AUTHORS

Jamil Baz
Managing Director
Head of Client Solutions
and Analytics

Josh Davis
Managing Director
Portfolio Manager

Steve Sapra
Executive Vice President
Client Solutions and Analytics

**Jerry Tsai**
Vice President
Quantitative Research Analyst

Normane Gillmann
Quantitative Research Analyst

A Theoretical Framework for Equity-Defensive Strategies

Baz et al. (2017) note that investors have come to accept four general approaches to diversifying and helping to mitigate equity risk: 1) long Treasuries, 2) trend-following, 3) tail risk hedging and 4) alternative risk premia diversifiers, such as carry and value strategies.[1] In that paper, the authors show that each of these strategies has merit for inclusion in a comprehensive risk mitigation portfolio, as combining positive expected return strategies that are negatively correlated with equities not only provides higher degree of confidence with respect to mitigating equity risk, but also may deliver higher returns. We expand upon their research here and develop a comprehensive theoretical framework for defensive portfolio construction.

One reason investors hold these types of investments is the belief that they will provide some degree of protection in down equity markets. The key trade-off among these strategies is the "cost" of each approach versus the certainty that it will work in times of crisis.[2] We can, therefore, frame them in the context of a theoretical risk mitigation frontier, as shown in Exhibit 1, where the y-axis is the expected return and the x-axis represents the uncertainty that a strategy will work in an equity risk-off event.

Exhibit 1: Risk/return trade-off for defensive strategies

Source: PIMCO

1 These include currency, fixed income and commodity carry strategies, and other equity risk premia factors, like quality and value.
2 The term "certainty," which is used throughout this paper, does not imply or connote a guaranteed outcome, but rather is used to express the relative confidence with which a strategy can be effective.

Tail risk hedging is contractually linked to equity market performance, providing a direct hedge with a high degree of certainty in severe downturns. However, this high degree of confidence comes at the cost of several headwinds, including negative exposure to the well-established equity risk premium and volatility risk premium, which often results in negative expected returns.

Long Treasuries are perhaps the most prominent example of a simple and robust positive-expected-return equity-risk-mitigation strategy. Baz et al. (2018) show that, since the mid-1990s, being passively long Treasuries has consistently delivered a negative average correlation with equities. Furthermore, the authors show that exposure to U.S. Treasuries has often provided investors with positive absolute returns during large equity declines, particularly when those declines are associated with a contraction in economic activity. Historically, however, the stock-bond correlation has moved considerably over time (Campbell et al. 2017), which increases uncertainty about using Treasuries for diversification. Furthermore, Wright (2011) and Adrian et al. (2017) show that the term premium has compressed toward zero over the past several decades, casting further doubt on the asset's forward-looking return-generating potential.

Trend-following strategies benefit from persistent trends in prices across major markets. Fung and Hsieh (2001) show that these strategies exhibit payoff profiles that resemble long-volatility strategies. Due to dynamic trading in the underlying assets, they are not subject to the same implied volatility headwinds faced by option-based strategies such as tail risk hedging, making trend-following an efficient way to capture the long-volatility profile. However, this efficiency comes at the cost of uncertainty around returns and equity risk mitigation, as trend-following may deliver negative returns when market reversals are abrupt and/or frequent. Historically, trend-following has delivered positive returns in equity market drawdowns, but there is substantial variation around this average.

Alternative risk premia strategies are broadly defined by having 1) limited exposure to traditional equity and bond risk premia, 2) a clear economic rationale for their existence, 3) empirical validation based on historical data, and 4) implementation that typically requires the use of leverage, via shorting and derivatives. Given that these strategies are designed to exhibit limited exposure to equity markets, they

are a natural candidate for risk mitigation portfolios. However, although the correlation of alternative risk premia with equities tends to be near zero, there is wide variation in the conditional correlations of these strategies with equities. As a result, there is less certainty that these strategies will diversify when equity returns are negative.

1. UNDERSTANDING RISK IN DEFENSIVE STRATEGIES

The goal of a defensive portfolio is effectively twofold. First, the allocation to defensive assets should have positive expected returns in aggregate. Second, it should provide some degree of protection if there is a significant impairment to the overall portfolio. Such impairment is almost always related to large equity market corrections. However, as discussed in the previous section, defensive assets embody a range of risk/return trade-offs, in which higher-returning strategies generally have a lower expected effectiveness in terms of protecting against equity market sell-offs. This raises the obvious question: What is the right risk measure for defensive strategies?

While Exhibit 1 shows a stylized example of the risk/return trade-off, we need a more precise measure of risk for this cohort of investment styles. Obvious measures are metrics that describe each strategy's covariance with equity risk. Assets with a greater and more reliable negative covariation with equities should generally provide more effective protection against broad equity market sell-offs. Equity beta, downside beta and conditional returns are all reasonable measures of equity market covariation. However, the problem is more complicated than this because such relationships may not be linear. For example, certain strategies may provide little protection against modest equity downturns but substantial protection in large market drawdowns. Out-of-the-money (OTM) equity puts are an example of such an asset. In modestly declining markets, OTM puts respond only modestly, but their protection can ramp up significantly in large sell-offs. In technical terms, we need to assess each strategy's convexity.

To assess the overall risk properties of defensive strategies, accounting for both small and large market moves, we consider the following time-series regression on each strategy's return:

$$r_{i,t} = \alpha + \beta_{i,1}r_{m,t} + \frac{1}{2}\beta_{i,2}\left(r_{m,t}^2\big|r_{m,t} < 0\right) +$$

$$\frac{1}{2}\beta_{i,3}\left(r_{m,t}^2\big|r_{m,t} \geq 0\right) + \varepsilon_{i,t} \qquad (1)$$

The word protection used throughout this paper does not imply or connote a guarantee that the strategy will prevent loss, but instead is used to express the relative confidence with which a strategy can effectively mitigate downside risks.

where r_i is strategy i's excess return over cash, r_m is the excess return on the S&P 500 index and ε_i is an uncorrelated residual term. $\beta_{i,1}$ captures each strategy's general (linear) sensitivity to the equity market, while $\beta_{i,2}$ and $\beta_{i,3}$ capture its sensitivity to large negative and positive equity market movements. The latter two parameters measure the strategy's statistical convexity in up and down markets.[3]

We estimate Equation 1 for 1) long Treasuries, 2) trend-following, 3) tail risk hedging, 4) carry and 5) value strategies using daily return data from 1 March 1994 to 31 December 2018. The carry and value strategies are constructed using currencies, rates and commodities. Within each asset class, the carry strategy takes long positions in assets with high carry and short positions in assets with low carry, while the value strategy takes long positions in assets with high value and short positions in assets with low value. Details of all backtests are provided in Appendix 1. All strategies, except for tail risk hedging, can be levered up or down depending on the investor's needs and we therefore scale each to have 5% annualized volatility.

Exhibit 2 shows the results for estimating Equation 1. β_1 and β_2 are likely to be of the most interest to potential defensive strategy investors since they speak directly to downside protection with respect to equity risk. β_1 is negative and highly statistically significant for long Treasuries and tail-risk-hedging strategies. As expected, the t-statistics on tail-risk-hedging strategies are materially higher, given that they provide a direct hedge against equity risk. Furthermore, all strategies except carry and value are characterized by positive convexity in down markets, as measured by β_2. For trend-following strategies, it has been well-documented that they are generally able to capture significantly down-trending equity market momentum.[4]

Exhibit 2: Estimated beta and convexity of defensive strategies

|	α	β_1	β_2	β_3	R^2
Long Treasuries	0.006%	-0.06	0.85	0.20	5.8%
t-stat	1.49	-11.08	2.49	0.63	
Trend-following	0.012%	-0.01	3.33	-2.64	6.6%
t-stat	3.18	-2.11	10.16	-8.90	
Tail risk hedging	-0.011%	-0.20	7.83	-6.10	86.1%
t-stat	-5.65	-76.90	47.21	-40.69	
Carry	0.021%	-0.01	-0.79	0.16	0.1%
t-stat	5.39	-1.24	-2.36	0.53	
Value	0.013%	0.00	-1.22	1.12	0.7%
t-stat	3.25	-0.06	-3.47	3.52	

Hypothetical example for illustrative purposes only.
Source: PIMCO and Bloomberg as of December 2018. Regressions are estimated using daily unannualized data from 1 March 1994 to 31 December 2018.

