction coming from the myopic demand. Even in the extreme case where all the increases in volatility is due to an increase in the volatility of mean-reverting shocks, long-horizon investors still benefit from volatility timing. The key for this result is that the high persistence of the price-dividend ratio in the data implies that mean-reversion takes a long-time. As a result, investors must have very long horizons in order to safely “ignore” volatility coming from mean-reverting shocks.

Importantly, there is yet no empirical evidence on how mean-reversion changes with volatility. In the natural case where the amount of mean-reversion is constant the long-run invests exactly as the short-horizon investor. Thus, our analysis shows that long-run investors should ignore volatility at their own peril.

5. Discussion and implications for economic theory

An obvious caveat to our results is that not all investors can engage in volatility timing in equilibrium – market clearing means the average investor has to hold the market portfolio. Suppose instead that security prices are such that investors are happy with not volatility timing. What, then, can we say about those investors?

First, it is worth sketching the implications of our empirical results for equilibrium

theories of time-varying risk-premia (e.g. Campbell and Cochrane, 1999; Barberis et al., 2001; Bansal and Yaron, 2004; He and Krishnamurthy, 2012; Wachter, 2013). Our results speak to these theories because their aim is to explain the level and variation in risk-premia through a combination of time-variation in the quantity and the price of risk. Our empirical work allows us to study whether the joint dynamics these models rely on is consistent with the data.

Specifically, Figure 1 shows that in the data compensation per unit of risk must be lower when volatility is high, i.e. the market effective risk aversion, $E_t[R_{t+1}]/\sigma_{R,t}^2$, is lower in periods of high volatility.²⁶ In the leading calibrations of the aforementioned models this relation is either positive or flat, so that effective risk aversion either roughly doesn't change with stock market volatility or actually rises in bad times when volatility is also high.

The challenge is the following. Suppose R_{t+1} is the excess return on the asset with highest conditional Sharpe ratio (e.g. the market portfolio), then it immediately follows that the equity risk premium is given by

$$E_t[R_{t+1}] \approx \sigma_{m,t}\sigma_{R,t}, \quad (16)$$

where $\sigma_{m,t}$ is the volatility of the pricing kernel or stochastic discount factor (sdf). These models achieve variation in the equity premium by inducing time-variation in the sdf volatility. From (16) we see that the degree of co-variation between $\sigma_{m,t}$ and $\sigma_{R,t}$ is critically important for the level and the amount of variation in the equity premium the model can explain.²⁷

Variation in the stock market volatility $\sigma_{R,t}$ is derived in equilibrium as a result of either variation in fundamental risk (e.g. Bansal and Yaron, 2004; Wachter, 2013) or risk-

²⁶For a formal connection between our strategy alpha and the fact that effective risk aversion is decreasing in volatility see Appendix E. Appendix E.1 also shows that the data allow us formally reject that this relationship is greater or equal to zero.

²⁷It is immediate from Equation 16 that $E[R_{t+1}] = cov(\sigma_{m,t}, \sigma_{R,t}) + E[\sigma_{m,t}]E[\sigma_{R,t}]$. Therefore the unconditional equity premium is increasing in this co-variance. For the volatility of the conditional equity premium it should be intuitive from the multiplicative functional form that the variance of the product of volatility is increasing in the co-variance between them. In the case the volatility processes are log-normally distributed this can be shown analytically.

aversion (e.g. Campbell and Cochrane, 1999; Barberis et al., 2001; He and Krishnamurthy, 2012), but all the leading calibrations predict a strong co-movement between $\sigma_{m,t}$ and $\sigma_{R,t}$. Variation in stock market risk acts as an amplifier of variation in the sdf volatility, which helps generate a large and volatile equity premium. In time-varying risk aversion models this amplification is so strong that the market effective risk aversion goes up with volatility. In standard calibrations of fundamental risk models, stock market risk moves roughly proportionally with the volatility of the sdf implying a constant effective risk-aversion. Neither case is consistent with our findings.

What, then, could explain our results? A definitive answer to this question is beyond the scope of this paper and left for future work. Nevertheless, we speculate a few possibilities.

First, our findings in the portfolio choice section suggests a model where most of the variation in volatility is driven by discount rate volatility (e.g. Campbell and Cochrane, 1999; He and Krishnamurthy, 2012) could work *if* the level of discount rates were not as tightly related to it's volatility as in these models. When most variation in volatility is driven by discount rate volatility, long-horizon investors find volatility timing somewhat less beneficial. However, in the leading asset pricing models, discount rate volatility is high when discount rates themselves are also high, leading to the tight relationship between volatility and expected returns. Second, the literature on parameter uncertainty (e.g Veronesi, 2000) can generate a negative relation between expected returns and volatility under certain parameters.

Another possibility is that investors have incorrect beliefs about risk and return. For example, agents might not update their beliefs about volatility quickly enough. This would explain why a sharp increase in realized volatility doesn't immediately illicit a response to sell even if expected returns do not rise quickly. This is potentially consistent with our impulse responses where expected returns rise slowly but the true expected volatility process rises quickly in response to a variance shock. In line with this view, Nagel et al. (2016) find that many households respond to volatility by selling stocks in 2008 and that this effect is larger for higher income households who may be more sophisticated traders. In the data, it appears the relatively lower income households respond

to volatility more slowly, selling much later in the crisis. However, we acknowledge that these or other explanations need to be considered in much more detail and be analyzed quantitatively before we can evaluate their success.

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. Thus, investors following our strategy would seem to behave in a “panicked” manner, and nevertheless earn superior risk-adjusted returns than investors that muster the will to follow the conventional buy-and-hold advice. We analyze both portfolio choice and general equilibrium implications of our findings. We find that both short and long-horizon investors can benefit from our volatility timing strategy. Furthermore, we show that our strategy performance is informative about the dynamics of effective risk-aversion, a key object for theories of time-varying risk premia.

References

- Adrian, T., Etula, E., and Muir, T. (2014). Financial intermediaries and the cross-section of asset returns. *Journal of Finance*, forthcoming.
- 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., Chen, J., and Xing, Y. (2006a). Downside risk. *Review of Financial Studies*, 19(4):1191–1239.
- Ang, A., Hodrick, R. J., Xing, Y., and Zhang, X. (2006b). The cross-section of volatility and expected returns. *Journal of Finance*, 61(1):259–299.
- Asness, C. S., Frazzini, A., and Pedersen, L. H. (2012). Leverage aversion and risk parity. *Financial Analysts Journal*, 68(1):47–59.

- Bansal, R. and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509.
- Barberis, N. (2000). Investing for the long run when returns are predictable. *The Journal of Finance*, 55(1):pp. 225–264.
- Barberis, N., Huang, M., and Santos, T. (2001). Prospect theory and asset prices. *The Quarterly Journal of Economics*, 116(1):1–53.
- Barroso, P. and Santa-Clara, P. (2015). Momentum has its moments. *Journal of Financial Economics*, 116(1):111–120.
- Black, F. (1972). Capital market equilibrium with restricted borrowing. *The Journal of Business*, 45(3):444–455.
- Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *The journal of political economy*, pages 637–654.
- Blanchard, O. J. (1985). Debt, deficits, and finite horizons. *The Journal of Political Economy*, 93(2):223–247.
- Boguth, O., Carlson, M., Fisher, A., and Simutin, M. (2011). Conditional risk and performance evaluation: Volatility timing, overconditioning, and new estimates of momentum alphas. *Journal of Financial Economics*, 102(2):363–389.
- Bollerslev, T., Hood, B., Huss, J., and Pedersen, L. H. (2016). Risk everywhere: Modeling and managing volatility. *working paper*.
- Bollerslev, T. and Todorov, V. (2011). Tails, fears, and risk premia. *The Journal of Finance*, 66(6):2165–2211.
- Breen, W., Glosten, L. R., and Jagannathan, R. (1989). Economic significance of predictable variations in stock index returns. *The Journal of Finance*, 44(5):1177–1189.
- Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial studies*, 22(6):2201–2238.
- Buffett, W. E. (2008). Buy American. I am. *The New York Times*, October 16th, 2008.
- Busse, J. A. (1999). Volatility timing in mutual funds: Evidence from daily returns. *Review of Financial Studies*, 12(5):1009–1041.
- 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.
- Campbell, J. Y. and Vuolteenaho, T. (2004). Bad beta, good beta. *American Economic Review*, pages 1249–1275.
- Carr, P. and Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, 22(3):1311–1341.
- 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.
- Cheng, I.-H. (2015). The expected returns to fear. *working paper*.
- Cochrane, J. H. (2008). Is now the time to buy stocks? *The Wall Street Journal*, November 12th, 2008.
- Constantinides, G. M., Jackwerth, J. C., and Savov, A. (2013). The puzzle of index option returns. *Review of Asset Pricing Studies*, 3(2):229–257.
- Daniel, K., Hodrick, R. J., and Lu, Z. (2015). The carry trade: Risks and drawdowns. *working paper*.
- Daniel, K. and Moskowitz, T. (2015). Momentum crashes. *Journal of Financial Economics*, forthcoming.
- Duffie, D. and Epstein, L. G. (1992). Stochastic differential utility. *Econometrica: Journal of the Econometric Society*, pages 353–394.
- Dybvig, P. H. and Ingersoll Jr, J. E. (1982). Mean-variance theory in complete markets. *Journal of Business*, pages 233–251.
- 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. (1989). Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics*, 25(1):23–49.
- Fama, E. F. and French, K. R. (1996). Multifactor explanations of asset pricing anomalies. *Journal of Finance*, 51(1):55–84.
- Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22.
- 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. (2012). Embedded leverage. Technical report, National Bureau of Economic Research.
- 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.
- Goetzmann, W., Ingersoll, J., Spiegel, M., and Welch, I. (2007). Portfolio performance manipulation and manipulation-proof performance measures. *Review of Financial Studies*, 20(5):1503–1546.
- He, Z. and Krishnamurthy, A. (2012). Intermediary asset pricing. *The American Economic Review*, forthcoming.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of financial studies*, 6(2):327–343.
- Hou, K., Xue, C., and Zhang, L. (2014). Digesting anomalies: An investment approach. *Review of Financial Studies*, forthcoming.
- Jagannathan, R. and Wang, Z. (1996). The conditional capm and the cross-section of expected returns. *Journal of Finance*, 51(1):3–53.
- Jensen, M. C., Black, F., and Scholes, M. S. (1972). The capital asset pricing model: Some empirical tests.
- 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.

