VOLATILITY MANAGED PORTFOLIOS

Alan Moreira
Tyler Muir

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

ABSTRACT

Managed portfolios that take less risk when volatility is high produce large alphas, substantially increase factor Sharpe ratios, and produce large utility gains for mean-variance investors. We document this for the market, value, momentum, profitability, return on equity, and investment factors in equities, as well as the currency carry trade. Volatility timing increases Sharpe ratios because changes in factor volatilities are not offset by proportional changes in expected returns. Our strategy is contrary to conventional wisdom because it takes relatively less risk in recessions and crises yet still earns high average returns. This rules out typical risk-based explanations and is a challenge to structural models of time-varying expected returns.

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, μ_t/σ_t^2. 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 and Sharpe ratios with respect to the original factors are substantial. For the market portfolio our strategy produces an alpha of 4.9%, an Appraisal ratio of 0.33, and an overall 25% increase in the buy-and-hold Sharpe ratio.

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 standard 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 representative agent models focused on the dynamics of risk premia. From the vantage point of these theories, the empirical pattern in Figure 1 implies that investor's willingness to take stock market risk must 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 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. A mean-variance investor would decrease his or her risk exposure by around 50% 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 helps to reconcile the evidence on countercyclical expected returns with the profitability of our strategy. Relatedly, we also show that our alphas slowly decline as the rebalancing period grows because current volatility is a weaker forecast for future volatility as we increase horizon.

We go through an extensive list of exercises to evaluate the robustness of our result.

---

[1]For 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`

[2]See 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.

We show that the typical investors can benefit from volatility timing even if subject to realistic transaction costs and hard leverage constraints. The strategy works just as well if implemented through options to achieve high embedded leverage, which further suggests that leverage constraints are unlikely to explain the high alphas of our volatility managed strategies. Consistent with these results, we show that our volatility managed strategy is different from 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).

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, that it is less exposed to volatility shocks than the original factors (ruling out explanations based on the variance risk premium), cannot be explained by downside market risk (Ang et al., 2006a; Lettau et al., 2014), disaster risk or jump risk, and that it outperforms not only using alpha and Sharpe ratios but also manipulation proof measures of performance (Goetzmann et al., 2007).

Once we establish that the profitability of our volatility managed portfolios is a robust feature of the data, we study the economic interpretation of our results in terms of utility gains, the behavior of the aggregate price of risk, and equilibrium models. First, we find that mean-variance utility gains from our volatility managed strategy are large, about 65% of lifetime utility. This compares favorably with Campbell and Thompson (2008), and a longer literature on return predictability, who find mean-variance utility benefits of 35% from timing expected returns.

Next we show more formally how the alpha of our volatility managed portfolio relates to the risk-return tradeoff. In particular, we show that $\alpha \propto -cov(\mu_t/\sigma_t^2, \sigma_t^2)$. Thus, consistent with Figure 1, the negative relationship between μ_t/σ_t^2 and conditional variance drives our positive alphas. The positive alphas we document across all strategies implies that the factor prices of risk, μ_t/σ_t^2, are negatively related to factor variances in each case. When the factors span the conditional mean variance frontier, this result tells us about the aggregate variation in the price of risk, i.e. it tells us about compensation for

risk over time and the dynamics of the stochastic discount factor. Formally, we show how to use our strategy alpha to construct a stochastic discount factor that incorporates these dynamics and that can unconditionally price a broader set of dynamic strategies with a large reduction in pricing errors.

Lastly, we contrast the price of risk dynamics we recover from the data with leading structural asset pricing theories. These models all feature a weakly positive correlation between μ_t/σ_t^2 and variance so that volatility managed alphas are either negative or near zero. 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. More specifically, the alphas of our volatility managed portfolios pose a challenge to macro-finance models that is statistically sharper than standard risk-return regressions which produce mixed and statistically weak results (see Glosten et al. (1993), Whitelaw (1994), Lundblad (2007), Lettau and Ludvigson (2003)).[3] Consistent with this view, we simulate artificial data from these models and show that they are able to produce risk-return tradeoff regressions that are not easily rejected by the data. However, they are very rarely able to produce return histories consistent with the volatility managed portfolio alphas that we document. Thus, the facts documented here are sharper challenges to standard models in terms of the dynamic behavior of volatility and expected returns.

The general equilibrium results and broader economic implications that we highlight also demonstrate why our approach differs from other asset allocation papers which use volatility because our results can speak to the evolution of the aggregate risk return trade-off. 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.[4] 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. Sec-

---

[3]See also related work by Bollerslev et al. (2016) and Tang and Whitelaw (2011).

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

tion 3 studies our strategy in more detail and provides various robustness checks. Section 4 shows formally the economic content of our volatility managed alphas. Section 5 discusses implications for structural asset-pricing models. Section 6 concludes.

2. Main 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. Finally, we 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 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.

2.2 Portfolio formation

We construct our volatility managed portfolios by scaling an excess return by the inverse of its conditional variance. Each month our strategy increases or decreases risk exposure to the portfolio according to variation in our measure of conditional variance. The managed portfolio is then

$$f_{t+1}^{\sigma} = \frac{c}{\hat{\sigma}_t^2(f)} f_{t+1},\tag{1}$$

5

where f_{t+1} is the buy-and-hold portfolio excess return, $\hat{\sigma}_t^2(f)$ is a proxy for the portfolio's conditional variance, and the constant c controls the average exposure of the strategy. For ease of interpretation, we choose c so that the managed portfolio has the same unconditional standard deviation as the buy-and-hold portfolio.[5]

The motivation for this strategy comes from the portfolio problem of a mean-variance investor that is deciding how much to invest in a risky portfolio (e.g. the market portfolio). IThe optimal portfolio weight is proportional to the attractiveness of the risk-return trade-off, i.e. $w_t^* \propto \frac{E_t[f_{t+1}]}{\hat{\sigma}_t^2(f)}$.[6] Motivated by empirical evidence that volatility is highly variable, persistent, and does not predict returns, we approximate the conditional risk-return trade-off by the inverse of the conditional variance. In our main results, we keep the portfolio construction even simpler by using the previous month realized variance as a proxy for the conditional variance,

$$\hat{\sigma}_t^2(f) = RV_t^2(f) = \sum_{d=1/22}^{1} \left(f_{t+d} - \frac{\sum_{d=1/22}^{1} f_{t+d}}{22} \right)^2 \tag{2}$$

An appealing feature of this approach is that it can be easily implemented by an investor in real time and does not rely on any parameter estimation. We plot the realized volatility for each factor in Figure 2. Appendix A.1 considers the use of more sophisticated variance forecasting models.

2.3 Empirical methodology

We run a time-series regression of the volatility managed portfolio on the original factors,

$$f_{t+1}^{\sigma} = \alpha + \beta f_{t+1} + \epsilon_{t+1}. \tag{3}$$

A positive intercept implies that volatility timing increases Sharpe ratios relative to the original factors. When this test is applied to systematic factors (e.g. the market port-

---

[5]Importantly 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.

[6]This is true in the univariate case but also in the multifactor case when factors are approximately uncorrelated.

folio) that summarize pricing information for a wide cross-section of assets and strategies, a positive alpha implies that our volatility managed strategy expands the mean-variance frontier. Our approach is to lean on the extensive empirical asset pricing literature in identifying these factors. That is, a large empirical literature finds that the factors we use summarize the pricing information contained in a wide set of assets and therefore we can focus on understanding the behavior of just these factors.

2.4 Single factor portfolios

We first apply our analysis factor by factor. The single factor alphas have economic interpretation when the individual factors describes well the opportunity set of investors or if these factors have low correlation with each other, i.e. each one captures a different dimension of risk. The single factor results are also useful to show the empirical pattern we document is pervasive across factors and that our result are uniquely driven by the time-series relationship between risk and return.

Table 1 reports the results from running a regression of the volatility managed portfolios on the original factors. We see positive, statistically significant intercepts (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] Finally, in the bottom of the Table, we show that these results are relatively unchanged when we control for the Fama-French three factors in addition to the original factor in every regression. Later sections discuss multifactor adjustments more broadly.

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

---

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

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

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 Great 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.

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 $\frac{\alpha}{\sigma_\varepsilon}$, 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 new Sharpe ratio is $SR_{new} = \sqrt{SR_{old}^2 + \left(\frac{\alpha}{\sigma_\varepsilon}\right)^2}$ 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 notable appraisal ratios across factors are: HML (0.20), Profitability (0.41), Carry (0.44), ROE (0.80), and Investment (0.32).

An alternative way to quantify the economic relevance of our results is from the perspective of a simple mean-variance investor. The percentage utility gain is

$$\Delta U_{MV}(\%) = \frac{SR_{new}^2 - SR_{old}^2}{SR_{old}^2}. \tag{4}$$

Our results imply large utility gains. For example, a mean-variance investors that can only trade the market portfolio can increase lifetime utility by 65% through volatility

---

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

timing. We extend these computations to long-lived investors and more general preferences in Moreira and Muir (2016). The extensive market timing literature provides a useful benchmark for these magnitudes. Campbell and Thompson (2008) estimate that the utility gain of timing expected returns is 35% of lifetime utility. Volatility timing not only generates gains almost twice as large, but also works across multiple factors.

2.5 Multifactor portfolios

We now extend our analysis to a multifactor environment. We first construct a portfolio by combining the multiple factors. We choose weights so that our multifactor portfolio is mean-variance efficient for the set of factors, and as such, the multifactor portfolio prices not only the individual factors but also the wide set of assets and strategies priced by them. We refer to this portfolio as multifactor mean-variance efficient (MVE). It follows that the MVE alpha is the right measure of expansion in the mean-variance frontier. Specifically, a positive MVE alpha implies that our volatility managed strategy increases Sharpe ratios relative to the best buy-and-hold Sharpe ratio achieved by someone with access to the multiple factors.