The importance of downside convexity is also illustrated in Exhibit 3, which shows the expected payoffs to each of the strategies as a function of the daily equity market return, based on the parameter estimates in Exhibit 2. For modestly negative equity market returns, these strategies show small differences in terms of expected returns. However, for large equity market drawdowns, the convexity of a long-put-option position begins to dominate, resulting in substantially higher expected returns for the tail risk hedging relative to other defensive strategies.[5]

Exhibit 3: Expected returns for defensive strategies versus downside equity return
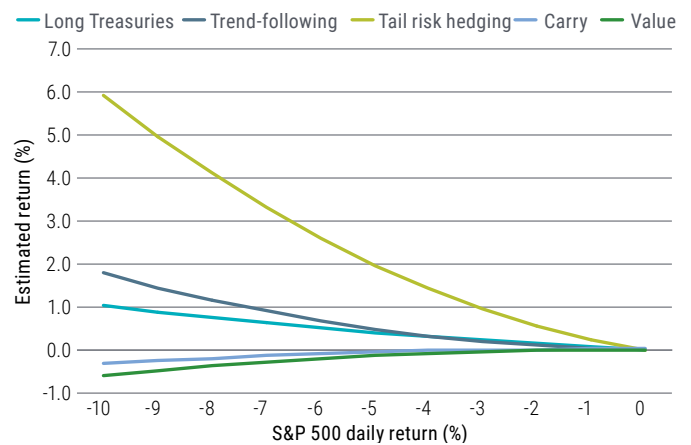
Source: PIMCO and Bloomberg as of December 2018. **Hypothetical example for illustrative purposes only.**

So far, the discussion has focused on estimates of conditional expected returns. A natural question that follows is how certain we are that the actual returns will be close to their predicted levels. To assess this, we turn to the r-squared statistics in Exhibit 2. While long Treasuries and trend-following strategies have negative equity betas and positive downside convexity, the r-squared values of these regressions are low at around 6%. That is, the actual returns on these strategies could be far from the positive predicted values, and therefore the equity hedge is far from guaranteed. This is even more true for carry and value strategies which have r-squared values around zero. Therefore, these strategies can be viewed as good equity diversifiers but not necessarily true hedges against equity market risk. Tail risk hedging, on the other hand, is characterized by a very high r-squared around 85%, implying that the actual returns are unlikely to deviate materially from their expected values. As such, the likelihood of tail risk hedging protecting against downside equity risk is significantly higher, making it the most

---

3 The reason for the ½ terms in front of the convexity components is to be consistent with the second-order Taylor series expansion. It does not change the economic significance in any way.

4 It can be shown that trend-following strategy long long-dated variance and short short-dated variance. Hence the strategy is exposed to crash risk as it means a spike in short-dated variance.

5 One should keep in mind that tail risk hedging has volatility around 6%, compared with 5% for the other strategies.

effective in the defensive strategy tool kit. However, this greater degree of protection comes at a cost: low or even negative average returns.

2. A THEORETICAL FRAMEWORK

In this section, we develop a theoretical framework for defensive portfolio construction. In many ways, it is not materially different from the standard Markowitz (1952) mean-variance optimization (MVO) problem, except that the choice of assets and their respective allocation are affected by the investor's need for downside equity market protection.

In our framework, investors care about both the conditional and the unconditional properties of the defensive portfolio. Specifically, investors optimize the standard risk/return trade-off between the portfolio's unconditional expected return and volatility while at the same time constraining the portfolio to have certain downside protection properties, measured by its conditional beta. The conditional beta refers to the asset's beta in down equity markets, which may be quite different from its beta in "normal" markets, depending on the asset class or investment style under consideration. As described in the previous section, the constraint on conditional beta is particularly relevant because asset classes and styles with the highest unconditional expected returns tend to have the weakest downside equity protection, and vice versa. This observation creates a risk/return trade-off in which a higher degree of downside protection must be "paid for" via lower returns.

Formally, we consider an investor who chooses a set of weights, w, to maximize the defensive portfolio's unconditional expected return subject to constraints on its unconditional volatility and conditional equity beta:

$$max_w \ E[R_P] = w'\mu \qquad (2)$$

$$0.5 \ w'\Sigma w = 0.5\sigma_p^2 \qquad (2.1)$$

$$w'\beta_c \leq \bar{\beta}_c \qquad (2.2)$$

where μ is a vector of unconditional expected excess returns, Σ is the unconditional covariance matrix, σ_p^2 is the targeted unconditional portfolio variance, β_c is a vector of conditional betas and $\bar{\beta}_c$ is the conditional beta target of the portfolio. Equations 2.1 and 2.2 thus represent constraints on the portfolio problem imposed by the investor, reflecting both the investor's willingness to assume volatility risk and the need for

downside equity market protection. By modeling the equity beta component of the problem as **conditional**, we are explicitly accounting for the fact that the asset classes in a defensive portfolio can have different conditional and unconditional behavior. We specifically exclude the market portfolio from the optimization problem and only consider a portfolio of defensive assets.[6]

A detailed solution to the problem is provided in Appendix 2. Appendix 2.1 shows properties for three special portfolios: 1) the unconstrained MVO portfolio w^{MVO}, which has a conditional beta β^{MVO} and Sharpe ratio SR_{MVO}, 2) the minimum-variance portfolio with a unit conditional beta (unit-beta portfolio), w^B, which has an unconditional variance $\sigma_{w^B}^2$ and a Sharpe ratio denoted by SR_{w^B} (hence, $\bar{\beta}_c w^B$ is the minimum variance portfolio with conditional beta $\bar{\beta}_c$) and 3) the zero-beta MVO portfolio ($w^{MVO} - \beta^{MVO}w^B$), which adjusts the weights of the unconstrained MVO proportionally to the unit-beta portfolio to achieve a zero net equity beta.

We are, of course, most interested in the case in which the conditional beta constraint (2.2) is binding. In this case, the optimal portfolio is given by

$$w = \bar{\beta}_c w^B + c(w^{MVO} - \beta^{MVO}w^B), \qquad (3)$$

where c is a constant. Equation (3) shows that the optimal portfolio is a weighted sum of the unit-beta portfolio w^B and the zero-beta MVO portfolio ($w^{MVO} - \beta^{MVO}w^B$). For intuition, the solution can be thought of as a two-step procedure: First, go short the minimum-variance unit-beta portfolio in order to achieve the conditional beta target, $\bar{\beta}_c$, then hold some amount in the zero-beta MVO portfolio to maximize the unconditional expected return. The degree of leverage, c, is a function of the overall unconditional volatility target, σ_p, and the fraction of the risk budget remaining after the conditional beta constraint has been satisfied.