- Lettau, M., Maggiori, M., and Weber, M. (2014). Conditional risk premia in currency markets and other asset classes. *Journal of Financial Economics*, 114(2):197–225.
- 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.
- Manela, A. and Moreira, A. (2013). News implied volatility and disaster concerns. *Available at SSRN 2382197*.
- Merton, R. C. (1973). An intertemporal asset pricing model. *Econometrica*, 41(5):867–887.
- Muir, T. (2013). Financial crises, risk premia, and the term structure of risky assets. *Yale School of Management Working Paper*.
- Nagel, S., Reck, D., Hoopes, J., Langetieg, P., Slemrod, J., and Stuart, B. (2016). Who sold during the crash of 2008-9? evidence from tax return data on daily sales of stock. *working paper*.
- Nakamura, E., Steinsson, J., Barro, R., and Ursúa, J. (2010). Crises and recoveries in an empirical model of consumption disasters. Technical report, National Bureau of Economic Research.
- Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1):1–28.
- Poterba, J. M. and Summers, L. H. (1988). Mean reversion in stock prices: Evidence and implications. *Journal of financial economics*, 22(1):27–59.
- Tang, Y. and Whitelaw, R. F. (2011). Time-varying sharpe ratios and market timing. *The Quarterly Journal of Finance*, 1(03):465–493.
- Veronesi, P. (2000). How does information quality affect stock returns? *The Journal of Finance*, 55(2):807–837.
- Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, forthcoming.
- Whitelaw, R. F. (1994). Time variations and covariations in the expectation and volatility of stock market returns. *The Journal of Finance*, 49(2):515–541.

7. Tables and Figures

Table 1: Volatility managed factor alphas. We run time-series regressions of each volatility managed factor on the non-managed factor $f_t^\sigma = \alpha + \beta f_t + \varepsilon_t$. The managed factor, f_t^σ , scales by the factors inverse realized variance in the preceding month $f_t^\sigma = \frac{c}{RV_{t-1}^2} f_t$. 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. All factors are annualized in percent per year by multiplying monthly factors by 12.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Mkt $^\sigma$	SMB $^\sigma$	HML $^\sigma$	Mom $^\sigma$	RMW $^\sigma$	CMA $^\sigma$	FX $^\sigma$	ROE $^\sigma$	IA $^\sigma$
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)			
Carry							0.71 (0.08)		
ROE								0.63 (0.07)	
IA									0.68 (0.05)
Alpha (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)	2.78 (1.49)	5.48 (0.97)	1.55 (0.67)
N	1,065	1,065	1,065	1,060	621	621	360	575	575
R ²	0.37	0.38	0.32	0.22	0.38	0.46	0.33	0.40	0.47
rmse	51.39	30.44	34.92	50.37	20.16	17.55	25.34	23.69	16.58

Table 2: Mean-variance efficient factors. We form unconditional mean-variance efficient (MVE) portfolios using various combinations of factors. These underlying factors can be thought of as the relevant information set for a given investor (e.g., an investor who only has the market available, or a sophisticated investor who also has value and momentum available). We then volatility time each of these mean-variance efficient portfolios and report alphas of regressing the volatility managed portfolio on the original MVE portfolio. The volatility managed portfolio simply scales the portfolio by the inverse of realized variance in the previous month, reducing exposure when variance is high and vice versa. We also report the annualized Sharpe ratio of the original MVE portfolio and the appraisal ratio of the volatility timed MVE portfolio, which tells us directly how much the volatility managed portfolio increases the investors Sharpe ratio relative to no volatility timing. The factors considered are the Fama-French three and five factor models, the momentum factor, and the Hou, Xue, and Zhang (2015) four factors (HXZ). Panel B reports the alphas of these mean-variance efficient combinations in subsamples where we split the data into three thirty year periods. Note some factors are not available in the early sample.

Panel A: Mean Variance Efficient Portfolios (Full Sample)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	FF3	FF3 Mom	FF5	FF5 Mom	HXZ	HXZ Mom
Alpha (α)	4.86 (1.56)	4.99 (1.00)	4.04 (0.57)	1.34 (0.32)	2.01 (0.39)	2.32 (0.38)	2.51 (0.44)
Observations	1,065	1,065	1,060	621	621	575	575
R-squared	0.37	0.22	0.25	0.42	0.40	0.46	0.43
rmse	51.39	34.50	20.27	8.28	9.11	8.80	9.55
Original Sharpe	0.42	0.69	1.09	1.20	1.42	1.69	1.73
Appraisal Ratio	0.33	0.50	0.69	0.56	0.77	0.91	0.91

Panel B: Subsample Analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	FF3	FF3 Mom	FF5	FF5 Mom	HXZ	HXZ Mom
α : 1926-1955	8.11 (3.09)	1.94 (0.92)	2.45 (0.74)				
α : 1956-1985	2.06 (2.82)	0.99 (1.43)	2.54 (1.16)	0.13 (0.43)	0.71 (0.67)	0.77 (0.39)	1.00 (0.51)
α : 1986-2015	4.22 (1.66)	5.66 (1.74)	4.98 (0.95)	40 1.88 (0.41)	2.65 (0.47)	3.03 (0.50)	3.24 (0.57)

Table 3: Persistence of volatility. We show that volatility is persistent for each individual factor by running a regression of realized volatility on lagged realized volatility. Realized 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 FX Carry factor which starts in 1983. MVE is the unconditional mean-variance efficient portfolio formed using the Fama-French three factors plus momentum. Standard errors are in parentheses and adjust for heteroscedasticity.

Variables	(1) $\ln RV_t^{Mkt}$	(2) $\ln RV_t^{SMB}$	(3) $\ln RV_t^{HML}$	(4) $\ln RV_t^{Mom}$	(5) $\ln RV_t^{RMW}$	(6) $\ln RV_t^{CMA}$	(7) $\ln RV_t^{MVE}$	(8) $\ln RV_t^{FX}$
$\ln RV_{t-1}^{Mkt}$	0.72 (0.02)							
$\ln RV_{t-1}^{SMB}$		0.78 (0.02)						
$\ln RV_{t-1}^{HML}$			0.78 (0.02)					
$\ln RV_{t-1}^{Mom}$				0.76 (0.02)				
$\ln RV_{t-1}^{RMW}$					0.78 (0.03)			
$\ln RV_{t-1}^{CMA}$						0.74 (0.03)		
$\ln RV_{t-1}^{MVE}$							0.69 (0.02)	
$\ln RV_{t-1}^{FX}$								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 4: Risk return tradeoff for each factor. 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 by asking whether increased volatility forecasts higher future returns. 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. MVE is the unconditional mean variance efficient portfolio formed using the Fama-French three factors plus momentum. Standard errors are in parentheses and adjust for heteroscedasticity.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	R_{t+1}^{Mkt}	R_{t+1}^{SMB}	R_{t+1}^{HML}	R_{t+1}^{Mom}	R_{t+1}^{RMW}	R_{t+1}^{CMA}	R_{t+1}^{MVE}	R_{t+1}^{FX}	R_{t+1}^{ROE}	R_{t+1}^{IA}
RV_t^{Mkt}	-0.04 (0.44)									
RV_t^{SMB}		0.38 (0.38)								
RV_t^{HML}			0.70 (0.58)							
RV_t^{Mom}				-1.10 (0.59)						
RV_t^{RMW}					0.87 (0.83)					
RV_t^{CMA}						1.19 (0.62)				
RV_t^{MVE}							-0.02 (0.35)			
RV_t^{FX}								-0.10 (0.67)		
RV_t^{ROE}									-0.90 (0.76)	
RV_t^{IA}										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 5: Recession betas by factor. We regress each scaled factor on the original factor and we include recession dummies $1_{rec,t}$ using NBER recessions which we interact with the original factors; $f_t^\sigma = \alpha_0 + \alpha_1 1_{rec,t} + \beta_0 f_t + \beta_1 1_{rec,t} \times f_t + \varepsilon_t$. This gives the relative beta of the scaled factor conditional on recessions compared to the unconditional estimate. Standard errors are in parentheses and adjust for heteroscedasticity. We find that $\beta_1 < 0$ so that betas for each factor are relatively lower in recessions.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Mkt $^\sigma$	HML $^\sigma$	Mom $^\sigma$	RMW $^\sigma$	CMA $^\sigma$	FX $^\sigma$	ROE $^\sigma$	IA $^\sigma$
MktRF	0.83 (0.08)							
MktRF $\times 1_{rec}$	-0.51 (0.10)							
HML		0.73 (0.06)						
HML $\times 1_{rec}$		-0.43 (0.11)						
Mom			0.74 (0.06)					
Mom $\times 1_{rec}$			-0.53 (0.08)					
RMW				0.63 (0.10)				
RMW $\times 1_{rec}$				-0.08 (0.13)				
CMA					0.77 (0.06)			
CMA $\times 1_{rec}$					-0.41 (0.07)			
Carry						0.73 (0.09)		
Carry $\times 1_{rec}$						-0.26 (0.15)		
ROE							0.74 (0.08)	
ROE $\times 1_{rec}$							-0.42 (0.11)	
IA								0.77 (0.07)
IA $\times 1_{rec}$								-0.39 (0.08)
Observations	1,065	1,065	1,060	621	621	362	575	575
R-squared	0.43	0.37	0.29	0.43	0.38	0.49	0.51	0.43
				0.38				0.49