We construct the MVE portfolio as follows. Let F_{t+1} be a vector of factor returns and b the static weights that produce the maximum in sample Sharpe ratio. We define the MVE portfolio as $f_{t+1}^{MVE} = b'F_{t+1}$. We then construct

$$f_{t+1}^{MVE,\sigma} = \frac{c}{\hat{\sigma}_t^2\left(f_{t+1}^{MVE}\right)} f_{t+1}^{MVE}, \tag{5}$$

where again c is a constant that normalizes the variance of the volatility managed portfolio to be equal to the MVE portfolio. Thus, our volatility managed portfolio only shifts the conditional beta on the MVE portfolio, but *does not* change the relative weights across individual factors. As a result, our strategy focuses uniquely on the time-series aspect of volatility timing.

In Table 2, we show that the volatility timed MVE portfolios have positive alpha with respect to the original MVE portfolios for all combinations of factors we consider including the Fama French three and five factors, or the Hou, Xue, and Zhang factors. This

9

finding is robust to including the momentum factor as well. Appraisal ratios are all economically large and range from 0.33 to 0.91. Note that the original MVE Sharpe ratios are likely to be overstated relative to the truth, since the weights are constructed in sample. Thus, the increase in Sharpe ratios we document are likely to be understated.[10]

We also analyze these MVE 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. Volatility varied far less in the 1956-1986 period, consistent with lower alphas during this time.

3. Understanding the profitability of volatility timing

In this section we investigate our strategy from several different perspectives. Each section is self-contained so a reader can easily skip across sections without loss.

3.1 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 3 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

---

[10]We thank Tuomo Vuolteenaho for this point.

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 volatility 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   Transaction costs

We show that our strategies survive transaction costs. These results are given in Table 4. 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 two strategies that cap the strategy's leverage at 1 and 1.5, respectively. 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 we do not have good measures of transaction costs for implementing

11

the original factors, much less their volatility managed portfolios.

3.3   Leverage constraints

In this section we explore the importance of leverage for our volatility managed strategy. We show that the typical investor can benefit from our strategy even under a tight leverage constraint.

Panel A of Table 5 documents the upper distribution of the weights in our baseline strategy for the volatility managed market portfolio. The median weight is near 1. The 75th, 90th, and 99th percentiles are 1.6, 2.6, and 6.4. Thus our baseline strategy uses modest leverage most of the time but does imply rather substantial leverage in the upper part of the distribution, when realized variance is low.

We explore several alternative implementations of our strategy. The first uses realized volatility instead of realized variance. This makes the weights far less extreme, with the 99th percentile around 3 instead of 6. Second, using expected variance from a simple AR(1), rather than realized variance, also reduces the extremity of the weights. Both of these alternatives keep roughly the same Sharpe ratio as the original strategy. Last, we consider our original strategy, but cap the weights to be below 1 or 1.5. Capturing a hard no-leverage constraint and a leverage of 50%, which is consistent with a standard margin requirement. Sharpe ratios do not change but of course the leverage constrained have lower alphas because risk weights are, on average, lower. Still alphas of all of these strategies are statistically significant.

Because Sharpe ratios are not a good metric to asses utility gains in the presence of leverage constraints, in Figure 4 we compute the utility gains for a mean-variance investor. Specifically, consider a mean-variance investor who follows a buy-and-hold strategy for the market with risk exposure $w = \frac{1}{\gamma}\frac{\mu}{\sigma^2}$ and an investor who times volatility by setting $w_t = \frac{1}{\gamma}\frac{\mu}{\sigma_t^2}$. For any risk aversion, γ, we can compute the weights and evaluate utility gains. Figure 4 shows a gain of around 60% for the market portfolio from volatility timing for an unconstrained investor.[11] With no leverage limit, percentage utility gains

---

[11]Note that 60% is slightly different from 65% that we obtain in the Sharpe ratio based calculation done in Section 2.4. The small difference is due to the fact that here we assuming that the mean-variance investor

are the same regardless of risk aversion because investors can freely adjust their average risk exposure.

Next, we impose a constraint on leverage, so that both the static buy-and-hold weight w and the volatility timing weight w_t must be less than or equal to 1 (no leverage) or 1.5 (standard margin constraint). We then evaluate utility benefits. For investors with high risk aversion this constraint is essentially never binding and their utility gains are unaffected. As we decrease the investors' risk aversion, however, we increase their target risk exposure and are more likely to hit the constraint. Taken to the extreme, an investor who is risk neutral will desire infinite risk exposure, and will hence do zero volatility timing, because w_t will always be above the constraint. To get a sense of magnitudes, Figure 4 shows that an investor whose target risk exposure is 100% in stocks (risk aversion $\gamma \approx 2.2$) and who faces a standard 50% margin constraint, will see a utility benefit of about 45%. An investor who targets a 60/40 portfolio of stocks and T-bills and faces a hard no-leverage constraint will have a utility benefit of about 50%. Therefore, the results suggest fairly large benefits to volatility timing even with tight leverage constraints.

For investors whose risk-aversion is low enough, our baseline strategy requires some way to achieve a large risk exposure when volatility is very low. To address the issue that very high leverage might be costly or unfeasible, we implement our strategy using options in the S&P 500 which provide embedded leverage. Of course, there may be many other ways to achieve a β above 1, options simply provide one example. 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 use the option portfolios to achieve our desired risk exposure. In Panel B of Table 5, 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, even for investors with relatively low risk aversion.[12]

only invests in the volatility managed portfolio, while in Section 2.4 the investors is investing in the optimal ex-post mean-variance efficient portfolio combination.

[12]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

Black (1972), Jensen et al. (1972) and Frazzini and Pedersen (2014) show that leverage constraints can distort the risk-return trade-off in the cross-section. The idea is that the embedded leverage of high beta assets make them attractive to investors that are leverage constrained. 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 investors willingness to take risk. While in theory leverage constraints could explain our findings, we find that most investors can benefit of volatility timing under very tight leverage constraints. Therefore constraints does not seem a likely explanation for our findings.

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

3.4   Contrasting with cross-sectional low-risk anomalies

In this section we show empirically that our strategy is also very different from strategies that explore a weak risk return trade-off in the *cross-section* of stocks, which are often attributed to leverage constraints.

The first strategy, popular among practitioners, is risk parity which 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. We follow Asness et al. (2012) and construct risk parity factors as $RP_{t+1} = b_t' f_{t+1}$ where $b_{i,t} = \frac{1/\tilde{\sigma}_t^i}{\sum_i 1/\tilde{\sigma}_t^i}$, and $\tilde{\sigma}_t^i$ is a rolling three year estimate of volatility for each factor (again exactly as in Asness et al. (2012)). This implies that, if the volatility of one factor increases 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, in Table 6 we include a risk parity factor as an additional 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.

options for relatively high levels of leverage.

(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.

3.5 Volatility co-movement

A natural question is whether one can implement our results using a common volatility factor. Because realized volatility is very correlated across factors, normalizing by a common volatility factor does not drastically change our results. To see this, we compute the first principal component of realized variance across all factors and normalize each factor by $\frac{1}{RV_t^{PC}}$.[13] This is in contrast to normalizing by each factors' own realized variance. Table 7 gives the results which are slightly weaker than the main results. For most factors the common volatility timing works about the same. However, it is worth noting that the alpha for the currency carry trade disappears. The realized volatility of the carry trade returns is quite different from the other factors (likely because it represents an entirely different asset class), hence it is not surprising that timing this factor with a common

---

[13]Using an equal weighted average of realized volatilities, or even just the realized volatility of the market return, produce similar results.

15

volatility factor from (mostly) equity portfolios will work poorly.

The strong co-movement among equities validates our approach in Section 2.5, where we impose a constant weight across portfolios to construct the MVE portfolio.

3.6   Horizon effects

We have implemented our strategy by rebalancing it once a month and running time-series regressions at the monthly frequency. A natural question to ask is if our results hold up at lower frequencies. Less frequent rebalancing periods might be interesting from the perspective of macro-finance models that are often targeted at explaining variation in risk premia and the price of risk at quarterly/yearly frequencies. They are also useful to better understand the dynamic relationship between volatility shocks, expected returns, and the price of risk. In particular, it allows us to reconcile our results with the well known empirical facts that movements in both stock-market variance and expected returns are counter-cyclical (French et al., 1987; Lustig and Verdelhan, 2012).

We start by studying the dynamics of risk and return through a vector autoregression (VAR) because it is a convenient tool to document how volatility and expected returns dynamically respond to a volatility shock over time. We run a VAR at the monthly frequency with one lag of (log) realized variance, realized returns, and the price to earning ratio (CAPE from Robert Shiller's website) and plot the Impulse Response Function to trace out the effects of a variance shock. We choose the ordering of the variables so that the variance shock can contemporaneously affect realized returns and CAPE.

Figure 5 plots the response to a one-standard deviation expected variance shock. While expected variance spikes on impact, this shock dies out fairly quickly, consistent with variance being strongly mean reverting. Expected returns, however, rise much less on impact but stay elevated for a longer period of time. Given the increase in variance but only small and persistent increase in expected return, the lower panel shows that it is optimal for the investor to reduce his portfolio exposure by 50% on impact because of an unfavorable risk return tradeoff. The portfolio share is consistently below 1 for roughly 12 months after the shock.

The lower persistence of volatility shocks implies the risk-return trade-off initially

deteriorates but gradually improves as volatility recedes through a recession. Thus, our results are not in conflict with the evidence that risk premia is counter-cyclical. Instead, after a large market crash such as October 2008, our strategy gets out of the market initially to avoid an unfavorable risk return tradeoff, but captures much of the persistent increase in expected returns by buying back in when the volatility shock subsides.

However, the estimated response of expected returns to a volatility shock should be read with caution, as return predictability regressions are poorly estimated. With this in mind we also study the behavior of our strategy at lower frequencies. Specifically, we form portfolios as before, using weights proportional to monthly realized variance, but now we hold the position for T months before rebalancing. We then run our time-series alpha test at the same frequency. Letting $f_{t \to t+T}$ be the cumulative factor excess returns from buying at the end of month T and holding till the end on month $t + T$, we run,

$$\frac{c}{\hat{\sigma}_t^2(f_{t+1})} f_{t \to t+T} = \alpha + \beta f_{t \to t+T} + \epsilon_{t+T}, \tag{6}$$

with non-overlapping data. Results are in Figure 6. We show alphas and appraisal ratios for the market and the MVE portfolios based on the Fama-French three factors and momentum factor. Alphas are statistically significant for longer holding periods but gradually decay in magnitude. For example, for the market portfolio, alphas are statistically different from zero (at the 10% confidence level) for up to 18 months. This same pattern holds up for the two MVE portfolios we consider.