As shown in Appendix 2.2, the expected return of the optimal portfolio is given by

$$E[R_p] = \mu'(\bar{\beta}_c w^B) + (\sigma_p^2 - \bar{\beta}_c^2\sigma_{w^B}^2)^{1/2}(SR_{MVO}^2 - SR_{w^B}^2)^{1/2}. \qquad (4)$$

A close inspection of Equation (4) yields some powerful insights. The first term is the unconditional expected excess return on the beta portfolio. This term arises due to the conditional beta constraint and directly corresponds to the first

6   The reason for excluding the market portfolio is that, when the market portfolio is part of the opportunity set, the optimal solution combines the MVO with a short position in equities. Though this may be a reasonable strategy, investors tend not to short equities in the defensive asset portfolio. Rather, the beta target is achieved only with defensive assets.

term in (3). Under the Capital Asset Pricing Model (CAPM), the expected excess return on the unit-beta portfolio will equal the expected equity excess return ($\mu' w^B = \mu_{\text{equity}}$); thus, the expected return on the beta portfolio, $\bar{\beta}_c \mu_{\text{equity}}$, would be negative for $\bar{\beta}_c < 0$. However, this term could be less negative or even positive when the CAPM does not hold (as we assume in our framework) since the investor can hold assets with both negative betas and positive unconditional expected returns. Nonetheless, even though the value of this term can be positive in our framework, it is useful to think of this first term as a cost, or "insurance premium", associated with achieving a negative conditional beta.[7]

The variance of the beta portfolio is $\bar{\beta}_c^2 \sigma_{w^B}^2$. Therefore, the second term in the parentheses measures the amount of the volatility budget remaining after the beta constraint has been satisfied. We refer to this term as the "risk budget" because it reflects the quantity of risk that can be allocated to the return-enhancing zero-beta MVO portfolio once the beta constraint has been achieved. If, for example, the investor desires to achieve a highly negative beta, this will have the effect of decreasing the amount of the risk budget that can be used for increasing overall returns.

The last term in parentheses is the difference between the squared unconditional Sharpe ratio of the unconstrained MVO and that of the beta-hedging portfolio. We label this term "efficiency" because it reflects the extent to which the defensive assets collectively embody diversifying risk/return properties relative to the pure hedging portfolio. Efficiency is scaled by the risk budget in Equation 4 because only variance that is not used to satisfy the targeted beta can be used to exploit the return-enhancing properties of the defensive assets.

This intuitive decomposition leads us to a general rule of thumb to describe the expected return on a portfolio of defensive assets:

$$E[R_p] = -insurance\ premium + (risk\ budget \times efficiency), \quad (5)$$

where the insurance premium is $-\mu'\left(\bar{\beta}_c w^B\right)$, the risk budget is $\left(\sigma_p^2 - \bar{\beta}_c^2 \sigma_{w^B}^2\right)^{1/2}$ and the efficiency is $\left(\text{SR}_{\text{MVO}}^2 - \text{SR}_{w^B}^2\right)^{1/2}$.

Equation 5 tells us that the defensive portfolio problem is ultimately about balancing the cost of hedging equity risk with the desire to generate positive unconditional returns. It shows that positive expected returns on a defensive portfolio can be achieved if one can allocate a sufficient amount of risk (risk budget) to assets with positive Sharpe ratios (efficiency) and one does not "overpay" for the equity hedging component (insurance premium). Under CAPM, the insurance component will generate a negative expected return contribution because $\bar{\beta}_c \mu_{\text{equity}}$ <0. However, according to Equation (5), we can do better than this when CAPM does not hold, as the insurance cost becomes less expensive and diversification adds return via the efficiency measure.[8]

3. EMPIRICAL RESULTS

To better understand the implications of our model, we consider five different asset classes and styles: long Treasuries, trend-following, tail risk hedging, carry and value. Exhibit 4 shows the excess return moment assumptions for each of these strategies. The unconditional moments for the S&P 500 and tail-risk-hedging strategy were calculated using quarterly data from first-quarter 1994 to fourth-quarter 2018.

For long Treasuries, trend-following, carry and value strategies, we believe it is more appropriate to use reasonable forward-looking views that account for the realities of today's markets. To reflect interest rates that are much lower today than over our historical sample, we assume a very modest Sharpe ratio of 0.07 (or an excess return of 0.4%) for long Treasuries.[9] For trend-following, carry and value strategies, we assume the unconditional Sharpe ratios to be 0.25, 0.50 and 0.25, respectively, which implies expected excess returns of 1.3%, 2.5% and 1.3%, respectively.[10]

For long Treasuries, trend-following, carry and value strategies, we assume a 5% annualized unconditional volatility, as they can generally be implemented using derivatives and therefore various volatility targets are relatively easy to achieve. Finally, we note the negative Sharpe ratio for tail risk hedging, which reflects the fact that passively buying puts has historically been a negative-returning strategy. The conditional moments for the S&P 500 and all strategies are calculated using quarters in which the S&P 500 excess returns are less than -3.75% (-15% annualized).

---

7 This term can be positive if the unit-beta portfolio loads heavily on negative beta and positive expected return strategies. In general, we should expect the insurance component to be greater than $\bar{\beta}_c \mu_{\text{equity}}$, but it can be negative if we use assets like put options. We can think of this term as a cost even if the beta portfolio has a positive expected return, as it still has a worse risk/return trade-off than the unconstrained MVO portfolio.

8 See Appendix 3 for a two-risky-asset example.

9 The Sharpe ratio of 0.07 reflects PIMCO's 5-year capital market assumption for long Treasuries.

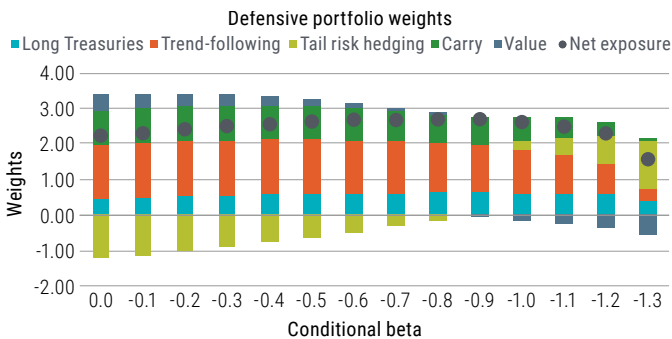10 See Exhibit 12 in Appendix 1 for the raw return moments.

Exhibit 4: Assumptions on conditional and unconditional excess return moments

	Unconditional moments			Conditional moments			
	Expected excess return	Volatility	Sharpe ratio	Volatility	Equity covariance	Equity beta	Equity correlation
S&P 500	5.2%	15.6%	0.33	10.3%	1.1%	1.00	1.00
Long Treasuries	0.4%	5.0%	0.07	6.7%	-0.4%	-0.42	-0.64
Tail risk hedging	-2.4%	5.4%	-0.45	7.1%	-0.7%	-0.64	-0.92
Trend-following	1.3%	5.0%	0.25	6.4%	-0.5%	-0.43	-0.70
Carry	2.5%	5.0%	0.50	5.7%	0.0%	-0.04	-0.06
Value	1.3%	5.0%	0.25	5.1%	0.3%	0.26	0.52

Source: PIMCO and Bloomberg as of December 2018. Data is quarterly from 1 March 1994 to 31 December 2018. Expected returns are in excess of cash. Conditional moments are based on S&P 500 quarterly excess returns less than -3.75%. Values in blue reflect overrides to the historical data.

Exhibit 5 shows the allocations to each of the assets/styles in Exhibit 4 as a function of the targeted conditional beta for a 10% volatility portfolio. Allocations on the left side of the exhibit reflect portfolios for which equity hedging is less of a consideration and thus have the least negative conditional betas. As expected, such portfolios are made up of higher-returning assets since the need for equity hedging is minimal. In fact, these portfolios actually hold a short position in the tail-risk-hedging portfolio (effectively selling puts) as a means of increasing the unconditional expected return. Moving right on the graph, the importance of equity hedging increases, and the impact of the conditional beta constraint becomes more pronounced. As a result, the investor allocates to more "reliable" hedging sources, such as long Treasuries and long puts, as the conditional beta constraint becomes increasingly negative.[11] Importantly, under this framework, as market conditions change investors can update their assumptions on expected returns and betas, etc., to rebalance the defensive portfolio weights.

equity hedge. Exhibit 6 shows the decomposition of the portfolio return based on Equations 4 and 5. With a zero conditional beta, the investor can avoid using any of their risk budget for equity hedging and as a result holds a portfolio that is perfectly proportional to the zero-beta MVO portfolio. In this case, all of the return comes from the combination of efficiency and risk budget described in Equation (5). As the conditional beta constraint becomes negative, the investor must use an increasingly large fraction of her risk budget purely for hedging purposes, resulting in a drag on expected return from the increasing insurance premium. On the far-right side of Exhibit 6, the unconditional return actually turns negative, as the majority of the portfolio is made up of "insurance assets" such as put buying. At the extreme, 100% of the investor's risk budget is used to pay the insurance premium and she holds a portfolio that is perfectly negatively proportional to the unit-beta portfolio. Because this portfolio is dominated by equity put options (the unit-beta portfolio is dominated by shorting put options), the expected return on the portfolio is negative.