Table 6: Time-series alphas controlling for risk parity factors. We run time-series regressions of each managed factor on the non-managed factor plus a risk parity factor based on Asness et al. (2012). The risk parity factor is given by $RP_{t+1} = b_t' f_{t+1}$ where $b_{i,t} = \frac{1/RV_t^i}{\sum_i 1/RV_t^i}$ and f is a vector of pricing factors. Volatility is measured on a rolling three year basis following Asness et al. (2012). We construct this risk parity portfolio for various combinations of factors. We then regress our volatility managed MVE portfolios from Table 2 on both the static MVE portfolio and the risk parity portfolio formed using the same factors, f , that make up the MVE portfolio. We find our alphas are unchanged from those found in the main text. In the last column, we show the alpha for the volatility managed betting against beta (BAB) portfolio to highlight that our time-series volatility timing is different from cross-sectional low risk anomalies. 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)
	Mkt	FF3	FF3 Mom	FF5	FF5 Mom	HXZ	HXZ Mom	BAB^σ
Alpha (α)	4.86 (1.56)	5.00 (1.00)	4.09 (0.57)	1.32 (0.31)	1.97 (0.40)	2.03 (0.32)	2.38 (0.44)	5.67 (0.98)
Observations	1,065	1,065	1,060	621	621	575	575	996
R-squared	0.37	0.23	0.26	0.42	0.40	0.50	0.44	0.33
rmse	51.39	34.30	20.25	8.279	9.108	8.497	9.455	29.73

Table 7: Leverage constraints and volatility timing. We study whether leverage is necessary to implement our volatility managed portfolios. We consider strategies that use embedded leverage in place of actual leverage for the market portfolio. Specifically, we look at investing in a portfolio of options on the S&P500 index. The portfolio is an equal-weighted average of 6 in the money call options with maturities of 60 and 90 days and moneyness of 90, 92.5, and 95. The beta of this portfolio is about 7. Thus, any time our strategy prescribes leverage to achieve high beta, we invest in this option portfolio to achieve our desired beta rather than using actual leverage. We then compare the no leverage volatility timed portfolio to the standard volatility managed portfolio studied in the main text in terms of Sharpe ratio, alpha, and appraisal ratio. Finally, we consider another option strategy that also sells in the money puts as well as buys calls to again achieve our desired beta. The puts we use mirror the moneyness of the calls and again have 60 and 90 day maturities. We take an equal weighted average across these 12 portfolios. The sample used for all numbers in the table is April, 1986 to January 2012 based on data availability and we only study the volatility timing of the market portfolio. Our option data comes from Constantinides et al. (2013).

Volatility Timing and Leverage Constraints				
	Buy and hold	Alternative volatility managed strategies		
		With leverage	No leverage (calls)	No leverage (calls + puts)
Sharpe Ratio	0.39	0.59	0.54	0.60
α	–	4.03	5.90	6.67
$s.e.(\alpha)$	–	(1.81)	(3.01)	(2.86)
β	–	0.53	0.59	0.59
Appraisal Ratio	–	0.44	0.39	0.46

Table 8: Transaction costs of volatility timing. We evaluate our volatility timing strategy for the market portfolio when including transaction costs. We consider alternative strategies that still capture the idea of volatility timing but significantly reduce trading activity implied by our strategy. Specifically, we consider using inverse volatility instead of inverse variance, using expected rather than realized variance, and only trading when variance is above its long run average. For expected variance, we run an AR(1) for log variance to form our forecast. We report the average absolute change in monthly weights ($|\Delta w|$), expected return, and alpha of each of these alternative strategies. 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
$\frac{1}{RV_t^2}$	Realized Variance	0.73	9.47%	4.86%	4.77%	3.98%	3.63%	56bps
$\frac{1}{RV_t}$	Realized Vol	0.38	9.84%	3.85%	3.80%	3.39%	3.21%	84bps
$\frac{1}{E_t[RV_{t+1}^2]}$	Expected Variance	0.37	9.47%	3.30%	3.26%	2.86%	2.68%	74bps
$\frac{1}{\max(E[RV_t^2], RV_t^2)}$	RV Above Mean	0.10	9.10%	2.20%	2.19%	2.08%	2.03%	183bps

Table 9: Calibration for portfolio choice exercise. This Table shows targeted moments in the data and in the three alternative models analyzed in Figure 7. Model based quantities are found by simulating monthly data for 100 different 100 year histories. We report the median of each moment across these 100 histories. RV_t represents realized standard deviation and time is given in months. The seven reported moments capture: average volatility, pe-year-ahead predictability of stock returns, auto-correlation of log variance at the monthly frequency, one-year ahead predictability of returns using realized variance, the volatility of log variance, the correlation between returns and the innovation in variance, and the persistence of the price-dividend ratio.

Moment	Data	(a)	(b)	(c)
$E[RV_{t+1}]\sqrt{12}$	0.160	0.180	0.235	0.164
$corr(pd_t, R_{t \rightarrow t+12})^2$	0.060	0.041	0.020	0.068
$corr(\log(RV_{t+1}), \log(RV_t))^2$	0.500	0.397	0.027	0.506
$corr(R_{t \rightarrow t+12}, RV_t^2)^2$	0.000	0.008	-0.005	0.001
$stdev(\log(RV_t^2))$	1.050	0.665	0.369	0.922
$corr(R_t, RV_t - RV_{t-1})$	-0.240	0.001	0.001	0.003
$corr(pd_{t+12}, pd_t)$	0.90	0.897	0.892	0.890

Figure 1: Sorts on previous month's volatility. We use the monthly time-series of realized volatility to sort the following month's returns into five buckets. The lowest, "low vol," looks at the properties of returns over the month *following* the lowest 20% of realized volatility months. We show the average return over the next month, the standard deviation over the next month, and the average return divided by variance. Average return per unit of variance represents the optimal risk exposure of a mean variance investor in partial equilibrium, and also represents "effective risk aversion" from a general equilibrium perspective (i.e., the implied risk aversion, γ_t , of a representative agent needed to satisfy $E_t[R_{t+1}] = \gamma_t \sigma_t^2$). The last panel shows the probability of a recession across volatility buckets by computing the average of an NBER recession dummy. Our sorts should be viewed analogous to standard cross-sectional sorts (i.e., book-to-market sorts) but are instead done in the time-series using the past months realized volatility.