These results are broadly consistent with the VAR in that alphas decrease with horizon. However, empirically volatility seems to be more persistent at moderate or long horizons than implied by it's very short-term dynamics. For example, the estimated VAR dynamics implies volatility has a near zero 12 month auto-correlation, while the non-parametric estimate is larger than 0.2. This means the alphas decline more slowly than the VAR suggests.

The economic content of the long-horizon alphas is similar to the monthly results. These results imply that even at lower frequencies there is a negative relation between variance and the price of risk (see Section 4).

3.7 Additional analysis

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). We also find our volatility managed factors are 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.

4. Theoretical framework

In this section we provide a theoretical framework to interpret our findings. We start by making the intuitive point that our alphas are proportional to the co-variance between variance and the factor price of risk. We then impose more structure to derive aggregate pricing implications.

We get sharper results in continuous time. Consider a portfolio excess return dR_t with expected excess return μ_t and conditional volatility σ_t^2. Construct the volatility managed version of this return exactly as in Equation (1), i.e. $dR_t^\sigma = \frac{c}{\sigma_t^2} dR_t$, where c is a normalization constant. The α of a time-series regression of the volatility managed portfolio dR_t^σ on the original portfolio dR_t is given by

$$\alpha = E[dR_t^\sigma]/dt - \beta E[dR_t]/dt. \tag{7}$$

Using that $E[dR_t^\sigma]/dt = cE\left[\frac{\mu_t}{\sigma_t^2}\right]$, $\beta = \frac{c}{E[\sigma_t^2]}$, and $cov\left(\frac{\mu_t}{\sigma_t^2}, \sigma_t^2\right) = E[\mu_t] - E\left[\frac{\mu_t}{\sigma_t^2}\right]E[\sigma_t^2]$, we obtain a relation between alpha and the dynamics of the price of risk μ_t/σ_t^2,

$$\alpha = -cov\left(\frac{\mu_t}{\sigma_t^2}, \sigma_t^2\right)\frac{c}{E[\sigma_t^2]} \tag{8}$$

Thus, our α is a direct measure of the comovement between the price of risk and variance. In the case where expected returns and volatility move together, i.e. $\mu_t = \gamma \sigma_t^2$, then trivially $\alpha = 0$. Intuitively, by avoiding high volatility times you avoid risk, but if the risk-return tradeoff is strong you also sacrifice expected returns, leaving the volatility timing strategy with zero alpha.

In contrast, when expected returns are constant or independent of volatility, Equation (8) implies $\alpha = c \frac{E[\mu_t]}{E[\sigma_t^2]} J_\sigma$, where $J_\sigma = \left(E[\sigma_t^2] E\left[\frac{1}{\sigma_t^2} \right] - 1 \right) > 0$ is a Jensen's inequality term that is increasing in the volatility of volatility. This is because the more volatility varies, the more variation there is in the price of risk that the portfolio can capture. Thus, the alpha of our strategy is increasing in the volatility of volatility and the unconditional compensation for risk.

The profitability of our strategy can also be recast in term of the analysis in Jagannathan and Wang (1996) because we are testing a strategy with zero conditional alpha using an unconditional model.[14] The above results provide an explicit mapping between volatility managed alphas and the dynamics of the price of risk for an individual asset.

4.1 The aggregate price of risk

While the above methodology applies to any return – even an individual stock – the results are only informative about the broader price of risk in the economy if applied to systematic sources of return variation. Intuitively, if a return is largely driven by idiosyncratic risk, then volatility timing will not be informative about the broader price of risk in the economy.[15] In this section we show how our volatility managed portfolios, when applied to systematic risk factors, recover the component of the aggregate price of risk variation driven by volatility.

Let $dR = [dR_1, ..., dR_N]'$ be a vector of returns, with expected excess return μ_t^R and covariance matrix Σ_t^R. The empirical asset pricing literature shows that exposures to a few factors summarize expected return variation for a larger cross-section of assets and

---

[14]See also Appendix A.6.1 where we show how to explicitly recover from our strategy alpha the strength of the conditional relationship between risk and return.

[15]See example in Appendix A.6.2.

strategies captured by dR_t. We formalize our interpretation of this literature as follows:

Assumption 1. *Let return factors $dF = [dF_1, ..., dF_I]$, with dynamics given by μ_t and Σ_t, span the unconditional mean-variance frontier for static portfolios of $d\widetilde{R} = [dR; dF_t]$, and the conditional mean-variance frontier for dynamic portfolios of $d\widetilde{R}$. Define the process $\Pi_t(\gamma_t)$ as*

$$\frac{d\Pi_t(\gamma_t)}{\Pi_t(\gamma_t)} = -r_t dt - \gamma_t'(dF_t - E_t[dF_t]), \tag{9}$$

then there exists a constant price of risk vector γ^u such that $E[d(\Pi_t(\gamma^u)\widetilde{R})] = 0$ holds for any static weights w, and there is a γ_t^* process for which $E[d(\Pi_t(\gamma_t^*)w_t\widetilde{R})] = 0$ holds for any dynamic weights w_t.

This assumption says that unconditional exposures to these factors contain all relevant information to price the static portfolios R, but one also needs information on the price of risk dynamics to properly price dynamic strategies of these assets.

We focus on the case where the factor covariance matrix is diagonal, $\Sigma_t = diag([\sigma_{1,t}...\sigma_{I,t}])$, i.e. factors are uncorrelated, which is empirically a good approximation for the factors we study.[16] In fact, many of the factors are constructed to be nearly orthogonal through double sorting procedures. Given this structure, we can show how our strategy alphas allows one to recover the component of the price of risk variation driven by volatility.

Implication 1. *The factor i price of risk is $\gamma_{i,t}^* = \frac{\mu_{i,t}}{\sigma_{i,t}^2}$ and $\gamma_i^u = \frac{E[\mu_{i,t}]}{E[\sigma_{i,t}^2]}$. Decompose factor excess returns as $\mu_t = b\Sigma_t + \zeta_t$, where we assume $E[\zeta_t|\Sigma_t] = \zeta_t$. Let $\gamma_{i,t}^\sigma = E[\gamma_{i,t}^*|\sigma_{i,t}^2]$ be the component of price of risk variation driven by volatility, and α_i be factor i volatility managed alpha, then*

$$\gamma_{i,t}^\sigma = \gamma_i^u + \frac{\alpha_i}{c_i}J_{\sigma,i}^{-1}\left(\frac{E[\sigma_{i,t}^2]}{\sigma_{i,t}^2} - 1\right), \tag{10}$$

and the process $\Pi_t(\gamma_t^\sigma)$ is a valid SDF for $d\widetilde{R}_t$ and volatility managed strategies $w(\Sigma_t)$, i.e. $E\left[d\left(\Pi_t(\gamma_t^\sigma)w(\Sigma_t)\widetilde{R}_t\right)\right] = 0.$[17]

---

[16] Appendix A.6.5 deals with the case where factor are correlated.

[17] Formally, $\gamma_t^\sigma = [\gamma_{1,t}^\sigma...\gamma_{I,t}^\sigma]$, and the strategies $w(\Sigma_t)$ must be adapted to the filtration generated by Σ_t, self-financing, and satisfy $E[\int_0^T ||w(\Sigma_t)\Sigma_t||^2 dt] < \infty$ (see Duffie (2010))

Equation (10) shows how volatility managed portfolio alphas allow us to reconstruct the variation in the price of risk due to volatility. The volatility implied price of risk has two terms. The term γ^u is the unconditional price of risk, the price of risk that prices static portfolios of returns dR_t. It is the term typically recovered in cross-sectional tests. The second is due to volatility. It increases the price of risk when volatility is low with this sensitivity increasing in our strategy alpha. Thus, volatility managed alphas allow us to answer the fundamental question of how much compensation for risk moves as volatility moves around.

Tracking variation in the price of risk due to volatility can be important for pricing. Specifically, $\Pi(\gamma_t^g)$ can price not only the original assets unconditionally, but also volatility based strategies of these assets.[18] Thus, volatility managed portfolios allow us to get closer to the true price of risk process γ_t^*, and as a result, closer to the unconditional mean-variance frontier, a first-order economic object. In Appendix A.6.4 we show how one can implement the risk-adjustment embedded in model $\Pi(\gamma_t^g)$ by adding our volatility managed portfolios as a factor.

We finish this section by providing a measure of how "close" $\Pi(\gamma_t^g)$ gets to $\Pi(\gamma_t^*)$ relative to the constant price of risk model $\Pi(\gamma^u)$. Recognizing that $E[(d\Pi(\gamma_t^a) - d\Pi(\gamma_t^b)) dR_t]$ is the pricing error associated with using model b when prices are consistent with a, it follows that the volatility of the difference between models, $\mathcal{D}_{b-a} \equiv Var(d\Pi(\gamma_t^a) - d\Pi(\gamma_t^b))$, provides an upper bound on pricing error Sharpe ratios (see Hansen and Jagannathan (1991)). It is thus a natural measure of distance. For the single factor case, we obtain

$$\mathcal{D}_{u-\sigma} = \left(\frac{\alpha}{c}\right)^2 E[\sigma_t^2] J_\sigma^{-1} \tag{11}$$

$$\mathcal{D}_{u-\zeta} = \frac{Var(\zeta_t)}{E[\sigma_t^2]} \tag{12}$$

$$\mathcal{D}_{u-*} = \left(\frac{\alpha}{c}\right)^2 E[\sigma_t^2] J_\sigma^{-1} + \frac{Var(\zeta_t)}{E[\sigma_t^2]}(J_\sigma + 1). \tag{13}$$

---

[18]For example, Boguth et al. (2011) argues that a large set of mutual fund strategies involve substantial volatility timing. Our volatility managed portfolio provides a straightforward method to risk-adjust these strategies. This assumes of course that investors indeed understand the large gain from volatility timing and nevertheless find optimal not to trade.