Exhibit 5: Optimal defensive portfolio weights

Defensive portfolio weights

■ Long Treasuries ■ Trend-following ■ Tail risk hedging ■ Carry ■ Value ● Net exposure

Source: PIMCO

Exhibit 6: Decomposition of unconditional expected excess returns

Unconditional expected return decomposition

■ Insurance premuim ■ Efficiency x Budget ◆ Total

Source: PIMCO

As expected, more negative conditional betas result in portfolios with lower unconditional returns, reflecting the "price" investors must pay to increase the size and effectiveness of the

11 While trend-following has a larger weight than that of long Treasuries, the position on the latter is more stable across various beta targets. For example, as the beta target moves from 0 to -1, the weight on trend-following drops by around 20% while that on long Treasuries increases by around 33%.

Exhibits are hypothetical examples for illustrative purposes only.

Exhibit 7: Defensiveness/cost trade-off

Source: PIMCO and Bloomberg as of December 2018

A natural question about the optimal defensive portfolios is how much benefit we get from forming portfolios relative to using individual defensive strategies. To evaluate this, Exhibit 7 plots the defensiveness versus the cost for the optimal portfolios and for the five individual strategies. To compare across portfolios/strategies with different volatilities, we use the unconditional Sharpe ratio as a measure of cost and conditional correlation as a measure of defensiveness.[12]

Indeed, the benefit of forming portfolios is sizable. Consider the trend-following strategy, which is the closest to the frontier, as an example. This strategy has a conditional correlation of -0.70 and an unconditional Sharpe ratio of 0.25. In contrast, the optimal portfolio with a beta target of -0.9 achieves a correlation of -0.68 but an unconditional Sharpe ratio of 0.36.[13] Hence, the optimized portfolio achieves a roughly 40% improvement in Sharpe ratio, with a similar overall level of defensiveness compared to trend-following on a stand-alone basis.

Thus far, we have considered the salient properties of the optimal defensive portfolio given a conditional beta target. The conditional beta, of course, is calculated based on covariance/correlation estimates and thus is subject to estimation error. To assess the uncertainty around these estimates, we calculate the 95% confidence intervals around the conditional correlations.[14]

This is shown in Exhibit 8, which illustrates that even after targeting a specific equity beta, the defensive properties of these portfolios are far from guaranteed. This is particularly true for portfolios derived from less negative beta constraints. For example, Exhibit 8 shows that to obtain a negative conditional correlation with 95% confidence, the conditional beta target would need to be below -0.6. Likewise, targeting a zero conditional correlation could produce realized correlations approximately between -0.5 and +0.5 – perhaps a wider range than some investors are comfortable with. If the investor desires to obtain a high degree of defensiveness with high confidence – say, a conditional correlation below -0.5 with 95% confidence – the investor needs to target a beta below -1.1.

Exhibit 8: 95% confidence interval for conditional correlation estimates
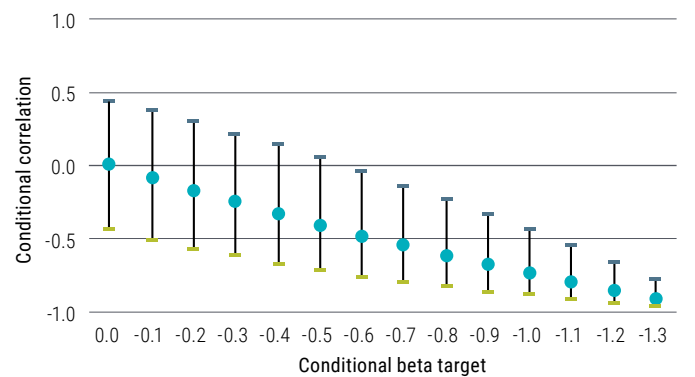
Source: PIMCO

12 Given the optimal weights for a beta target (β_i), we can calculate the portfolio's conditional volatility, and the conditional Pearson correlation is β_i x (conditional volatility of S&P 500) ÷ (conditional volatility of portfolio).

13 We should note that it is not guaranteed that the individual strategies will lie below the optimal portfolio curve in the unconditional Sharpe ratio-conditional correlation space. This is because we have a conditional beta and not a conditional volatility target in the optimization problem. Hence, if the conditional volatility of the optimal portfolio is too high, then the conditional correlation could be less negative, even if the conditional beta is more negative.

14 We first convert conditional beta to conditional correlation, then the confidence interval of conditional correlation is constructed using Fisher transformation. The $(1 − \alpha)$% confidence interval for is $\left[\tanh \left(\operatorname{artanh}(r) − \frac{z_{\alpha/2}}{\sqrt{n-3}} \right), \tanh \left(\operatorname{artanh}(r) + \frac{z_{\alpha/2}}{\sqrt{n-3}} \right) \right]$, where r is the sample correlation and n the sample size. tanh and artanh are the hyperbolic tangent and the inverse hyperbolic tangent functions, respectively.

Exhibits are hypothetical examples for illustrative purposes only.