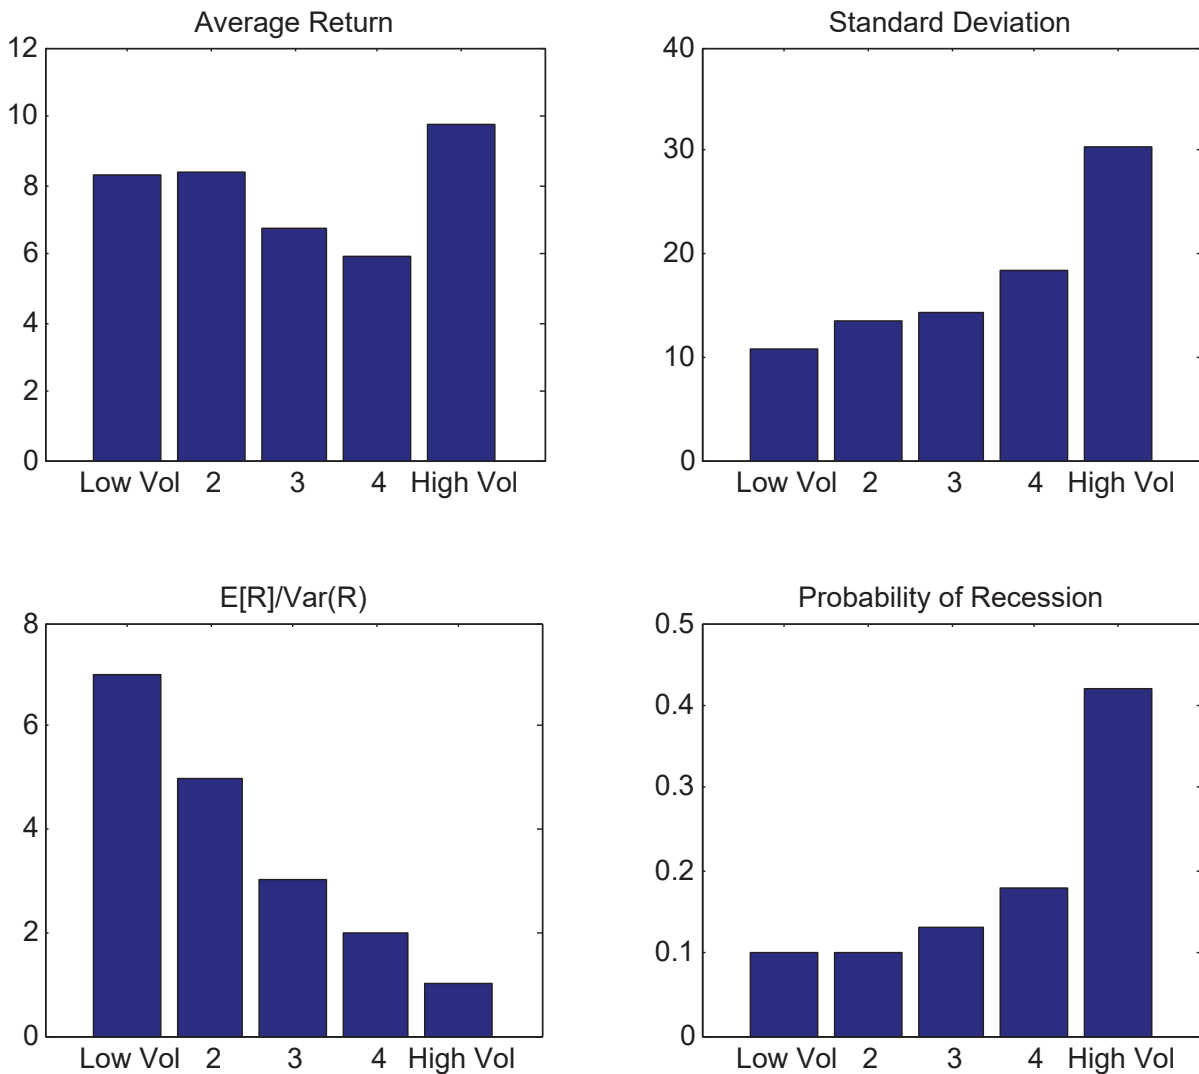


Figure 2: Time-series of volatility by factor. This figure plots the time-series of the monthly volatility of each individual factor. We emphasize the common co-movement in volatility across factors and that volatility generally increases for all factors in recessions. Light shaded bars indicate recessions and show a clear business cycle pattern in volatility.




Figure 3: Cumulative returns to volatility timing for the market return. The top panel 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. The lower left panel plots rolling one year returns from each strategy and the lower right panel shows the drawdown of each strategy.

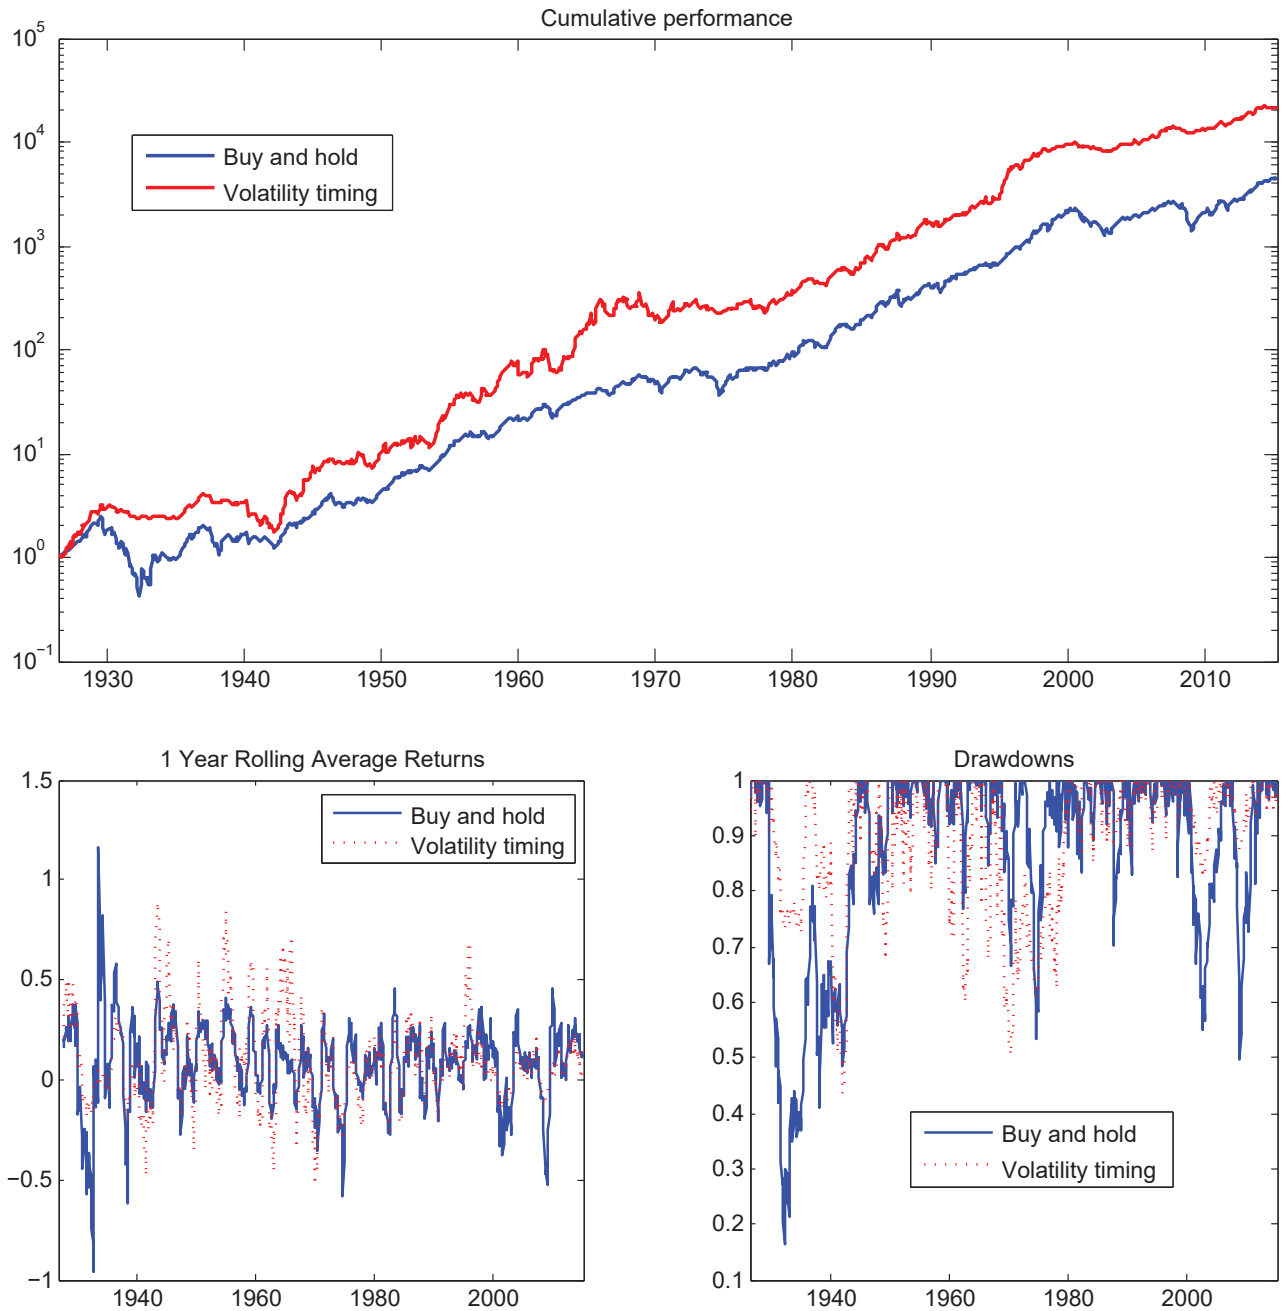


Figure 4: Dynamics of the risk return tradeoff. The figure plots the impulse response of the expected return and variance of the market portfolio for a shock to the variance. The x-axis is months. The last panel gives the portfolio choice implications for the differential movements in expected returns 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 and whose demand is proportional to $E_t[R_{t+1}]/var_t[R_{t+1}]$. 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 described in the text which uses 3 lags of log variance. We compute impulse responses using a VAR with 3 monthly lags of each variable.

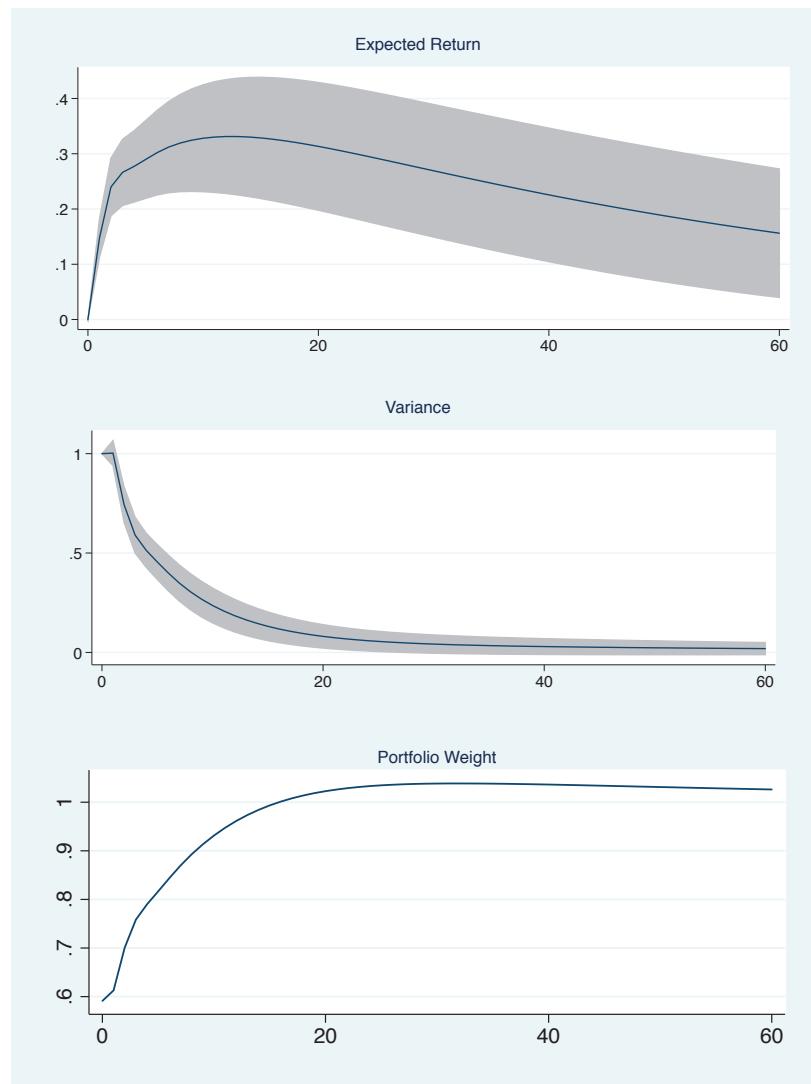


Figure 5: Volatility timing performance during large market downturns. The figure plots the performance of our volatility timing strategy compared to a buy-and-hold strategy for the market return during specific episodes of market turmoil where there was also a large stock market drop.