Equation (11) shows that the distance between models u and σ grows with alpha. It implies that the maximum excess Sharpe ratios decrease proportionally with the strategy alpha when you move from the constant price of risk model u to the model σ that incorporates variation in the price of risk driven by volatility. This is similar in spirit to Nagel and Singleton (2011) who derive general optimal managed portfolios based on conditioning information to test unconditional models against. Analogously, Equation (12) accounts for variation in the expected return signal ζ_t, but ignores volatility information. Equation (13) shows the total difference between the constant price of risk model u and the true $*$ model.

To have a sense of magnitudes, we assume that the market portfolio satisfies Assumption 1 and plug in numbers for the market portfolio. Notice that $\mathcal{D}_{u-\sigma}$ is the volatility managed market's appraisal ratio squared which measures the expansion of the MVE frontier for the managed strategy. We measure all the quantities in (11)-(13) but $Var(\zeta_t)$, which is tightly related to return predictability R-square. We use the estimate from Campbell and Thompson (2008) who obtain a number around 0.06.[19] We obtain $\mathcal{D}_{u-\sigma} = 0.33^2 = 0.11$, $\mathcal{D}_{u-\zeta} = 0.06$, and $\mathcal{D}_{u-*} = 0.11 + 0.06 * 3.2 = 0.29$. Accounting for only time-variation in volatility can reduce squared pricing error Sharpe ratios by approximately 0.11/0.29=38%, compared with 0.06/0.29=21% for time-variation in expected returns, with the large residual being due to the multiplicative interaction between them.

This shows that accounting for time-variation in the price of risk driven by volatility seems at least as important, perhaps even more important, than accounting for variation in the price of risk driven by expected returns.

5. General equilibrium implications

We start this section by showing that the high Sharpe ratios of our volatility managed portfolios pose a new challenge to leading macro-finance models. We then discuss potential economic mechanisms that could generate our findings.

---

[19] A range in this literature would put an upper bound around 13% for the R-square at the yearly horizon, see Kelly and Pruitt (2013). Notice also that $Var(\zeta_t)$ is actually below $Var(\mu_t)$ so these are strong upper bounds.

5.1 Macro-finance models

Our empirical findings pose a challenge to macro-finance models that is statistically sharper than standard risk-return regressions. In fact, many equilibrium asset pricing models have largely ignored the risk-return tradeoff literature, which runs regressions of future returns on volatility, because the results of that literature are ambiguous and statistically weak (see Glosten et al. (1993), Whitelaw (1994), Lundblad (2007), Lettau and Ludvigson (2003)).[20]

We show the statistical power of our approach by studying the predictions of four leading equilibrium asset pricing models; the habits model (Campbell and Cochrane, 1999), long run risk model (Bansal et al., 2009), time-varying rare disasters model (Wachter, 2013), and intermediary-based asset pricing model (He and Krishnamurthy, 2012). Specifically, we calibrate each model according to the original papers and simulate stock market return data for a sample of equal length to our historical sample.

We first run the following standard risk-return tradeoff regression in simulated data from each model

$$R_{t+1} = a + \gamma \hat{\sigma}_t^2 + \epsilon_{t+1} \tag{14}$$

We plot the histogram of the coefficient γ across simulations in each model and compare this to the actual point estimate from this regression in the data for the market portfolio. Results are shown in Figure 7.

We then construct our volatility managed portfolios, exactly as described in Section 2.2. We compute alphas and appraisal ratios in the model simulated data and again compare to the actual data for the market portfolio.

The contrast between our approach and the return forecasting approach is striking. Because return predictability regressions are poorly estimated, all models frequently generate return histories consistent with the weak risk-return trade-off estimated in the data. However, no model comes close to reproducing our findings in terms of alphas or appraisal ratios. For example, Bansal and Yaron (2004) generate alphas as high as in the data only in 0.2% of the simulated samples. The other three models do even worse in matching

---

[20]See also related work by Bollerslev et al. (2016) and Tang and Whitelaw (2011)

our estimates. This highlights that our volatility managed portfolios pose a statistically sharper challenge to these models than the standard risk-return tradeoff literature.

Notably, in the models alphas are either near zero or negative on average. This is equivalent to the statement that $cov(\gamma_t^2, \sigma_t^2) \geq 0$ in each of these models where $\gamma_t = E_t[R_{t+1}]/\sigma_t^2$ can be thought of as effective risk aversion. The models generally feature a weakly positive covariance between effective risk aversion and variance because they typically have risk aversion either increasing or staying constant in bad economic times when volatility is also high. The positive alphas we document empirically suggest this covariance would need to be strongly negative.

5.2 What 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.

The easiest, but least plausible, explanation is that investors willingness to take risk is negatively related to volatility. That is, investors choose not to volatility time because they are more risk-averse during low volatility periods. A more nuanced explanation is that non-traded wealth becomes less volatile when financial market volatility is high. We find this explanation also unappealing, as volatility tends to be high in recessions, when macro-economic uncertainty is high. A more plausible explanation is that volatility driven by learning about structural parameters might be priced differently than when driven by standard forms of risk (e.g Veronesi, 2000).

One intuitive explanation is that investors are slow to trade or to update their beliefs about volatility. This could explain why a sharp increase in realized volatility doesn't immediately illicit a response to sell. This explanation would also be consistent with our impulse responses where expected returns rise slowly but the true expected volatility process rises and mean-reverts quickly in response to a variance shock. In line with this view, Nagel et al. (2016) find that lower income households, who may be less sophisticated investors, respond to volatility more slowly to volatility through the 2008 crisis.

A final possibility is that the composition of shocks changes with volatility. In a companion paper (Moreira and Muir, 2016) we show that long-horizon investors can find

volatility timing somewhat less beneficial if increases in volatility are driven by discount rate volatility. That is, the increase in volatility is due to a increase in the volatility of shocks that eventually mean-revert. Intuitively, long-horizon investors are less scared of discount rate volatility because they can wait until the shocks eventually mean-revert. The open challenge is to develop a plausible equilibrium mechanism where discount-rate volatility is not tightly related to the level of discount-rates.

We acknowledge that these explanations need to be considered in much more detail and be analyzed quantitatively before we can evaluate their success, and we leave this task to future work.

6. Conclusion

Volatility managed portfolios offer large risk-adjusted returns and are easy to implement in real time. Because volatility doesn't strongly forecast future returns, factor Sharpe ratios are improved by lowering risk exposure when volatility is high and increasing risk exposure when volatility is low. Our strategy is contrary to conventional wisdom because it takes relatively less risk in recessions and crises yet still earns high average returns. We analyze both portfolio choice and general equilibrium implications of our findings. We find utility gains from volatility timing for mean-variance investors of around 65%, much larger than utility gains that focus on timing expected returns. 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

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., Kiku, D., and Yaron, A. (2009). An empirical evaluation of the long-run risks model for asset prices. Technical report, National Bureau of Economic Research.

Bansal, R. and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509.

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.

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.

Buffett, W. E. (2008). Buy American. I am. *The New York Times*, October 16th, 2008.

Campbell, J. Y. and Cochrane, J. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251.

Campbell, J. Y. and 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.

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

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. (2010). *Dynamic asset pricing theory*. Princeton University Press.

Dybvig, P. H. and Ingersoll Jr, J. E. (1982). Mean-variance theory in complete markets. *Journal of Business*, pages 233–251.

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.

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.

Hansen, L. P. and Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):pp. 225–262.

He, Z. and Krishnamurthy, A. (2012). Intermediary asset pricing. *The American Economic Review*, forthcoming.

Hou, K., Xue, C., and Zhang, L. (2014). Digesting anomalies: An investment approach. *Review of Financial Studies*, forthcoming.

Jagannathan, R. and Wang, Z. (1996). The conditional capm and the cross-section of expected returns. *Journal of Finance*, 51(1):3–53.

Jensen, M. C., Black, F., and Scholes, M. S. (1972). The capital asset pricing model: Some empirical tests.

Kelly, B. and Pruitt, S. (2013). Market expectations in the cross-section of present values. *The Journal of Finance*, 68(5):1721–1756.

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.

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*.

Moreira, A. and Muir, T. (2016). Portfolio choice, volatility, and investment horizon. *working paper*.

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 sold during the crash of 2008-09? evidence from tax return data on daily sales of stock. *working paper*.

Nagel, S. and Singleton, K. J. (2011). Estimation and evaluation of conditional asset pricing models. *The Journal of Finance*, 66(3):873–909.

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.

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^σ, scales by the factors inverse realized variance in the preceding month $f_t^\sigma = \frac{c}{RV_{t-1}^2} f_t$. In Panel B, we include the Fama-French three factors as additional controls in this regression. The data is monthly and the sample is 1926-2015 for Mkt, SMB, HML, and Mom, 1963-2015 for RMW and CMA, 1967-2015 for ROE and IA, and 1983-2015 for the FX Carry factor. Standard errors are in parentheses and adjust for heteroscedasticity. All factors are annualized in percent per year by multiplying monthly factors by 12.