Exhibit 9: Portfolio performance historical drawdown periods

						Beta target									Individual (rescaled 10% unconditional vol)					S&P 500
	0	-0.1	-0.2	-0.3	-0.4	-0.5	-0.6	-0.7	-0.8	-0.9	-1	-1.1	-1.2	-1.3	LTE	TFE	TRH	Carry	Value	
3/1994	-12.7%	-12.8%	-12.7%	-12.6%	-12.4%	-12.2%	-11.8%	-11.3%	-10.7%	-10.0%	-9.0%	-7.7%	-5.9%	-1.2%	-4.3%	-10.4%	1.5%	-5.0%	-2.0%	-4.9%
9/1998	-0.2%	0.1%	0.5%	0.9%	1.3%	1.6%	2.0%	2.4%	2.8%	3.2%	3.6%	4.0%	4.4%	4.8%	5.8%	3.0%	5.0%	-1.9%	1.3%	-11.6%
9/1999	1.6%	1.6%	1.5%	1.5%	1.4%	1.4%	1.3%	1.2%	1.1%	1.0%	0.8%	0.6%	0.4%	-0.3%	-1.3%	-0.2%	0.3%	4.1%	1.2%	-7.8%
6/2000	5.7%	5.5%	5.2%	4.9%	4.5%	4.1%	3.6%	3.1%	2.6%	2.0%	1.3%	0.4%	-0.7%	-3.1%	-0.6%	0.3%	-2.0%	5.9%	6.8%	-4.5%
12/2000	6.9%	7.1%	7.3%	7.4%	7.4%	7.5%	7.4%	7.3%	7.2%	6.9%	6.5%	6.0%	5.2%	2.8%	4.9%	6.7%	1.7%	2.6%	2.4%	-9.6%
3/2001	4.9%	5.3%	5.8%	6.2%	6.6%	6.9%	7.2%	7.5%	7.7%	7.9%	7.9%	7.9%	7.7%	6.3%	0.0%	7.8%	7.1%	7.1%	0.7%	-13.3%
9/2001	1.8%	2.6%	3.4%	4.1%	4.9%	5.6%	6.3%	7.0%	7.7%	8.3%	8.9%	9.5%	10.0%	10.1%	5.2%	6.9%	13.5%	5.6%	6.6%	-15.8%
6/2002	7.3%	7.8%	8.1%	8.5%	8.8%	9.0%	9.2%	9.3%	9.3%	9.3%	9.1%	8.8%	8.1%	5.8%	5.0%	5.9%	6.3%	10.2%	3.7%	-14.1%
9/2002	-0.6%	0.8%	2.1%	3.4%	4.8%	6.1%	7.5%	8.9%	10.2%	11.6%	13.0%	14.4%	15.8%	17.3%	10.5%	9.0%	17.1%	4.9%	-5.4%	-18.0%
3/2003	4.9%	4.7%	4.5%	4.3%	4.0%	3.7%	3.4%	3.1%	2.6%	2.2%	1.6%	0.9%	0.1%	-1.9%	1.0%	4.3%	-0.8%	-1.9%	8.0%	-3.9%
12/2007	7.7%	7.7%	7.7%	7.7%	7.6%	7.5%	7.3%	7.0%	6.7%	6.2%	5.7%	5.0%	3.9%	1.1%	3.9%	4.2%	-0.4%	6.2%	1.5%	-5.1%
3/2008	9.5%	9.7%	9.9%	10.1%	10.2%	10.2%	10.1%	10.0%	9.8%	9.4%	8.9%	8.2%	7.1%	3.7%	2.8%	11.8%	2.0%	1.9%	0.9%	-10.7%
6/2008	8.0%	7.7%	7.3%	6.9%	6.4%	5.9%	5.3%	4.6%	3.9%	3.1%	2.1%	0.9%	-0.6%	-3.9%	-2.6%	1.4%	-1.9%	8.8%	9.5%	-3.9%
9/2008	-9.2%	-8.7%	-8.2%	-7.6%	-6.9%	-6.2%	-5.4%	-4.5%	-3.5%	-2.4%	-1.2%	0.2%	2.1%	6.1%	1.9%	1.4%	5.6%	-11.6%	-6.8%	-9.5%
12/2008	3.4%	5.4%	7.4%	9.4%	11.4%	13.3%	15.2%	17.1%	18.9%	20.7%	22.5%	24.2%	25.7%	26.4%	16.3%	22.1%	24.2%	0.6%	-8.4%	-23.0%
3/2009	-3.7%	-3.5%	-3.2%	-2.8%	-2.5%	-2.1%	-1.7%	-1.3%	-0.8%	-0.3%	0.3%	0.9%	1.7%	3.4%	-4.8%	0.1%	7.9%	3.1%	4.3%	-11.8%
6/2010	1.7%	2.0%	2.3%	2.6%	2.8%	3.1%	3.4%	3.6%	3.8%	4.0%	4.2%	4.3%	4.4%	4.1%	11.0%	0.2%	5.3%	1.4%	7.4%	-11.9%
9/2011	10.4%	11.4%	12.4%	13.3%	14.1%	14.9%	15.5%	16.1%	16.6%	17.0%	17.2%	17.2%	16.7%	13.8%	22.3%	9.6%	8.1%	8.6%	-5.2%	-14.4%
9/2015	2.8%	2.9%	3.0%	3.0%	3.0%	3.0%	3.0%	3.0%	2.9%	2.8%	2.7%	2.5%	2.1%	1.2%	4.6%	2.2%	1.8%	0.0%	5.2%	-7.0%
12/2018	-3.6%	-3.4%	-3.2%	-3.0%	-2.7%	-2.4%	-2.1%	-1.7%	-1.3%	-0.9%	-0.4%	0.2%	1.0%	2.5%	3.2%	1.7%	4.2%	-8.1%	4.2%	-14.5%
Average	2.3%	2.7%	3.1%	3.4%	3.7%	4.0%	4.3%	4.6%	4.9%	5.1%	5.3%	5.4%	5.5%	4.9%	4.2%	4.4%	5.3%	2.1%	1.8%	
Minimum	-12.7%	-12.8%	-12.7%	-12.6%	-12.4%	-12.2%	-11.8%	-11.3%	-10.7%	-10.0%	-9.0%	-7.7%	-5.9%	-3.9%	-4.8%	-10.4%	-2.0%	-11.6%	-8.4%	
Maximum	10.4%	11.4%	12.4%	13.3%	14.1%	14.9%	15.5%	17.1%	18.9%	20.7%	22.5%	24.2%	25.7%	26.4%	22.3%	22.1%	24.2%	10.2%	9.5%	

Source: PIMCO and Bloomberg as of December 2018. Each row represents a quarter in which the S&P 500 excess return was lower than -3.75%. Each column under "Beta target" shows portfolio returns using optimal weights with the specified beta target (Exhibit 5). Each strategy is scaled to have 5% unconditional volatility. The columns under "Individual" show the historical return for each individual strategy, rescaled to 10% unconditional volatility for comparison. LTE stands for long Treasuries, TFE for trend-following and TRH for tail risk hedging. Due to data availability, results for first-quarter 1994 are based only on March 1994 data.

Finally, to better understand how much variability one might expect in a portfolio of defensive assets, we calculate the returns of the optimal portfolios in Exhibit 5 for the 20 quarters in which the excess return on the S&P 500 was less than -3.75%. Exhibit 9 shows these results.[15] Indeed, the average conditional return increases nearly monotonically with the conditional beta target and the minimum and maximum returns go from largest to smallest moving left to right. These same measures can help to understand the risk/return properties of individual defensive strategies. For example, tail risk hedging has a much more consistent payoff in negative equity markets than, say, carry or value. This is precisely why the most negative conditional beta portfolios will have relatively large allocations to tail-risk-hedging strategies. Though by no means perfect, these results indicate that constraining the defensive portfolio to more negative equity beta levels materially increases the payoff in stressed markets.

15 Exhibit 9 should not be considered a backtest of the defensive strategy because the historical performance is based on the current optimal portfolio weights, not what would have been held in the past. Nonetheless, these values are useful in understanding when, and to what extent, defensive strategies have been effective historically.

4. EXTENSIONS

In practice, we could further refine the asset selection for each strategy to potentially improve its performance. In this section, we show two such refinements – one for the fixed income (long Treasuries) part of the strategy and one for the trend-following portion.

In the first example, we replace the passive long Treasuries index with an active swap position. Specifically, we leverage a constant-maturity five-year swap to achieve the same duration as the long Treasury index. Historically, the swap strategy has delivered a much higher Sharpe ratio than the index return while preserving the beneficial downside equity property of passive Treasuries.[16] Over our sample, the unconditional Sharpe ratio of the swap strategy is 0.8 while that of the long Treasury index is 0.4. These two series also have high correlation – around 0.80 – both conditionally and unconditionally (Davis and Fuenzalida 2019). We repeat our optimization problem replacing long Treasuries with the active swap strategy and haircut the expected unconditional Sharpe ratio to 0.14 so that it is double that of long Treasuries. The swap strategy has a conditional volatility of 6.5% and a conditional equity beta of -0.53 (or a conditional equity correlation of -0.84).

Exhibit 10 shows the portfolio allocations. Because of the better return property of the active swap strategy, the optimization now allocates more to fixed-income and the unconditional Sharpe ratio-conditional correlation frontier shifts upward.