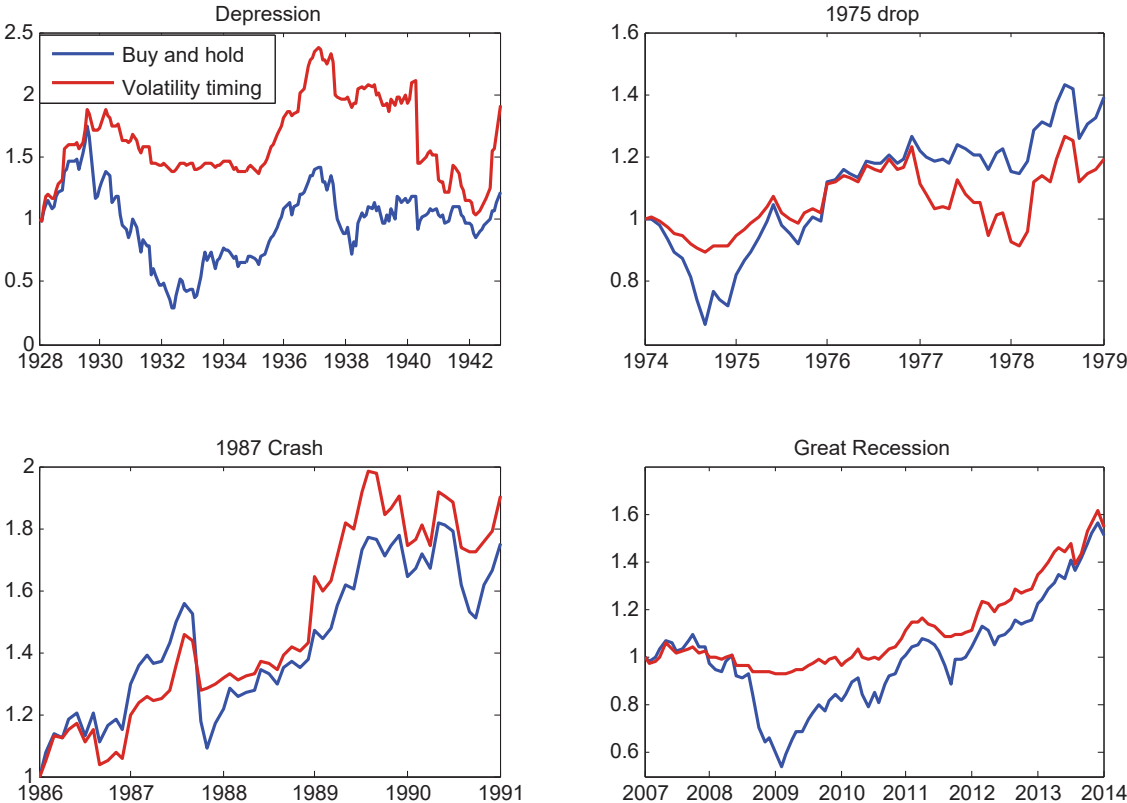



Figure 6: Intuition for why changes in the volatility of mean-reverting shocks matters less for long horizon investors. The figure gives a stylized description of an increase in the volatility of a mean-reverting shock. For illustration, we consider a shock to returns that is fully mean-reverting the next period. Thus, a negative return next period is perfectly offset by a positive return of equal magnitude the following period, and vice versa. Regardless of the volatility of this mean-reverting shock, 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.

Panel A:



Panel B:

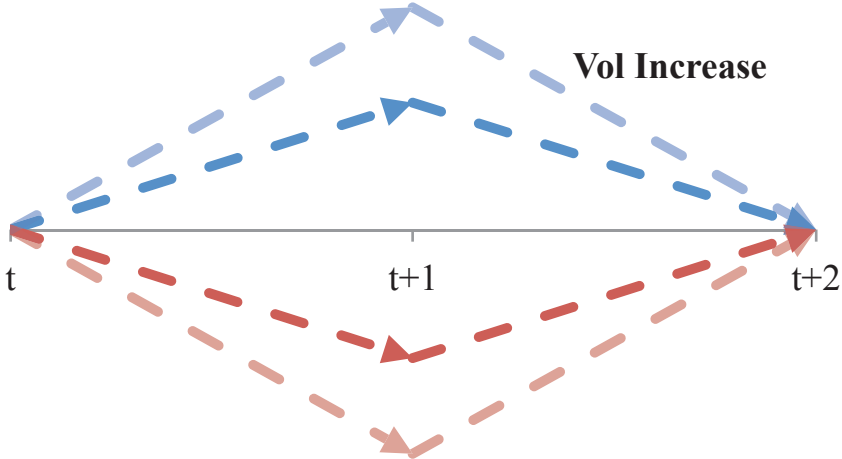



Figure 7: Optimal volatility timing by investor horizon. The Figure show the coefficient b of a linear approximations to the optimal portfolio policy for calibrations (a)-(c) in Table 9. Specifically, we estimate Equation 15 by running regressions (in data from numerical simulations) of the true optimal policy evaluated at the unconditional expected return, $w(\mu_x, \sigma_t^2)$, on $\frac{\mu_x}{\gamma\sigma_t^2}$ as we vary volatility $\sigma_t^2 = D'Dy_t + F'F$. A short-horizon mean-variance investor will have a coefficient of $b = 1$ so b measures the relative extent of volatility timing compared to a short-horizon mean-variance investor. In the blue line (calibration (a)) all variation in volatility is driven by mean-reverting shocks, in the red line (calibration (b)) it is driven only by permanent shocks, and the yellow line (calibration (c)) is the case where the share of permanent and mean-reverting shocks is constant.



• **Appendix: Not intended for publication**

A. Additional empirical results

This subsection performs various robustness checks of our main result. A reader who is less concerned with the robustness of our main fact can skip this subsection.

A.1 Using expected variance in place of realized variance

Table 10 shows the results when, instead of scaling by past realized variance, we scale by the expected variance from our forecasting regressions where we use three lags of realized log variance to form our forecast. This offers more precision but comes at the cost of assuming that 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 has a clear appeal from the perspective of practical implementation. Other variance forecasting methods behave similarly, e.g., Andersen and Bollerslev (1998).

A.2 International data

As an additional robustness check, we show that our results hold for the stock market indices of 20 OECD countries. 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. The volatility managed index has a higher Sharpe ratio than the passive strategy in 80% of cases. These results are detailed in Figure 8 of our Appendix. Note that this is a strong condition – a portfolio can have positive alpha even when its Sharpe ratio is below that of the non-managed factor.

A.3 Betting against beta controls

Table 11 gives the alphas of our volatility managed factors when we include the BAB factor of Frazzini and Pedersen. As we can see from the Table, the results are identical to those in the main text. Moreover, the BAB factor does not appear significant – meaning it is not strongly correlated with our volatility managed portfolios. This again highlights that our strategy is quite different from this cross-sectional low risk anomaly.

A.4 Multivariate analysis

In this section we study whether some of the single-factor volatility timing strategies are priced by other aggregate factors. Consistent with Table 2, Tables 12 and 13 show that the scaled factors expand the mean variance frontier of the existing factors because the

appraisal ratio of HML, RMW, Mom are positive and large when including all factors. Notably, the alpha for the scaled market portfolio is reduced when including all other factors. Thus, the other asset pricing factors, specifically momentum, contain some of the pricing information of the scaled market portfolio. For an investor who only has the market portfolio available, the univariate results are the appropriate benchmark; in this case, the volatility managed market portfolio does have large alpha. For the multivariate results (i.e., for an investor who has access to all factors) the relevant benchmark is the MVE portfolio, or “tangency portfolio”, since this is the portfolio investors with access to these factors will hold (within the set of static portfolios). We find that the volatility managed version of each of the different mean variance efficient portfolios has a substantially higher Sharpe ratio and large positive alpha with respect to the static factors.

A.5 An alternative performance measure and simulation exercises

So far, we have focused on time-series alphas, Sharpe ratios, and appraisal ratios as our benchmark for performance evaluation. This section considers alternative measures and discusses some statistical concerns. We also conduct simulations to better evaluate our results.