	Panel A: Univariate regressions								
	(1) Mkt^σ	(2) SMB^σ	(3) HML^σ	(4) Mom^σ	(5) RMW^σ	(6) CMA^σ	(7) FX^σ	(8) ROE^σ	(9) IA^σ
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
	Panel B: Alphas also controlling for Fama-French 3 factors								
Alpha (α)	5.45 (1.56)	-0.33 (0.89)	2.66 (1.02)	10.52 (1.60)	3.18 (0.83)	-0.01 (0.68)	2.54 (1.65)	5.76 (0.97)	1.14 (0.69)

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 scales the portfolio by the inverse of the portfolios' realized variance in the previous month. 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) Mkt	(2) FF3	(3) FF3 Mom	(4) FF5	(5) FF5 Mom	(6) HXZ	(7) HXZ Mom
Alpha (α)	4.86	4.99	4.04	1.34	2.01	2.32	2.51
	(1.56)	(1.00)	(0.57)	(0.32)	(0.39)	(0.38)	(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) Mkt	(2) FF3	(3) FF3 Mom	(4) FF5	(5) FF5 Mom	(6) HXZ	(7) HXZ Mom
α: 1926-1955	8.11	1.94	2.45				
	(3.09)	(0.92)	(0.74)				
α: 1956-1985	2.06	0.99	2.54	0.13	0.71	0.77	1.00
	(2.82)	(1.43)	(1.16)	(0.43)	(0.67)	(0.39)	(0.51)
α: 1986-2015	4.22	5.66	4.98	1.88	2.65	3.03	3.24
	(1.66)	(1.74)	(0.95)	(0.41)	(0.47)	(0.50)	(0.57)

Table 3: 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) Mkt^σ	(2) HML^σ	(3) Mom^σ	(4) RMW^σ	(5) CMA^σ	(6) FX^σ	(7) ROE^σ	(8) IA^σ
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	32 0.38	0.49	0.51	0.43	0.49

Table 4: 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 variance, using expected rather than realized variance, and using our original inverse realized variance but limiting risk exposure to be below 1 (i.e., no leverage) or 1.5. 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 | $\lvert \Delta w \rvert$ | $E[R]$ | 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
$min\left(\frac{c}{RV_t^2}, 1\right)$	No Leverage	0.16	5.61%	2.12%	2.10%	1.93%	1.85%	110bps
$min\left(\frac{c}{RV_t^2}, 1.5\right)$	50% Leverage	0.16	7.18%	3.10%	3.08%	2.91%	2.83%	161bps

Table 5: Volatility Timing and Leverage. Panel A shows several alternative volatility managed strategies and the corresponding alphas, Sharpe ratios, and distribution of weights used in each strategy. The alternative strategies include using inverse volatility instead of variance, using expected rather than realized variance, and using inverse realized variance but limiting risk exposure to be below 1 (i.e., no leverage) or 1.5. For expected variance, we run an AR(1) for log variance to form our forecast. In particular, we focus on upper percentiles of weights to determine how much leverage is typically used in each strategy. In each case we focus only on the market portfolio. In Panel B, 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 using either just call options or using both calls and puts. 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 7. Any time our strategy prescribes leverage to achieve high beta, we invest in this option portfolio to achieve our desired beta. We then compare the performance of the embedded leverage volatility timed portfolio to the standard volatility managed portfolio studied in the main text. Finally, we consider an option strategy that also sells in the money puts (with same moneyness as before) as well as buys calls to again achieve our desired beta. The sample used for Panel B is April, 1986 to January 2012 based on data from Constantinides et al. (2013).

<div align="center">Volatility Timing and Leverage</div>

Panel A: Weights and Performance for Alternative Volatility Managed Portfolios								
					Distribution of Weights w			
w_t	Description	α	Sharpe	Appraisal	P50	P75	P90	P99
$\frac{1}{RV_t^2}$	Realized Variance	4.86 (1.56)	0.52	0.34	0.93	1.59	2.64	6.39
$\frac{1}{RV_t}$	Realized Volatility	3.30 (1.02)	0.53	0.33	1.23	1.61	2.08	3.36
$\frac{1}{E_t[RV_{t+1}^2]}$	Expected Variance	3.85 (1.36)	0.51	0.30	1.11	1.71	2.38	4.58
$min\left(\frac{c}{RV_t^2},1\right)$	No Leverage	2.12 (0.71)	0.52	0.30	0.93	1	1	1
$min\left(\frac{c}{RV_t^2},1.5\right)$	50% Leverage	3.10 (0.98)	0.53	0.33	0.93	1.5	1.5	1.5

Panel B: Embedded Leverage Using Options: 1986-2012				
			Vol Timing With Embedded Leverage	
	Buy and hold	Vol Timing	Calls	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 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 $b_{i,t} = \frac{1/\tilde{\sigma}^i_t}{\sum_i 1/\tilde{\sigma}^i_t}$ 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) Mkt	(2) FF3	(3) FF3 Mom	(4) FF5	(5) FF5 Mom	(6) HXZ	(7) HXZ Mom	(8) BAB^σ
Alpha (α)	4.86	5.00	4.09	1.32	1.97	2.03	2.38	5.67
	(1.56)	(1.00)	(0.57)	(0.31)	(0.40)	(0.32)	(0.44)	(0.98)
N	1,065	1,065	1,060	621	621	575	575	996
R^2	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: Normalizing by Common Volatility. We construct managed volatility strategies for each factor using the first principal component of realized variance across all factors. Each factor is thus normalized by the same variable, in contrast to our main results where each factor is normalized by that factors' past realized variance. We run time-series regressions of each managed factor on the non-managed factor.

	(1) Mkt^σ	(2) SMB^σ	(3) HML^σ	(4) Mom^σ	(5) RMW^σ	(6) CMA^σ	(7) FX^σ	(8) ROE^σ	(9) IA^σ
Alpha (α)	4.22	0.24	3.09	11.00	1.16	-0.22	-1.28	4.21	1.24
	(1.49)	(0.83)	(0.96)	(1.70)	(0.81)	(0.66)	(1.21)	(1.00)	(0.61)
N	1,061	1,061	1,061	1,060	622	622	362	576	576
R^2	0.42	0.45	0.36	0.33	0.44	0.51	0.64	0.47	0.56
rmse	49.31	28.74	33.87	46.57	19.11	16.67	18.49	22.13	15.06