Exhibit 10: Optimal portfolio and Sharpe ratio-correlation frontier
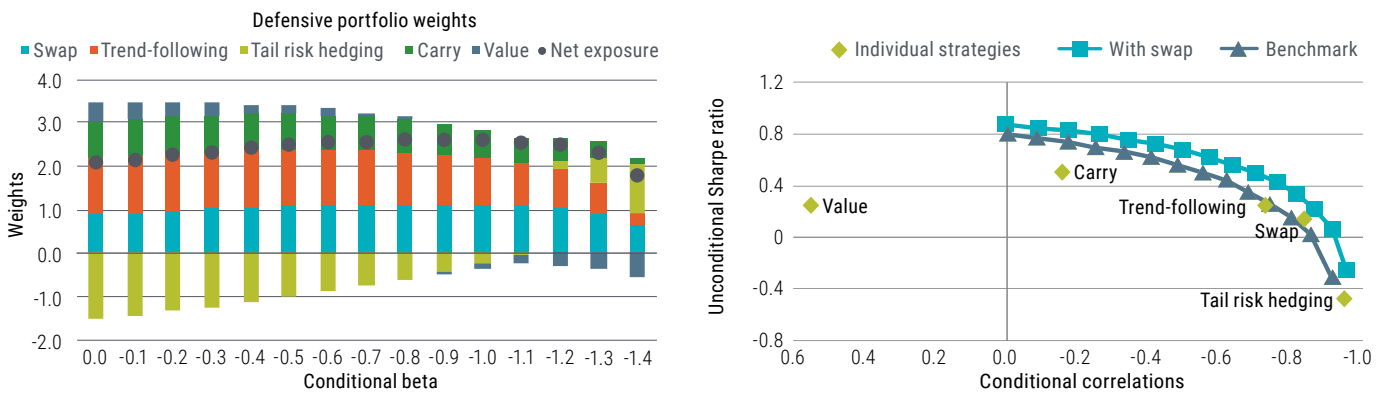
Source: PIMCO and Bloomberg as of December 2018

In the next example, we refine the trend-following strategy by filtering it to exclude positions that have a correlation with the S&P 500 higher than 0.2. That is, we do not hold long (short) assets that are too positively (negatively) correlated with equity. Perhaps unsurprisingly, imposing this constraint decreases the historical unconditional Sharpe ratio for the trend-following strategy from 0.66 to 0.47. However, it greatly improves the strategy's performance during equity drawdowns; the filtered trend-following strategy has a conditional correlation with the S&P 500 of -0.83, compared with -0.70 for the benchmark version.

Exhibit 11: Optimal portfolio and Sharpe ratio-correlation frontier

Source: PIMCO and Bloomberg as of December 2018. **Exhibits are hypothetical examples for illustrative purposes only.**

16 Historically, leveraging the five-year swap rate to match the duration of long Treasuries has delivered the highest Sharpe ratio of a 60/40 portfolio most of the time.

Exhibit 11 displays the portfolio allocation, as well as the Sharpe ratio-correlation frontier. We set the unconditional Sharpe ratio for the filtered trend-following at 0.15 to reflect the fact that, historically, its unconditional Sharpe is lower than its unfiltered counterpart. However, as we can see in Exhibit 11, even with the lower unconditional return assumption, the filtered trend-following still offers good balance between return and defensiveness. As a result, we observe material allocations to the defensive trend-following strategy across most of the conditional beta spectrum. Furthermore, because the filtered trend-following strategy has a much better defensive property, the Sharpe ratio-correlation frontier shifts outward relative to the benchmark case.

5. CONCLUDING REMARKS

We first examined the general diversification and hedging properties for various well-accepted equity defensive assets, noting a critical trade-off between a strategy's defensiveness and its return potential. This is a key point that investors must understand: It will only be possible to obtain high and reliable negative conditional betas if the investor is willing to give up some overall return for the increased certainty. We then proposed a theoretical framework based on this risk/return trade-off for optimal defensive portfolio construction. The model was evaluated using data on five common strategies. The empirical results highlighted the benefit of using a portfolio approach rather than simply relying on individual strategies. Finally, we showed that refining certain strategies may help to achieve better Sharpe ratios while at the same time maintaining, or even enhancing, those strategies' downside properties.

REFERENCES

Adrian, Tobias, Richard K. Crump, and Emanuel Moench. "Pricing the Term Structure with Linear Regressions." *Journal of Financial Economics*, October 2013

Baz, Jamil, Josh Davis, and Graham A. Rennison. "Hedging for Profit: A Novel Approach to Diversification." PIMCO Quantitative Research and Analytics, 2017

Baz, Jamil, Steve Sapra, and German Ramirez. "Stock, Bonds and Causality." *Journal of Portfolio Management* (forthcoming)

Campbell, John Y., Adi Sunderam, and Luis M. Viceira. "Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds." *Critical Finance Review*, September 2017

Davis, Josh and Cristian Fuenzalida. Forthcoming article.

Fung, William and David A. Hsieh. "The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers." *Review of Financial Studies*, Summer 2001

Markowitz, Harry M. "Portfolio Selection." *Journal of Finance*, March 1952

Wright, Jonathan H. "Term Premia and Inflation Uncertainty: Empirical Evidence from an International Panel Dataset." *American Economic Review*, June 2011

TECHNICAL APPENDIX

Appendix 1: Backtests for trading strategies

Long Treasuries:
PnL is calculated daily using the return on the Bloomberg Barclays US Long Treasury Total Return Index (Unhedged).

Tail risk hedging:
PnL is calculated as follows: The underlying portfolio, V_t, earns cash rate (daily compounding). Each month on option expiration day, use the fund in the underlying portfolio to buy $V_t/(12*S_t)$ contracts of 10% out-of-the-money put options with a one-year expiration (S_t denotes the S&P 500 index value on option expiration date t). At each point in time, the option portfolio consists of 12 put options, with expiration times ranging from one to 12 months. The PnL is calculated daily using the total portfolio, which includes the underlying portfolio and the options portfolio.

Carry:
Target a 1% scale based on the 260-day volatility for each of the three asset classes. For currency, at the beginning of each month, rank G10 currencies against USD (AUD, CAD, CHF, EUR, GBP, JPY, NOK, NZD and SEK) based on their annualized real carry. Long the top-third-highest-yielding currencies, and short the bottom-third-lowest-yielding currencies.

For rates, rank G6 10-year swaps (AUD 10Y, CAD 10Y, EUR 10Y, GBP 10Y, JPY 10Y and USD 10Y) by carry plus roll-down per year of duration. Long the top-third-highest carry and short the bottom-third-lowest carry plus roll-down per year of duration.

For commodity, calculate carry using commodity forwards for 21 commodities (aluminum, Brent, cocoa, coffee, copper, corn, cotton, crude, gasoline, gold, heating oil, lead, low sulphur gasoil, natural gas, nickel, platinum, silver, soybeans, sugar, wheat and zinc). At the beginning of each month, rank the commodities based on their annualized nominal carry. Long those with the top-third-highest carry, and short those with the bottom-third-lowest carry.

Value:
Target a 1% scale based on the 260-day volatility for each of the three asset classes. For currency, at the beginning of each month, rank the G10 currencies against USD (AUD, CAD, CHF, EUR, GBP, JPY, NOK, NZD and SEK) based on their purchasing power parity. Long the lowest-third currencies, and short the highest-third currencies.

For rates, rank G6 10-year swaps (AUD 10Y, CAD 10Y, EUR 10Y, GBP 10Y, JPY 10Y and USD 10Y) by their real yields, then long the highest third and short the lowest third.

For commodity, divide the assets into four sectors: livestock (live cattle and lean hogs), grain (corn, soybeans and wheat), soft (cocoa , coffee, cotton and sugar) and petroleum (Brent, crude, gasoline, heating oil and low sulphur gasoil). At each point in time for each commodity, compute the five-year rolling risk-adjusted return of the front rolling contract. Within each sector, long the third with the lowest risk-adjusted return and short the third with the highest risk-adjusted return.

Trend-following:
Target a 1% scale based on the 260-day volatility for each of the four asset classes: equity (Euro Stoxx 50, S&P 500, Nikkei 225, FTSE 100, ASX 200 and Swiss Market Index); commodity (aluminum, Brent, cocoa, coffee, copper, corn, cotton, crude, gasoline, gold, heating oil, lead, low sulphur gasoil, natural gas, nickel, platinum, silver, soybeans, sugar, wheat and zinc); currency (AUD, CAD, CHF, EUR, GBP, JPY, NOK, NZD and SEK against USD); and rates (10-year swaps for AUD, CAD, EUR, GBP, JPY and USD). For each asset at each point in time, calculate

the difference between the last value of the excess return index relative to its 250-day moving average. Assuming the signal has a normal distribution, apply the normal CDF and convert the signal to a strength measure between 0 and 1. Apply a response function to map the measure of the strength to the position, the maximum of which is scaled to have 1% volatility. The response function used gradually enters a trend; at the very extremes, the size of the position declines to "take profits."