In our simulations, we consider a world where the price of risk is constant $E_t[R_{t+1}] = \gamma Var_t[R_{t+1}]$ and choose parameters to match the average equity premium, average market standard deviation, and the volatility of the market standard deviation. We model volatility as lognormal and returns as conditionally lognormal. Using these simulations we can ask, if the null were true that the risk return tradeoff is strong, what is the probability we would see the empirical patterns we document in the data (alphas, Sharpe ratios, etc.).

First, we study the manipulation proof measure of performance (henceforth MPPM) from Goetzmann et al. (2007). This measure is useful because, unlike alphas and Sharpe ratios, it can’t be manipulated to produce artificially high performance. This manipulation could be done intentionally by a manager, say by decreasing risk exposure if they had experienced a string of lucky returns, or through a type of strategy that uses highly non-linear payoffs. Essentially, the measure is based on the certainty equivalent for a power utility agent with risk aversion ranging from 2 to 4 and evaluates their utility directly. We choose risk aversion of 3, although our results aren’t sensitive to this value. We find the market MPPM to be 2.48% and the volatility managed market portfolio MPPM to be 4.33%, so that the difference between the two is 1.85% per year. This demonstrates that even under this alternative test which overcomes many of the potential shortcomings of traditional performance measure, we find our volatility managed strategy beats the buy and hold portfolio.

It is useful to consider the likelihood of this finding in relation to the null hypothesis that the price of risk is constant. In our simulations, we can compute the MPPM measure of the scaled market portfolio and compare it to the market portfolio MPPM. We find that the volatility managed MPPM beats the market MPPM measure only 0.2% of the time.

Hence, if the price of risk isn't moving with volatility it is highly unlikely that the MPPM measure would favor the volatility managed portfolio. Using these simulations, we can also ask the likelihood we would observe an alpha as high as we see in the data. The median alpha in our simulations with a constant price of risk is about 10 bps and the chance of seeing an alpha as high as we see empirically (4.86%) is essentially zero.

A.6 Are volatility managed portfolios option like?

At least since Black and Scholes (1973), it is well known that under some conditions option payoffs can be replicated by dynamically trading the reference asset. Since our strategy is dynamic, a plausible concern is that our strategy might be replicating option payoffs. A large literature discusses potential issues with evaluating strategies that have a strong option like return profile.

We discuss each of the potential concerns and explain why it does not apply to our volatility managed portfolios. First, a linear asset pricing factor model where a return is a factor implies a stochastic discount factor that can be negative for sufficiently high factor return realizations (Dybvig and Ingersoll Jr, 1982). Thus, there are states with a negative state price, which implies an arbitrage opportunity. A concern is that our strategy may be generating alpha by implicitly selling these negative state-price states. However, empirically this cannot be the source of our strategy alpha, as the implied stochastic discount factor is always positive in our sample.²⁸

Second, the non-linearity of option like payoffs can make the estimation of our strategy's beta challenging. Because some events only happen with very low probability, sample moments are potentially very different from population moments. This concern is much more important for short samples. For example, most option and hedge fund strategies for which such biases are shown to be important have no more than 20 years of data; on the other hand we have 90 years of data for the market portfolio. In Figure 9 we also look at kernel estimates of the buy-and-hold and volatility managed factor return distributions. No clear pattern emerges; if anything, the volatility managed portfolio appears to have less mass on the left tail for some portfolios.

Third, another concern is that our strategy loads on high price of risk states; for example, strategies that implicitly or explicitly sell deep out of the money puts can capture the expected return resulting from the strong smirk in the implied volatility curve. Note that our strategy reduces risk exposure after a volatility spike, which is typically associated with low return realizations, while one would need to increase exposure following a low return realization to replicate the sale of a put option. Mechanically our strategy does exactly the opposite of what a put selling strategy would call for. This also implies that our strategy will typically have less severe drawdowns than the static portfolio, which

²⁸For example, for the market factor the implied SDF can be written as $\approx 1/R_t^f - b(R_{t+1}^m - R_t^f)$, where empirically $b = E[R_{t+1}^m - R_t^f] / \text{Var}(R_t^e) \approx 2$. In our sample the highest return realization is 38% so that the SDF is never negative.

accords with our Figure 3.

Another more general way of addressing the concern that our strategy's alpha is due to its option-like returns is to use the manipulation proof measure of performance (MPPM) proposed in Goetzmann et al. (2007). We find that the volatility managed MPPM is 75% higher than the market MPPM. Using simulations we show that a volatility managed portfolio would beat the market (as measured by MPPM) only 0.2% of the time if the risk-return trade-off was constant. This is again another piece of evidence that our strategy increases Sharpe ratios by simply avoiding high risk times and does not load on other unwanted risks.

Overall, there is no evidence that our volatility managed portfolios generate option-like returns.

B. Unconditional alpha as a measure of the weakness of the risk return trade-off

We start by considering the possibility that there is a strong relation between risk and return in the time series, so that $\mu_t = \gamma\sigma_t^2$. This means increases in variance are compensated proportionally by an increase in expected returns as would happen naturally in a general equilibrium model with a representative agent with risk aversion γ . Notice in this case a mean-variance investor would not want to time volatility, but would simply keep a constant weight in the risky asset. In this case, our strategy which sets $z_t = 1/\sigma_t^2$ will tradeoff two aspects. On one hand $E[\beta_t] > \beta$, making the first term positive. This captures the reduction in unconditional risk due to volatility timing. This comes from the fact that our strategy increases β_t when variance is low and decreases β_t when variance is high, making the average conditional beta much higher than the unconditional beta. However, in a world where the risk-return trade-off is constant, the reduction in risk comes at the expense of lower expected returns, because $cov(\beta_t, \mu_t) = cov(1/\sigma_t^2, \mu_t) = \gamma cov(1/\sigma_t^2, \sigma_t^2) = < 0$. Volatility timing in this case reduces risk exposure exactly when expected returns are greatest. This lowers our unconditional alpha. We find that when the risk-return trade-off is constant, the unconditional alpha totals to around 10 basis points per annum as the two effects essentially offset each other. More specifically, one can show that the unconditional alpha is $\mu (var(R)/E[\sigma_t^2] - 1)$. This residual term arises because the unconditional model uses unconditional variance rather than expected conditional variance, but this difference turns out to be empirically tiny. Thus, if the risk-return tradeoff is strong, our strategy would not produce significant risk adjusted returns.

In the second case we consider there is no relation between risk and return in the time-series. In particular, suppose there is no risk return tradeoff so that $cov(1/\sigma_t^2, \mu_t) \approx 0$ and variance is roughly uncorrelated with expected returns. Then the alpha becomes $\alpha = \mu (E[\beta_t] - \beta)$. In other words, the unconditional alpha stems completely from volatility timing – it increases risk exposure when volatility is low and the risk return tradeoff is favorable and vice versa. For our strategy, we find $E[\beta_t] - \beta = 0.6$. This spread in

average conditional betas and the unconditional betas captures how well our strategy can volatility time. The more volatility our proxy for expected volatility can predict, the larger this spread. With the sample mean of the market excess returns $\mu = 0.08$ and our strategy beta profile for the market portfolio, this formula gives an $\alpha = 4.8\%$.

Our empirical estimate for the volatility managed market α is identical to this number at 4.86%. Thus our result is consistent with a lack of risk return tradeoff. Our strategy would not have produced a large positive alpha if the risk return tradeoff were indeed strong. This is also confirmed in simulations we run where we build in a strong relation between risk and return. In these cases, the unconditional alpha is always near zero and the chances of seeing an alpha as large as we see in the data are very small, occurring about 0.1% of the time. Importantly equation (2) makes it transparent that what is required for our strategy to produced alpha is that the relationship between risk and return is weak enough, not necessarily zero.

C. Portfolio choice for a mean-variance investor

We expand on the analysis in Section 4.1 and 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,

$$\begin{aligned} w(Y_t, x_t) &= \frac{1}{\gamma} \frac{\mu + x_t}{Y_t}; \\ E[w(Y_t, x_t)r_{t+1}] &= \frac{1}{\gamma} \left(\frac{S^2 + R_r^2}{1 - R_r^2} E[Y_t]E[Y_t^{-1}] + 2\mu \text{cov}(x_t, Y_t^{-1}) + \text{cov}(x_t^2, Y_t^{-1}) \right), \end{aligned} \quad (\text{C.0.17})$$

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

The second term we can construct directly from our estimates in Table 4, using that $\text{cov}(x_t, Y_t^{-1}) = -\beta \text{Var}(Y_t^{-1})$. 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

$$\text{cov}(x_t^2, Y_t^{-1}) \geq -1\sigma(x_t^2)\sigma(Y_t^{-1}) = -1\sqrt{2}\sigma_x^2\sigma(Y_t^{-1}), \quad (\text{C.0.18})$$

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

obtain,

$$E[w(Y_t, x_t)r_{t+1}] \geq \frac{1}{\gamma} \left(\frac{S^2 + R_r^2}{1 - R_r^2} E[Y_t]E[Y_t^{-1}] - 2\beta\mu\sigma^2(Y_t^{-1}) - \sqrt{2}\sigma_x^2\sigma(Y_t^{-1}) \right). \quad (\text{C.0.19})$$

Plugging numbers for σ_x consistent with a monthly R_r^2 of 0.43%, and $\sigma(Y_t^{-1})$ and β as implied by the variance model that uses only a lag of realized variance ($R_y^2 = 53\%$), 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.

D. Portfolio choice for a long-horizon investor

In this section we describe in detail the portfolio problem of a long-lived investor.