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 lagged realized volatility.
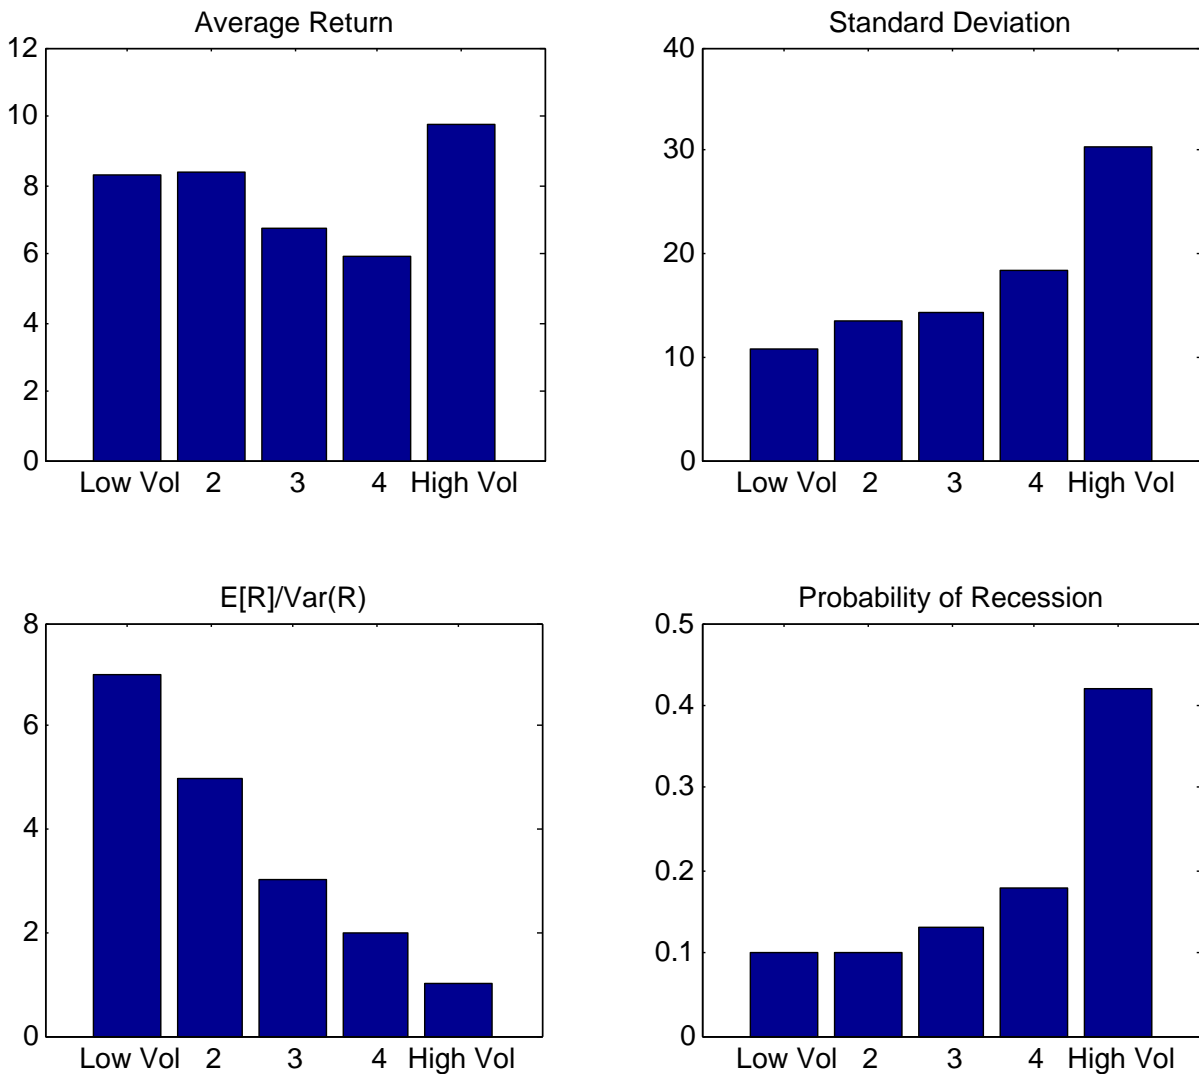
Figure 2: Time-series of volatility by factor. This figures plots the time-series of the monthly volatility of each individual factor. We emphasize the common co-movement in volatility across factors and that volatility generally increases for all factors in recessions. Light shaded bars indicate NBER 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.
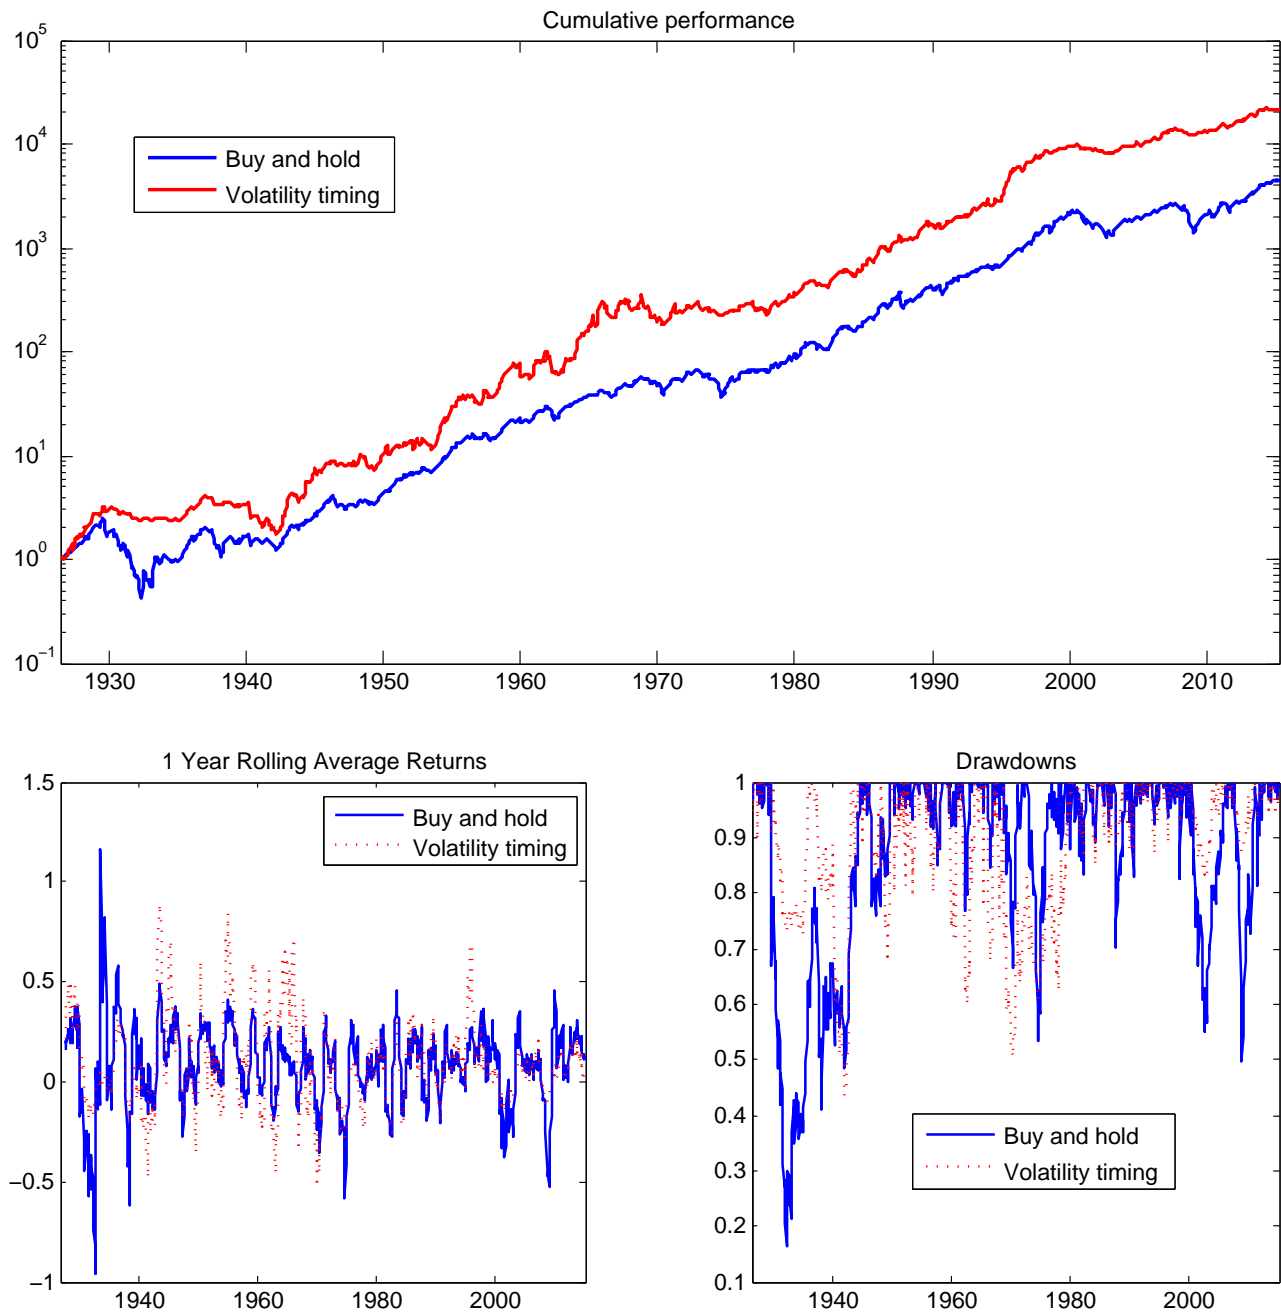
Figure 4: Utility Benefits and Leverage Constraints. We plot the empirical percentage utility gain $\Delta U\%$ for a mean-variance investor going from a buy-and-hold portfolio to a volatility managed portfolio. Specifically $U = E[w_t R_{t+1}] - \frac{1}{2}\gamma var(w_t R_{t+1})$. We compute unconditional buy-and-hold weights as $w = \frac{1}{\gamma}\frac{\mu}{\sigma^2}$ and volatility managed weights as $w_t = \frac{1}{\gamma}\frac{\mu}{\sigma_t^2}$. The x-axis denotes the targeted unconditional weight w as we vary investors risk aversion γ and represents the desired unconditional exposure to equities. The black line shows the percentage increase in utility $(U(w_t)/U(w) - 1)$ when our weights, w, are unrestricted and shows that in this case the utility gain doesn't depend on risk aversion. The red and blue lines impose leverage constraints of zero leverage and 50% leverage (consistent with a standard margin constraint), respectively. We evaluate the utility percentage increases $U(min(w_t, \bar{w}))/U(min(w, \bar{w})) - 1$ with $\bar{w} = \{1, 1.5\}$. Numbers presented are for the market return.

Figure 5: Dynamics of the risk return tradeoff. The figure plots the impulse response of the expected variance and expected return of the market portfolio for a shock to the realized variance. The x-axis is in years. The last panel gives the portfolio choice implications for a mean-variance investor who sets their risk exposure proportional to $E_t[R_{t+1}]/var_t[R_{t+1}]$. The units are percentage deviations from their average risk exposure. We compute impulse responses using a VAR of realized variance, realized returns, and the cyclically-adjusted price-to-earnings ratio (CAPE) from Robert Shiller. We include 2 lags of each variable. Bootstrapped 95% confidence bands are shown in dashed lines.
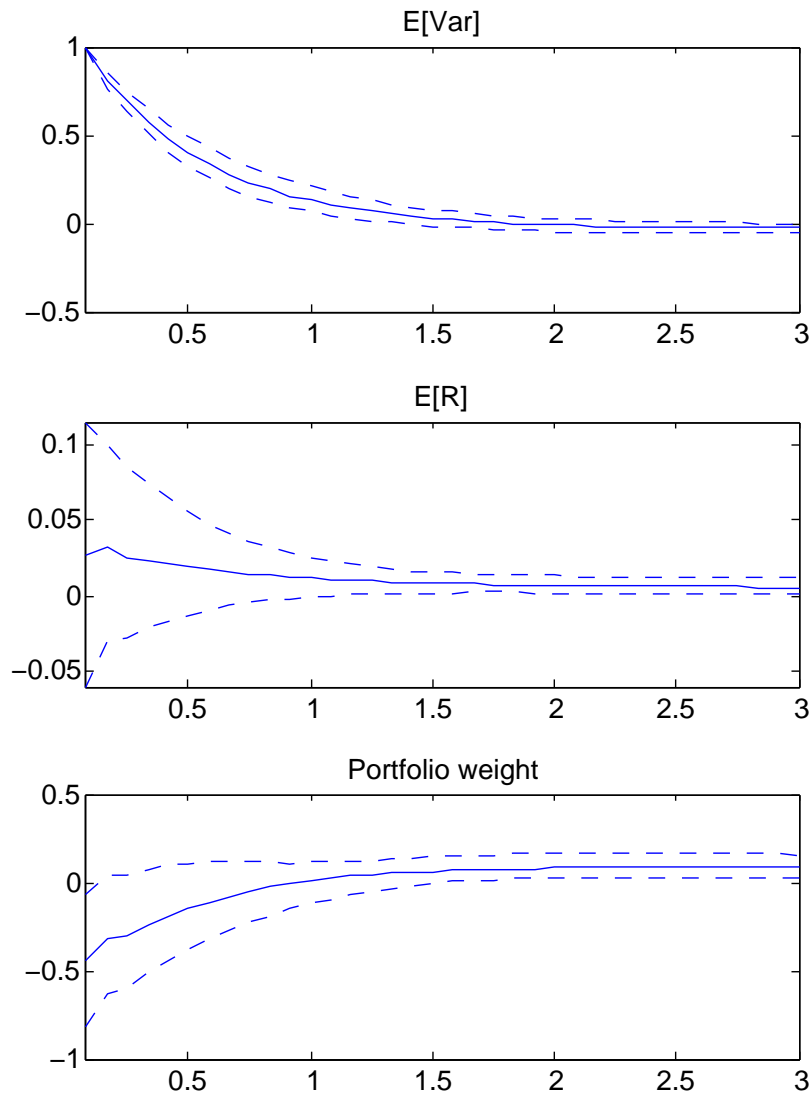
Figure 6: Results by holding period horizon. This figures plots alphas and appraisal ratios by holding period horizon given in years on the x-axis. We compute scaled portfolios using the inverse of monthly realized variance and plot the alphas and appraisal ratios for different rebalancing horizons. All numbers are annualized for ease of interpretation. The first panel does this for the market return, the middle panel uses the MVE portfolio formed from the Fama-French three factors and the lower panel adds the momentum factor. We include 90% confidence bands for alphas in dashed lines.
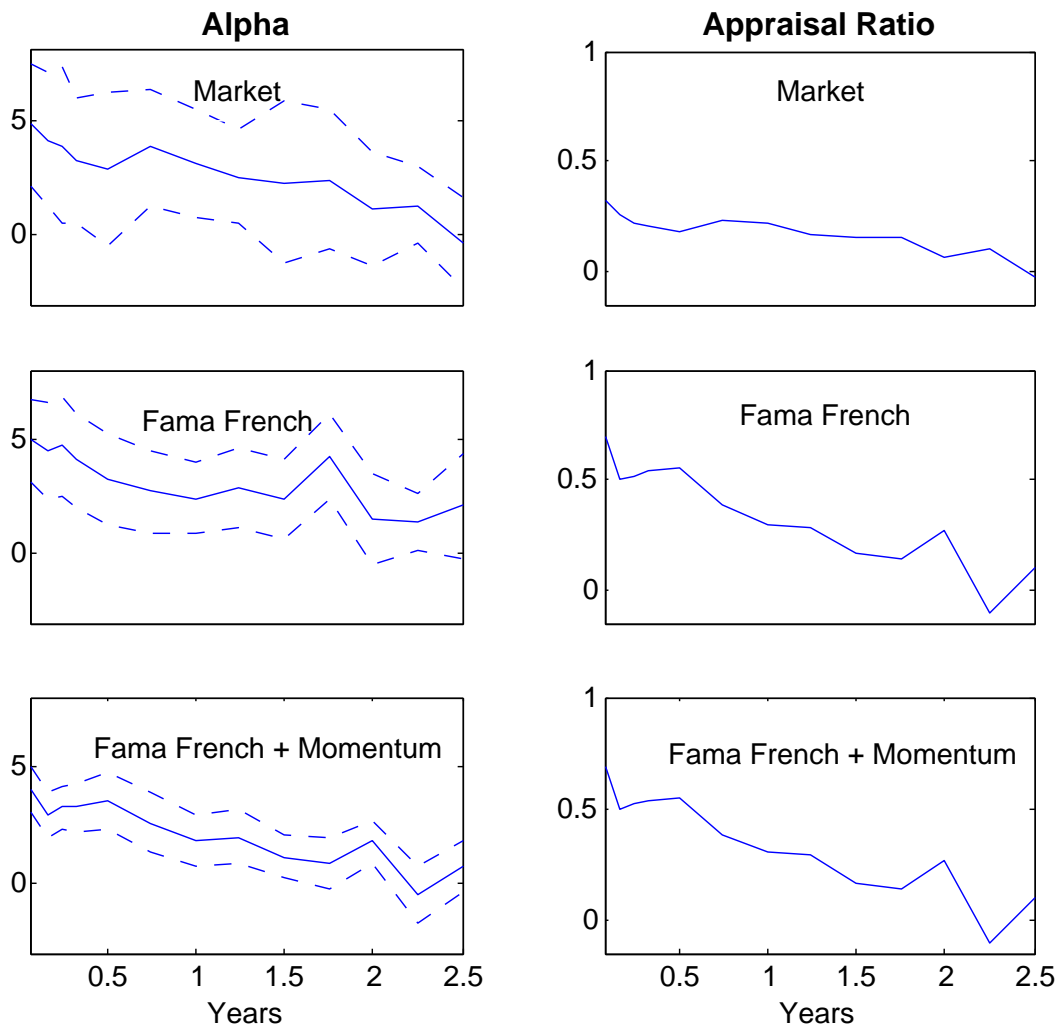
Figure 7: Equilibrium models and volatility timing