Historical unconditional excess return moments of various strategies (see Exhibit 12):

Exhibit 12: Historical unconditional return moments

	Historical unconditional moments		
	Expected excess return	Volatility	Sharpe ratio
S&P 500	5.2%	15.6%	0.33
Long Treasuries	2.0%	5.0%	0.40
Tail risk hedging	-2.4%	5.4%	-0.45
Trend-following	3.3%	5.0%	0.66
Carry	5.1%	5.0%	1.01
Value	3.1%	5.0%	0.63

Source: PIMCO and Bloomberg as of date December 2018

Appendix 2: Model solution

Appendix 2.1: MVO, unit-beta and zero-beta MVO portfolios

Unconstrained mean-variance optimization (MVO): First, we consider a standard Markowitz MVO portfolio. The optimization is given by

$$max_w \, E[R_P] = w'\mu$$
$$0.5 \, w'\Sigma w = 0.5\sigma_p^2 \, .$$

The Lagrangian is given by

$$\mathcal{L} = w'\mu - 0.5\gamma_{\text{MVO}}\big(\, w'\Sigma w - \sigma_p^2\big).$$

The first-order condition with respect to w is $w^{MVO} = \gamma_{\text{MVO}}^{-1}\Sigma^{-1}\mu$. Substituting this into the volatility constraint yields $\gamma_{\text{MVO}} = \sigma_p^{-1}(\mu'\Sigma^{-1}\mu)^{1/2}$. Hence:

$$w^{MVO} = \gamma_{\text{MVO}}^{-1}\Sigma^{-1}\mu = \sigma_p \frac{\Sigma^{-1}\mu}{(\mu'\Sigma^{-1}\mu)^{1/2}}. \qquad (A.1)$$

Note that the expected return of the MVO is $\mu'w^{MVO} = \sigma_p(\mu'\Sigma^{-1}\mu)^{1/2}$ and the Sharpe ratio of the MVO portfolio (or, in fact, any portfolio that is proportional to the MVO portfolio, $c \cdot w^{MVO}$), is given by

$$SR_{MVO} = (\mu'\Sigma^{-1}\mu)^{1/2}. \qquad (A.2)$$

Unit-beta portfolio: Next we consider an optimization problem that minimizes variance subject to a unit-conditional-beta target:

$$min_w \ \frac{1}{2} \ w'\Sigma w$$

$$w'\beta_c = 1.$$

The solution to this problem is given by

$$w^B = \frac{\Sigma^{-1}\beta_c}{\beta_c'\Sigma^{-1}\beta_c}.$$

Note that this portfolio has unconditional expected excess return:

$$\mu^B = \mu'w^B = \frac{\mu'\Sigma^{-1}\beta_c}{\beta_c'\Sigma^{-1}\beta_c},$$

unconditional variance:

$$\sigma_{w^B}^2 = (\beta_c'\Sigma^{-1}\beta_c)^{-1}$$

and an unconditional Sharpe ratio:

$$SR_{w^B} = \frac{\mu'\Sigma^{-1}\beta_c}{(\beta_c'\Sigma^{-1}\beta_c)^{1/2}}. \qquad (A.3)$$

It also follows that the minimum-variance portfolio for any target $\hat{\beta}$ is $\hat{\beta}w^B$, which has an unconditional expected excess return $\hat{\beta}\mu^B = \frac{\mu'\Sigma^{-1}\beta_c}{\beta_c'\Sigma^{-1}\beta_c}\hat{\beta}$, unconditional variance $\hat{\beta}^2\sigma_{w^B}^2 = \hat{\beta}^2(\beta_c'\Sigma^{-1}\beta_c)^{-1}$ and unconditional Sharpe ratio $\frac{\mu'\Sigma^{-1}\beta_c}{(\beta_c'\Sigma^{-1}\beta_c)^{1/2}} = SR_{w^B}$.

Zero-beta MVO portfolio: Last, we consider an optimization problem subject to a zero-conditional-beta target:

$$max_w \ w'\mu - \frac{\gamma_{MVO}}{2} \ w'\Sigma w$$

$$w'\beta_c = 0,$$

with $\gamma_{MVO} = \sigma_p^{-1}(\mu'\Sigma^{-1}\mu)^{1/2}$, the multiplier in the MVO problem. The Lagrangian is given by

$$\mathcal{L} = w'\mu - \frac{\gamma_{MVO}}{2} \ w'\Sigma w - \lambda_{ZB} \ w'\beta_c.$$

The first-order condition with respect to w is

$w = \gamma_{MVO}^{-1}\Sigma^{-1}\mu - \lambda_{ZB}\gamma_{MVO}^{-1}\Sigma^{-1}\beta_c$. Substituting this into the beta constraint $w'\beta_c = 0$ yields $\lambda_{ZB} = \frac{\mu'\Sigma^{-1}\beta_c}{\beta_c'\Sigma^{-1}\beta_c}$.

Therefore, this zero-beta portfolio is given by

$$w^{ZB} = \gamma_{MVO}^{-1}\Sigma^{-1}\mu - \frac{\Sigma^{-1}\beta_c}{\beta_c'\Sigma^{-1}\beta_c}\gamma_{MVO}^{-1}\mu'\Sigma^{-1}\beta_c.$$

Note that $\gamma_{MVO}^{-1}\mu'\Sigma^{-1}\beta_c = w^{MVO}\beta_c \equiv \beta^{MVO}$, beta of the unconstrained MVO portfolio. We can thus rewrite the zero-beta portfolio as

$$w^{ZB} = w^{MVO} - \beta^{MVO}w^B.$$

Appendix 2.2: Solution to the beta-constrained optimization problem

The Lagrangian for the optimization problem (2) − (2.2) is given by

$$\mathcal{L} = w'\mu - \gamma(0.5 \ w'\Sigma w - 0.5\sigma_p^2) - \lambda(\ w'\beta_c - \bar{\beta}_c). \quad (A.4)$$

The optimal conditions are given by

$$\mu - \gamma\Sigma w - \lambda\beta_c = 0 \qquad (A.5)$$

$$w'\Sigma w - \sigma_p^2 = 0 \qquad (A.6)$$

$$w'\beta_c - \bar{\beta}_c \leq 0, \ \lambda \geq 0, \ \lambda(\ w'\beta_c - \bar{\beta}_c) = 0. \qquad (A.7)$$

If the beta constraint does not bind, then $\lambda = 0$ and (2) − (2.2) reduces to a standard unconstrained MVO problem, as shown in Appendix 2.1. On the other hand, if the beta constraint binds, then $\lambda > 0$ and (A.5) can be rewritten as

$$w = \gamma^{-1}\Sigma^{-1}\mu - \gamma^{-1}\lambda\Sigma^{-1}\beta_c. \qquad (A.8)$$

Substituting this into the binding beta constraint $w'\beta_c - \bar{\beta}_c = 0$ in (A.7), we can solve for λ as a function of γ:

$$\lambda = \frac{\gamma^{-1}\mu'\Sigma^{-1}\beta_c - \bar{\beta}_c}{\gamma^{-1}\beta_c'\Sigma^{-1}\beta_c}. \qquad (A.9)$$

We can then rewrite the optimal weight in terms of γ only:

$$w = \gamma^{-1}\Sigma^{-1}\mu - (\ \gamma^{-1}\mu'\Sigma^{-1}\beta_c - \bar{\beta}_c)\frac{\Sigma^{-1}\beta_c}{\beta_c'\Sigma^{-1}\beta_c}. \quad (A.10)$$