D.1 Investment opportunity set

The investment opportunity set is described by Equations (10), (11), and (12). The volatility process follows a CIR process as in Heston (1993), 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 shocks to expected returns and volatility, 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 shocks to expected returns only have a transitory effect on asset prices implies $G[3] = -D[3]\kappa_x$ and $H[3] = -F[3]\kappa_x$. The initial drop (increase) in prices is exactly compensated with an increase (decrease) in future expected returns. In this precise sense these shocks are mean-reverting shocks. The vector L has the first entry equal to zero. The second entry captures pure volatility shocks that are contemporaneously unrelated to mean-reverting shocks (i.e. shocks to expected returns x), 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 shocks to prices. In the first entry we have the exposure to shocks unrelated to volatility or expected returns, and in the second entry shocks related to volatility. In the third entry we have the shocks to mean-reverting shocks. Shocks that by construction mean-revert in the long run.

D.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:

$$J = E_t \left[\int_t^\infty f(C_s, J_s) ds \right] \quad (D.2.20)$$

$$f(C, J) = \rho \frac{1-\gamma}{1-\psi^{-1}} J \times \left[\left(\frac{C}{((1-\gamma)J)^{\frac{1}{1-\gamma}}} \right)^{1-\psi^{-1}} - 1 \right] \quad (D.2.21)$$

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

$$f(C, J) = \rho(1-\gamma)J \times \left[\log(C) - \frac{\log((1-\gamma)J)}{1-\gamma} \right]. \quad (D.2.22)$$

The investor maximizes D.2.20 subject to his intertemporal budget constraint 13 and the evolution of state variables x_t and y_t .

D.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

$$\begin{aligned} 0 &= \sup_{w,C} f(C_t, J_t) + [w_t x_t W_t + r W_t - C_t] J_W + \frac{1}{2} w_t^2 W_t^2 J_{WW} (y_t D' D + F' F) \\ &+ \kappa_y (\mu_y - y_t) J_y + \frac{1}{2} Y_t J_{yy} L' L + \kappa_x (\mu_x - x_t) J_x + \frac{1}{2} J_{xx} (y_t G' G + H' H) \\ &+ y_t J_{xy} G' L + w_t W_t J_{xW} (G' D y_t + F' H) + w_t W_t J_{yW} L' D Y_t \end{aligned} \quad (D.3.23)$$

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,

$$w_t = \frac{x_t}{\gamma (D' D y_t + F' F)} + V_x \frac{G' D y_t + F' H}{\gamma (D' D y_t + F' F)} + V_y \frac{L' D y_t}{\gamma (D' D y_t + F' F)}. \quad (D.3.24)$$

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.

D.4 Analysis

We study CRRA preferences with risk aversion of 5 and 10, and Esptein 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-life²⁹ of utility weights ranges from 5 year to 30 years.

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 16 (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-squared across specifications in a given row where the row denotes investor horizon. Overall, we see that the linear relation fits the optimal policy extremely well.

Table 16(a) shows the case where all time-variation in volatility is due to mean-reverting shocks, i.e. increases in volatility are associated with greater-than-proportional increases in the amount of mean-reversion. The linear policy is an almost perfect description of

²⁹Specifically we map horizon into ρ as follows: For a given horizon T_j , look for ρ_j such that $\frac{\int_0^{T_j} e^{-\rho_j \times t} dt}{\int_0^{\infty} e^{-\rho_j \times t} dt} =$

the optimal policy (the R-squared is close to one). The long-horizon investors respond less aggressively than a short-horizon investor to changes in volatility (weight on the volatility managed portfolio is less than one). Nevertheless they still trade quite a bit. Coefficients are always quite high, implying that the dynamic investor invests quite a bit on the volatility managed portfolio. Even though shocks eventually mean-revert, they take a long time. The long time mean-reversion takes make mean-reverting shocks still somewhat risky to investors with fairly long (but not extreme) investment horizons.

Table 16(b) shows the case where all time-variation in volatility is driven by permanent shocks and increases in volatility are associated with less-than-proportional increases in the amount of mean-reversion. The long-horizon investor responds more aggressively than the mean-variance investor (weight on the volatility managed portfolio is higher than one). Intuitively, when volatility increases, the proportion of mean-reverting shocks decreases, making stocks relatively riskier for a long-horizon investor. Thus, the long-horizon investor responds by reducing his portfolio allocation more than proportionally with the increase in variance.

Table 16(c) shows the case where the proportion of mean-reverting shocks is constant. This is a natural benchmark which is typically assumed in the literature (Campbell and Vuolteenaho, 2004). Furthermore, we are not aware of any work that measures how the amount of mean-reversion co-moves with volatility. 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. Note that the presence of mean-reverting 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 short-horizon investor behavior.

In Tables 17 (d) and (e) we study the effect of introducing contemporaneous correlation between return realizations and volatility innovations. Chacko and Viceira (2005) studies 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 17 (d) is the case where the correlation between returns and volatility is entirely due to the mean-reverting component. That is volatility and expected returns shock are positively related in this case. Table 17 (e) analyzes the other extreme where the correlation between returns and volatility is due to the permanent component. In both cases we see that the introduction of this correlation did not meaningfully change how investors respond to changes in volatility. As it was the case for calibration (c) the long-horizon investor responds roughly in the same fashion as the short-horizon investors

E. The dynamics of the price of risk

We now show that the fact that our volatility managed portfolios generate risk-adjusted returns relative to original factors, implies that the price of risk, $b_t = \mu_t/\sigma_t^2$, has a negative co-variance with volatility. That is, risk premia increases less than proportionally with stock market variance. Formally this can be seen in the following result

Proposition 1 (Alpha and the dynamics of the price of risk). *Let μ_t be the factor expected excess return, σ_t^2 the factor conditional variance, α be the intercept of a time-series regression of the volatility managed portfolio returns on the buy-and-hold returns, $\beta_t = \frac{1}{\sigma_t^2 \zeta_t}$ the strategy conditional beta, where ζ_t is measurement error with $E[\zeta_t] = 1$, and β the unconditional strategy beta, then the co-variance between the price of risk $b_t = \mu_t/\sigma_t^2$ and factor variance σ_t^2 is*

$$\text{cov}(b_t, \sigma_t^2) = - \left(\frac{\alpha - E[\mu_t](E[\beta_t] - \beta)}{E[\beta_t]} + \frac{E[b_t]}{E[1/\sigma_t^2]} \left(E[\sigma_t^2]E[1/\sigma_t^2] - 1 \right) \right), \quad (\text{E.0.25})$$

and it follows that the co-variance between the price and quantity of risk $\text{cov}(b_t, \sigma_t^2)$ is decreasing in the volatility managed portfolio alpha.

Proof. This follows immediately from the definition of unconditional alpha,

$$\alpha = E[\beta_t \mu_t] - \beta E[\mu_t]. \quad (\text{E.0.26})$$

Using that $E[\beta_t \mu_t] = \text{cov}(b_t, \mu_t) + E[\beta_t]E[\mu_t]$, $\mu_t = b_t \sigma_t^2$, and some manipulation we obtain

$$\text{cov}(b_t, \sigma_t^2) = - \frac{\alpha - E[\mu_t](E[\beta_t] - \beta)}{E[\beta_t]} - \frac{E[\beta_t]E[b_t]E[\sigma_t^2] - E[\beta_t b_t \sigma_t^2]}{E[\beta_t]}. \quad (\text{E.0.27})$$

We now use that $\beta_t = \frac{1}{\sigma_t^2 \zeta_t}$ and measurement error ζ_t is independent of σ_t^2 , b_t , and μ_t to obtain the result. \square

Note that $E[\sigma^2]E[1/\sigma_t^2] - 1 \geq 0$, and will typically be greater than zero if variance is time-varying. For the market, the term $\alpha - E[\mu_t](E[\beta_t] - \beta)$ is roughly zero. This means that the increase in average beta is completely reflected in alpha.

E.1 The covariance between the price of risk and variance is negative

Table 14 present estimates for the co-variance between the price of risk and variance $\text{cov}(b_t, \sigma_t^2)$ using the relationship in Equation (E.0.25). We compute standard errors and confidence intervals using bootstrap. The point estimate for the covariance is negative for all factors besides SMB. Again, this is to a large extent the result of the large positive

alphas of our volatility timing strategy. In most cases the 95% confidence intervals do not include 0, with the exception of SMB and CMA. Therefore, for most factors the data easily rejects a value of the covariance above 0, and imply negative point estimate in almost all cases.

These results show that volatility timing works because expected returns do not increase nearly as much as required to keep the price of risk constant. The negative relation between volatility and the slope of the mean-variance frontier implies an strategy that takes more risk when volatility is low will typically produce higher a Sharpe ratio.

Figure 8: Increase in volatility managed Sharpe ratios by country. 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 country 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 taken 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.