The figure plots the distribution of moments recovered from 1000 simulations of 100 year samples for the different equilibrium models. The dashed line shows the point estimate in the historical sample. The left panel shows the alpha of the volatility managed strategy, the middle panel the appraisal ratio of the volatility managed strategy, and the right panel shows the coefficient in a predictive regression of the market excess return on the previous months realized variance. Moments are recovered by replicating in the simulations exactly the same exercise we do in the data. In the first row we show the habits model of Campbell and Cochrane (1999), in the second row the rare disaster model of Wachter (2013), in the third row the long run risk model of Bansal and Yaron (2004), and in the last row the intermediary based model of He and Krishnamurthy (2012). Simulations are done using the original papers parameter calibrations.
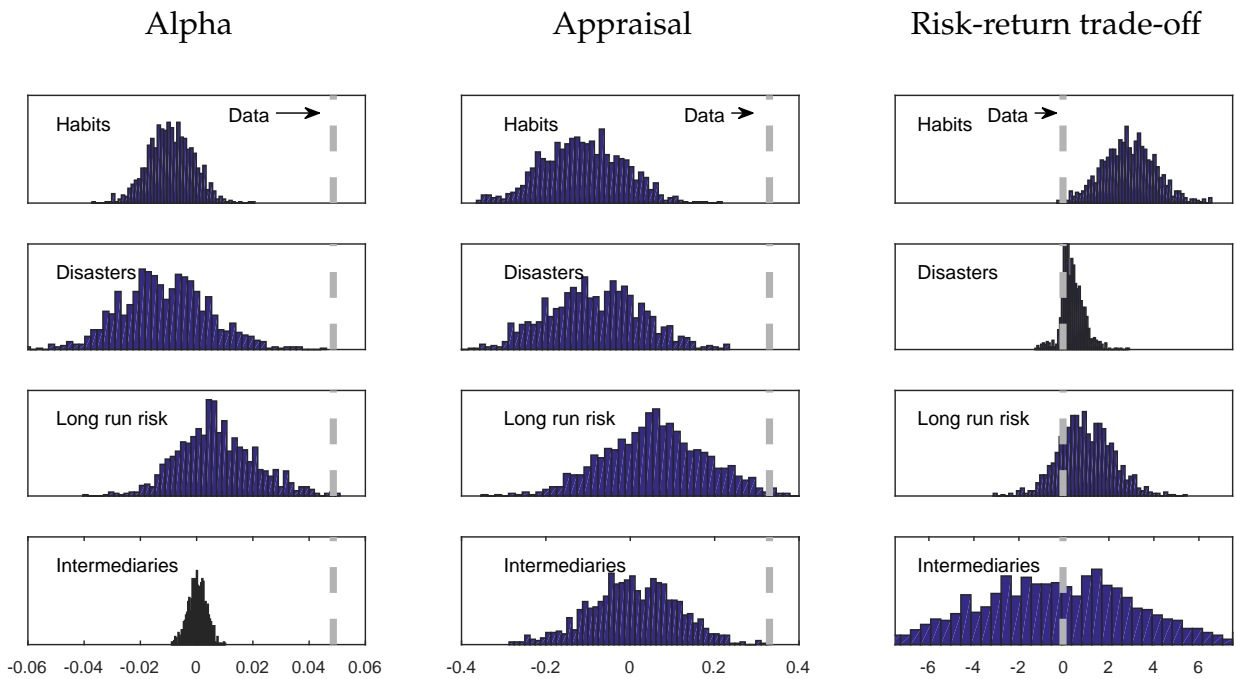
. Appendix: Not intended for publication

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 8 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 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 3), 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 unconditional 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.

Betting against beta controls: Table 9 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.

Multivariate analysis: We study whether some of the single-factor volatility timing strategies are priced by other aggregate factors. Consistent with Table 2, Tables 10 and 11 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.4 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.5 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.[21]

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 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.

---

[21]For 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 $b = E[R_{t+1}^m - R_t^f]/Var(R_t^e) \approx 2$. In our sample the highest return realization is 38% so that the SDF is never negative.

A.6 Theoretical framework: proofs and extensions

A.6.1 Conditional risk-return trade-off

We decompose variation in expected returns in terms of a component due to volatility and an orthogonal component, $\mu_t = b\sigma_t^2 + \zeta_t$, for a constant b. We assume that the process ζ_t that satisfies $E[\zeta_t|\sigma_t] = E[\zeta_t]$. The coefficient b represents the conditional risk-return trade off. Then

$$\alpha = cJ_\sigma(\gamma - b),\tag{A.6.15}$$

and alpha is positive if and only if $b < \gamma$, which means the conditional risk-return tradeoff is weaker than the unconditional risk-return tradeoff. Moreover, the weaker the conditional risk-return tradeoff, b, the higher the alpha.

A.6.2 Individual stocks

Consider a simple example where the CAPM holds, and the market portfolio dF_t has constant expected returns and variance. Consider a individual stock R with returns $dR_t = (r_t dt + \mu_{R,t})dt + \beta_R(dF_t - E_t[dF_t]) + \sigma_{R,t}dB_{R,t}$ where $dB_{R,t}$ shocks are not priced. Equation (??) implies that the volatility managed alpha is

$$\alpha_R \propto -cov\left(\sigma_{R,t}^2, \frac{1}{\beta_R^2\sigma_F^2 + \sigma_{R,t}^2}\right),\tag{A.6.16}$$

which is positive if $\beta_R > 0$ or negative if $\beta_R < 0$, but CAPM alphas are always zero.

While volatility timing can "work" for any asset with positive expected returns for which volatility is forecastable but doesn't predict returns, the alphas become economically interesting when studying systematic factors.

A.6.3 Proof of implication 1

Recognize that the fact that $\Pi(\gamma^u)$ must price factors F unconditionally immediately imply $\gamma_i^u = E[\mu_{i,t}]/E[\sigma_{i,t}]$. Analogously the fact that $\Pi(\gamma_t^*)$ must price factors F conditionally imply $\gamma_{i,t}^* = \mu_{i,t}/\sigma_{i,t}$. We can then write $\gamma_{i,t}^* = b + \zeta_t/\sigma_{i,t}$, which conditional expectation is $E[\gamma_{i,t}^*|\Sigma_t] = b + E[\zeta_t]/\sigma_{i,t}$.

We now use result (A.6.15) to substitute b and $E[\zeta_t]$ out. Specifically we use that

$$b = \gamma_i^u - \alpha_i(c_iJ_{\sigma,i})^{-1},\tag{A.6.17}$$
$$E[\zeta_t] = E[\mu_{i,t}] - bE[\sigma_{i,t}^2],\tag{A.6.18}$$

to obtain Equation (10).