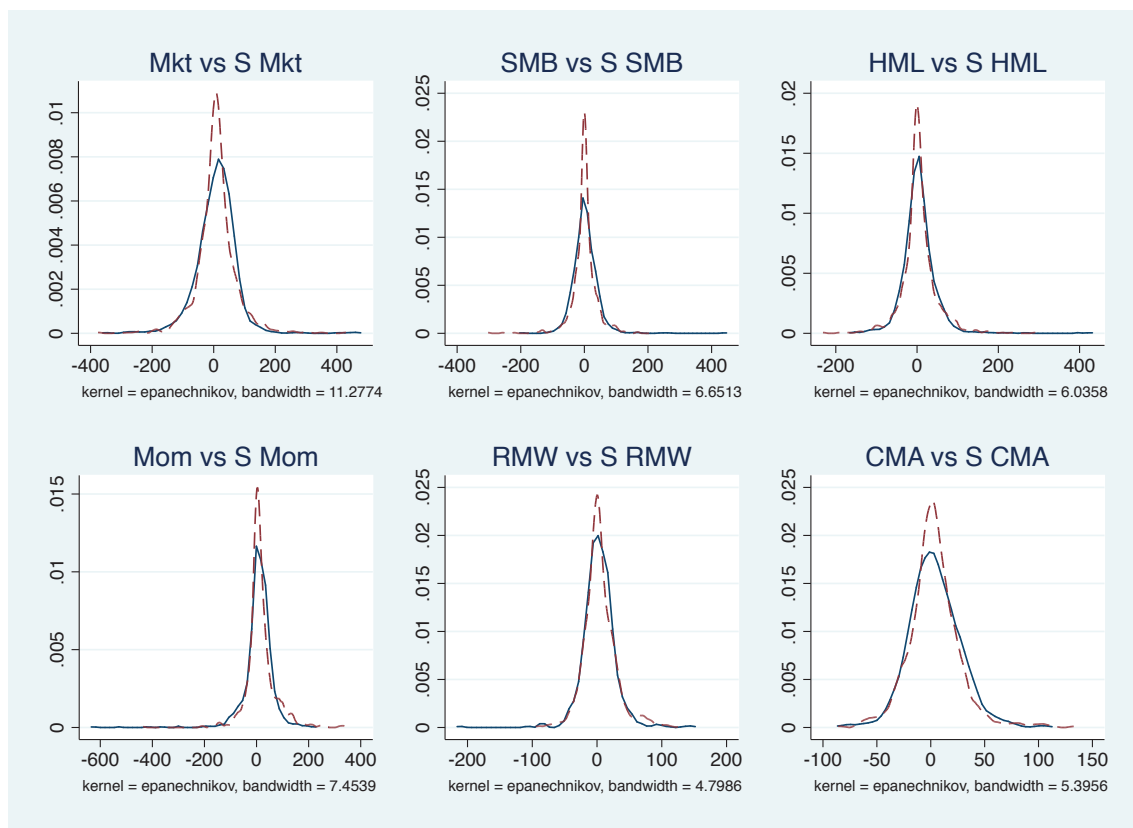


Figure 9: Distribution of volatility managed factors. The figure plots the full distribution of scaled factors (S) vs non-scaled factors estimated using kernel density estimation. The scaled factor, f_t^σ , scales by the factors inverse realized variance in the preceding month $f_t^\sigma = \frac{c}{RV_{t-1}^2} f_t$. In particular, for each panel we plot the distribution of f_t (solid line) along with the distribution of $\frac{c}{RV_{t-1}^2} f_t$ (dashed line).



B. Additional Tables

Table 10: Alphas when using expected rather than realized variance. 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 managed factor, f^σ , scales by the factors inverse realized variance in the preceding month $f_{t+1}^\sigma = \frac{c}{E_{t-1}[RV_t^2]} f_t$. 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. All factors are annualized in percent per year by multiplying monthly factors by 12.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Mkt $^\sigma$	SMB $^\sigma$	HML $^\sigma$	Mom $^\sigma$	RMW $^\sigma$	CMA $^\sigma$	MVE $^\sigma$	FX $^\sigma$
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 11: Time-series alphas controlling for betting against beta factor. 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 (2014). The managed factor, f^{σ} , scales by the factors inverse realized variance in the preceding month $f_t^{\sigma} = \frac{c}{RV_{t-1}^2} f_t$. 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.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt $^{\sigma}$	SMB $^{\sigma}$	HML $^{\sigma}$	Mom $^{\sigma}$	RMW $^{\sigma}$	CMA $^{\sigma}$	MVE $^{\sigma}$
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 12: Alphas of volatility managed factors when controlling for other risk factors. We run time-series regressions of each managed factor on the 4 Fama-French Carhart factors. The managed factor, f^σ , scales by the factors inverse realized variance in the preceding month $f_t^\sigma = \frac{c}{RV_{t-1}^2} f_t$. 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)
	Mkt $^\sigma$	SMB $^\sigma$	HML $^\sigma$	Mom $^\sigma$	MVE $^\sigma$
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 13: Alphas of volatility managed factors when controlling for other risk factors. We run time-series regressions of each managed factor on the 6 Fama-French Carhart factors. The managed factor, f^σ , scales by the factors inverse realized variance in the preceding month $f_t^\sigma = \frac{c}{RV_{t-1}^2} f_t$. 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.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Mkt $^\sigma$	SMB $^\sigma$	HML $^\sigma$	Mom $^\sigma$	RMW $^\sigma$	CMA $^\sigma$	MVE $^\sigma$	MVE2 $^\sigma$
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 14: Covariance between price and quantity of risk. 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 (E.0.25) the Appendix. 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 15: Calibration for portfolio choice exercise. This Table shows targeted moments in the data and in the three alternative models analyzed in Tables 16 and 17. Model based quantities are found by simulating monthly data for 100 different 100 year histories. We report the median of each moment across these 100 histories. RV_t represents realized standard deviation and time is given in months. The seven reported moments capture: average volatility, pe-year-ahead predictability of stock returns, auto-correlation of log variance at the monthly frequency, one-year ahead predictability of returns using realized variance, the volatility of log variance, the correlation between returns and the innovation in variance, and the persistence of the price dividend ratio.

Moment	Data	(a)	(b)	(c)	(d)	(e)
$E[RV_{t+1}]\sqrt{12}$	0.160	0.180	0.235	0.164	0.163	0.163
$corr(pd_t, R_{t \rightarrow t+12})^2$	0.060	0.041	0.020	0.068	0.060	0.059
$corr(\log(RV_{t+1}), \log(RV_t))^2$	0.500	0.397	0.027	0.506	0.445	0.501
$corr(R_{t \rightarrow t+12}, RV_t^2)^2$	0.000	0.008	-0.005	0.001	0.025	-0.030
$stdev(\log(RV_t^2))$	1.050	0.665	0.369	0.922	0.750	0.922
$corr(R_t, RV_t - RV_{t-1})$	-0.240	0.001	0.001	0.003	-0.196	-0.175
$corr(pd_{t+12}, pd_t)$	0.90	0.897	0.892	0.890	0.886	0.889

Table 16: Optimal volatility timing by investor horizon. Panels (a)-(c) show linear approximations to the optimal portfolio policy for calibrations (a)-(c). In calibration (a) all variation in volatility is driven by the volatility of mean-reverting shocks, in calibration (b) it is driven only by the volatility of permanent shocks, and calibration (c) is the case where the share of permanent and mean-reverting shocks is constant. Specifically, we run regressions of the true optimal policy evaluated at the unconditional expected return, $w(\bar{x}, \sigma_t^2)$, on $\frac{\mu_x}{\gamma \sigma_t^2}$ in numerical simulations as we vary volatility $\sigma_t^2 = D'Dy_t + F'F$. A mean variance investor will have a coefficient of $b = 1$ so b measures the relative extent of volatility timing compared to a myopic investor. The last column reports the minimum R-squares across these regressions in a given row, and thus measures how closely the approximation matches the true optimal policy. See text for more details. The calibrations are in Table 9.

(a) Mean-reverting shocks									
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	0.73	0.52	0.64	0.43	0.44	0.34	0.33	0.27	1.00
20	0.73	0.52	0.65	0.46	0.49	0.39	0.36	0.33	1.00
10	0.74	0.54	0.69	0.52	0.58	0.50	0.48	0.47	1.00
5	0.76	0.56	0.74	0.61	0.69	0.63	0.63	0.63	1.00
0	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.00
(b) Permanent shocks									
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	1.86	2.28	2.15	2.36	2.34	2.44	1.00
20	1.67	2.18	1.79	2.18	2.01	2.21	2.19	2.24	1.00
10	1.60	2.12	1.63	1.96	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 share of permanent and mean-reverting shocks									
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	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
20	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
10	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
5	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
0	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.00

Table 17: Optimal volatility timing by investor horizon. Panels (d) and (e) show linear approximations to the optimal portfolio policy in the cases of a constant share of permanent and mean-reverting shocks. Panel (d) introduces a positive correlation between mean-reverting shocks and volatility shocks and Panel (e) introduces a correlation between permanent shocks and volatility shocks. Specifically, we run regressions of the true optimal policy evaluated at the unconditional expected return, $w(\bar{x}, \sigma_t^2)$, on $\frac{\mu_x}{\gamma \sigma_t^2}$ in numerical simulations as we vary volatility $\sigma_t^2 = D'Dy_t + F'F$. A mean variance investor will have a coefficient of $b = 1$ so b measures the relative extent of volatility timing compared to a myopic investor. The last column reports the minimum R-squares across these regressions in a given row, and thus measures how closely the approximation matches the true optimal policy. See text for more details. The calibrations are in Table 9.

(d)Volatility shocks positively correlated to mean-reverting 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	0.84	0.62	0.72	0.57	0.93	0.94	0.98	0.96	0.99
20	0.84	0.63	0.73	0.58	0.93	0.93	0.98	0.96	0.99
10	0.84	0.63	0.75	0.61	0.92	0.92	0.98	0.96	1.00
5	0.85	0.64	0.78	0.71	0.93	0.93	0.99	0.97	1.00
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 permanent 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	0.94	0.95	0.94	0.74	0.95	0.95	0.93	0.92	1.00
20	0.94	0.95	0.94	0.74	0.95	0.96	0.93	0.94	1.00
10	0.94	0.95	0.94	0.76	0.96	0.96	0.94	0.96	1.00
5	0.94	0.95	0.94	0.79	0.95	0.98	0.95	0.98	1.00
0	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.00