Now we show that the sdf $\Pi(\gamma_t^\sigma)$ prices all volatility based strategies. We need to show

that $E\left[d\left(\Pi_t(\gamma_t^\sigma)w(\Sigma_t)\widetilde{R}_t\right)\right] = 0$,

$$E\left[d\left(\Pi_t(\gamma_t^\sigma)w(\Sigma_t)\widetilde{R}_t\right)\right] = E[w(\Sigma_t)\mu_t] - E[\gamma_t^\sigma(dF_t - E_t[])w(\Sigma_t)\widetilde{R}_t] \quad \text{(A.6.19)}$$

Using that factors are on the conditional mean-variance frontier. It is sufficient to show that the expression holds for the factors themselves. Furthermore, it is sufficient to show that the pricing equation holds for each portfolio conditional on Σ_t information.

This yields,

$$
\begin{aligned}
E\left[d\left(\Pi_t(\gamma_t^\sigma)F_t\right)|\Sigma_t\right] &= E[\mu_t|\Sigma_t] - E[\gamma_t^\sigma(dF_t - E_t[dF_t])dF_t|\Sigma_t] && \text{(A.6.20)} \\
&= b\Sigma_t + E[\zeta_t] - \gamma_t^\sigma\Sigma_t && \text{(A.6.21)} \\
&= b\Sigma_t + E[\zeta_t] - (b + E[\zeta_t]\Sigma_t^{-1})\Sigma_t && \text{(A.6.22)} \\
&= 0, && \text{(A.6.23)}
\end{aligned}
$$

where in the last line we used that $\gamma_{i,t}^\sigma = E[\gamma_{i,t}^*|\Sigma_t] = b + E[\zeta_t]/\sigma_{i,t}$. This proves implication 1.

A.6.4 Spanning the unconditional mean-variance frontier with volatility managed portfolios

The price of risk in (7) is also the unconditional mean-variance-efficient portfolio from the perspective of an investor that can measure time-variation in volatility but not variation in ζ_t. It can be decomposed in terms of constant positions on the buy-and-hold factors and the volatility managed factors.

Implication 2. *The unconditional mean-variance-efficient portfolio spanned by conditional information on volatility can be replicated by a constant position of the factors and the volatility managed factors $[dF; dF^\sigma]$,*

$$E[w_i^*(\mu_t, \Sigma_t)|\Sigma_t] \propto \left[\gamma_i^u - \frac{\alpha_i}{c_i}J_{\sigma,i}^{-1}; \frac{\alpha_i}{c_i}J_{\sigma,i}^{-1}\frac{E[\sigma_{i,t}^2]}{c_i}\right]. \quad \text{(A.6.24)}$$

These weights are simple functions of our strategy alpha. Assuming the market portfolio is on the conditional mean-variance frontier, we can plug numbers for the market portfolio to have a sense of magnitudes. We get $[0.14; 0.86]$ for the weights on the market and the volatility managed market portfolio. Empirically, our volatility managed portfolio get close to be unconditionally mean-variance efficient because the relationship between return and volatility is so weak.

A.6.5 Correlated factors

Our approach can be easily be extended to the case factors are correlated. Let the factors variance-covariance matrix be block diagonal. It can be decomposed in N blocks as $\Sigma_t \Sigma_t' = diag\left(\left[H_1 \sigma_{1,t}^2 ..., H_N \sigma_{N,t}^2\right]\right)$, where $\sigma_{n,t}^2$ are scalars, H_n are constant full rank matrixes.

Given this factor structure in factor variances (see Section 3.5 to see that this is a good description of the data), we can apply our analysis to "block-specific" mean-variance efficient portfolios constructed as follows. For a block n, let $df_{n,t}$ be the vector of factor returns and $\mu_{n,t}$ be the vector of expected excess returns. Form MVE portfolios as $df_{n,t}^{MVE} \equiv rdt + \mu_{n,t}' H_n^{-1} (df_{n,t} - rdt)$, which is exactly the procedure we follow in Section 2.5.

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^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).
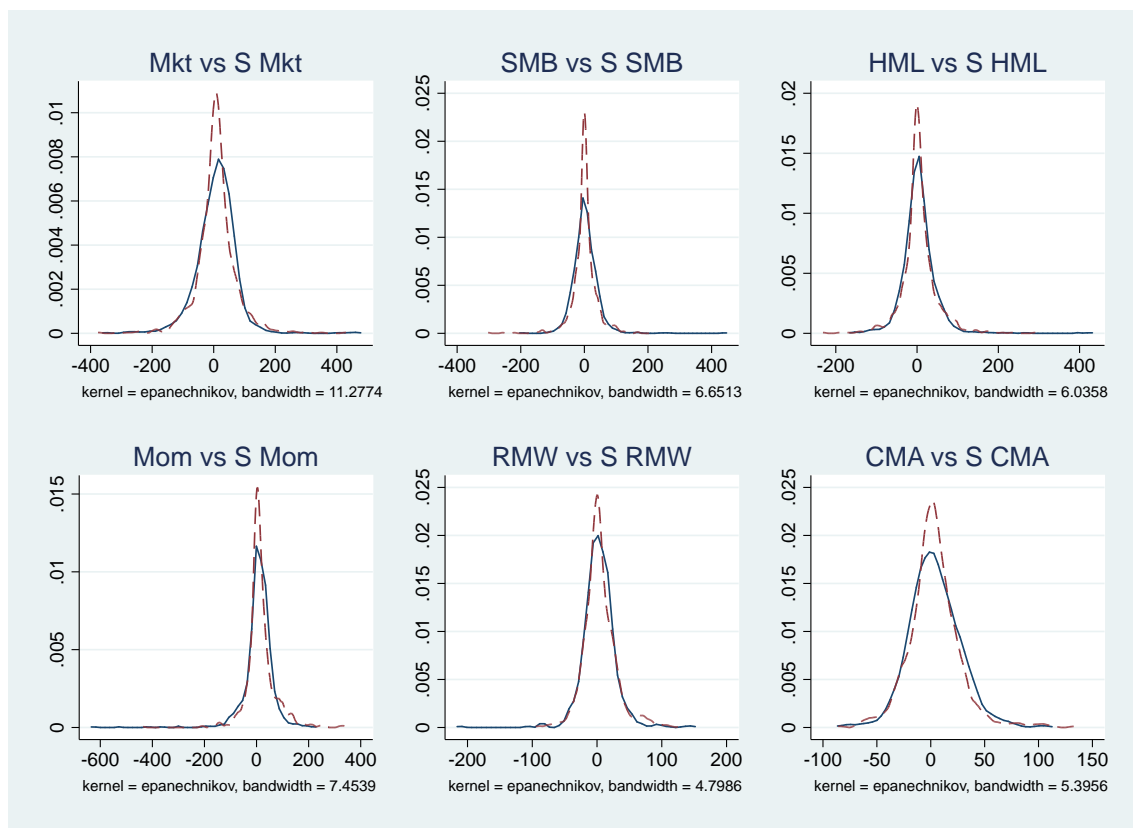
B. Additional Tables

Table 8: 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 $f^\sigma_{t+1} = \frac{c}{E_{t-1}[RV_t^2]} f_t$. 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) Mkt^σ	(2) SMB^σ	(3) HML^σ	(4) Mom^σ	(5) RMW^σ	(6) CMA^σ	(7) MVE^σ	(8) FX^σ
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 9: 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) Mkt^σ	(2) SMB^σ	(3) HML^σ	(4) Mom^σ	(5) RMW^σ	(6) CMA^σ	(7) MVE^σ
MktRF	0.60						
	(0.05)						
BAB	0.09	0.01	0.02	-0.07	-0.13	-0.06	0.04
	(0.06)	(0.05)	(0.05)	(0.04)	(0.02)	(0.02)	(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	-0.77	2.05	13.52	3.97	0.94	4.10
	(1.80)	(1.10)	(1.15)	(1.86)	(0.89)	(0.71)	(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 10: 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) Mkt^σ	(2) SMB^σ	(3) HML^σ	(4) Mom^σ	(5) MVE^σ
MktRF	0.70	-0.02	-0.10	0.16	0.23
	(0.05)	(0.01)	(0.02)	(0.03)	(0.02)
HML	-0.03	-0.02	0.63	0.09	0.08
	(0.05)	(0.04)	(0.05)	(0.05)	(0.02)
SMB	-0.05	0.63	-0.00	-0.10	-0.15
	(0.06)	(0.08)	(0.05)	(0.04)	(0.02)
Mom	0.25	0.01	0.06	0.54	0.30
	(0.03)	(0.03)	(0.04)	(0.05)	(0.02)
Constant	2.43	-0.42	1.96	10.52	4.47
	(1.60)	(0.94)	(1.06)	(1.60)	(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 11: 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) Mkt^σ	(2) SMB^σ	(3) HML^σ	(4) Mom^σ	(5) RMW^σ	(6) CMA^σ	(7) MVE^σ	(8) $\mathrm{MVE2}^\sigma$
MktRF	0.79	0.03	-0.06	0.12	0.02	0.02	0.26	0.23
	(0.05)	(0.03)	(0.03)	(0.04)	(0.02)	(0.01)	(0.03)	(0.02)
HML	0.11	0.09	1.03	0.15	-0.21	0.03	0.16	0.05
	(0.09)	(0.06)	(0.08)	(0.09)	(0.04)	(0.03)	(0.04)	(0.03)
SMB	0.02	0.75	-0.05	-0.12	-0.02	-0.03	-0.15	-0.09
	(0.05)	(0.05)	(0.04)	(0.07)	(0.03)	(0.02)	(0.03)	(0.02)
Mom	0.15	-0.01	0.05	0.64	-0.00	-0.02	0.32	0.23
	(0.03)	(0.03)	(0.03)	(0.08)	(0.02)	(0.02)	(0.03)	(0.02)
RMW	0.15	0.23	-0.56	-0.04	0.64	-0.18	0.01	0.04
	(0.06)	(0.07)	(0.08)	(0.08)	(0.06)	(0.04)	(0.04)	(0.03)
CMA	0.04	0.00	-0.28	-0.25	-0.00	0.63	-0.04	0.14
	(0.12)	(0.07)	(0.10)	(0.11)	(0.06)	(0.05)	(0.06)	(0.04)
Constant	0.18	-1.68	4.16	12.91	3.21	1.07	4.00	3.03
	(1.87)	(1.25)	(1.44)	(2.17)	(0.81)	(0.72)	(1.02)	(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