Because γ is a constant (the Lagrange multiplier), the first component of (A.10), $\gamma^{-1}\Sigma^{-1}\mu$, is simply a set of weights that is proportional to that of the unconstrained MVO (A.1) and can be expressed as cw^{MVO}. The first term in the parentheses can thus be expressed as $\gamma^{-1}\mu'\Sigma^{-1}\beta = \beta'(cw^{MVO}) = c\beta^{MVO}$, where $\beta^{MVO} = \beta'w^{MVO}$ is the beta of the unconstrained MVO portfolio and constant $c = \gamma_{MVO}/\gamma$. Finally, $\frac{\Sigma^{-1}\beta_c}{\beta_c'\Sigma^{-1}\beta_c}$ is the unit-beta portfolio w^B. Putting everything together, (A.10) becomes

$$w = \bar{\beta}_c w^B + c(w^{MVO} - \beta^{MVO}w^B). \qquad (A.11)$$

Substituting (A.10) into the variance constraint (A.6) to solve for γ:

$$\gamma = \left(\frac{\mu'\Sigma^{-1}\mu - \frac{(\mu'\Sigma^{-1}\beta_c)^2}{\beta_c'\Sigma^{-1}\beta_c}}{\sigma_p^2 - \bar{\beta}_c^{\,2}(\beta_c'\Sigma^{-1}\beta_c)^{-1}} \right)^{1/2}. \qquad (A.12)$$

Appendix 2.1 shows that $\mu'\Sigma^{-1}\mu = SR^2_{MVO}$ and $\frac{(\mu'\Sigma^{-1}\beta_c)^2}{\beta_c'\Sigma^{-1}\beta_c} = SR^2_{w^B}$. Therefore, the numerator is the difference between the squared Sharpe ratio of the MVO portfolio and that of the beta portfolio. Because the unconstrained MVO achieves the highest possible Sharpe ratio, the numerator is non-negative and equals zero when $\mu = constant * \beta_c$ (e.g., if CAPM holds and conditional beta equals unconditional beta). Appendix 2.1 also shows that $\bar{\beta}_c^{\,2}(\beta_c'\Sigma^{-1}\beta_c)^{-1}$ is the minimum variance for a portfolio with beta $\bar{\beta}_c$. Therefore, for a solution $\gamma > 0$ to exist, we need the expected excess returns not to be proportional to the conditional betas, and the beta target has to be achievable so $\sigma_p^2 > \bar{\beta}_c^{\,2}(\beta_c'\Sigma^{-1}\beta_c)^{-1}$. Using (A.9) together with $\gamma^{-1}\mu'\Sigma^{-1}\beta_c = c\beta^{MVO}$, we see that the condition for $\lambda > 0$ to exist, when $\gamma > 0$ exists, is that $c\beta^{MVO} - \bar{\beta}_c > 0$.

Substituting (A.12) into (A.10), we obtain the optimal weights, it then follows that

$$\mu'w = \gamma^{-1}\mu'\Sigma^{-1}\mu - \left(\gamma^{-1}\mu'\Sigma^{-1}\beta_c - \bar{\beta}_c \right) \frac{\mu'\Sigma^{-1}\beta_c}{\beta_c'\Sigma^{-1}\beta_c} \qquad (A.13)$$

we obtain the expected return of the portfolio:

$$\mu'w = \mu'\left(\frac{\Sigma^{-1}\beta_c}{\beta_c'\Sigma^{-1}\beta_c} \right)\bar{\beta}_c +$$
$$\left(\mu'\Sigma^{-1}\mu - \frac{(\mu'\Sigma^{-1}\beta_c)^2}{\beta_c'\Sigma^{-1}\beta_c} \right)^{1/2} \left(\sigma_p^2 - \frac{\bar{\beta}_c^2}{\beta_c'\Sigma^{-1}\beta_c} \right)^{1/2}.$$

Using notations introduced earlier,

$$\mu'w = \mu'\left(\bar{\beta}_c w^B\right) + \left(SR^2_{MVO} - SR^2_{w^B} \right)^{1/2}\left(\sigma_p^2 - \bar{\beta}_c^2 \sigma_{w^B}^2 \right)^{1/2}. \quad (A.14)$$

Appendix 2.3: CAPM

If CAPM holds and the conditional and unconditional betas are equal, then a portfolio with a negative beta $\bar{\beta}_c$ must have an expected return of $\mu_{equity}\bar{\beta}_c < 0$.

As we mentioned in Appendix 2.2, $\gamma > 0$ does not exist in this case; nonetheless, we can still use (A.14) to calculate expected returns. To see this, substitute $\mu = \mu_{equity}\beta_c$ into (A.2) and (A.3):

$$SR_{MVO} = (\mu'\Sigma^{-1}\mu)^{1/2} = \mu_{equity}(\beta_c'\Sigma^{-1}\beta_c)^{1/2}, \qquad (A.15)$$

and

$$SR_{w^B} = \frac{\mu_{equity}(\beta_c'\Sigma^{-1}\beta_c)}{(\beta_c'\Sigma^{-1}\beta_c)^{1/2}} = \mu_{equity}(\beta_c'\Sigma^{-1}\beta_c)^{1/2}. \quad (A.16)$$

That is, the Sharpe ratio of the beta portfolio equals that of the MVO. In addition,

$$\mu'w^B = \frac{\left(\mu_{equity}\beta_c\right)'\Sigma^{-1}\beta_c}{\beta_c'\Sigma^{-1}\beta_c} = \mu_{equity}. \qquad (A.17)$$

Plugging (A.15)–(A.17) into (A.14), we again obtain $\mu'w = \mu_{equity}\bar{\beta}_c$.

APPENDIX 3: AN EXAMPLE

Consider an example with two risky assets, X and Y. Assume the unconditional excess return, variance covariance matrix and conditional betas (or beta, in the case when CAPM holds) are given by

$$\mu = \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \sigma_X^2 & 0 \\ 0 & \sigma_Y^2 \end{bmatrix}, \quad \beta_c = \begin{bmatrix} -0.5 \\ 0 \end{bmatrix}.$$

First, we can see that the minimum-variance portfolio with a beta of 1 is one that shorts two units of X – that is, $w^B = [-2, 0]'$.

Now suppose our target beta is $\bar{\beta}_c = -1$, and assume the MVO portfolio has a beta greater than -1 so the constraint binds.

1. The first component of (A.14) is

$$\mu'(\bar{\beta}_c w^B) = \mu_X(-1)(-2) + \mu_Y(-1)(0) = 2\mu_X.$$

We can consider a few different cases:

CAPM holds: $\mu = \mu_{\text{equity}}\beta_c$, so $\mu_X = -0.5\mu_{\text{equity}}$ or $2\mu_X = -\mu_{\text{equity}}$ (as if we are only shorting equity).

CAPM fails and $\mu_X > 0$ (X is "unicorn"): The "insurance premium" is $2\mu_X > 0$.

CAPM fails and $0 > \mu_X > -0.5 * \mu_{\text{equity}}$: We still have to pay an insurance premium for the negative beta target, but not as much as if we are directly shorting equity.

2. For the Sharpe ratios:

$$\text{SR}_{\text{MVO}} = (\mu'\Sigma^{-1}\mu)^{1/2} = \left(\left(\frac{\mu_X}{\sigma_X}\right)^2 + \left(\frac{\mu_Y}{\sigma_Y}\right)^2\right)^{1/2}$$

and

$$\text{SR}_{w^B} = \frac{\mu'\Sigma^{-1}\beta_c}{(\beta_c'\Sigma^{-1}\beta_c)^{1/2}} = \frac{-0.5(\mu_X\sigma_X^{-2})}{0.5(\sigma_X^{-1})} = -\frac{\mu_X}{\sigma_X}.$$

Therefore:

$$\text{SR}_{\text{MVO}}^2 - \text{SR}_{w^B}^2 = \left(\frac{\mu_Y}{\sigma_Y}\right)^2.$$

CAPM holds: Then $\mu_Y = 0$ and $\text{SR}_{\text{MVO}}^2 - \text{SR}_{w^B}^2 = 0$. There is no diversification benefit, as we cannot increase the expected return without changing the beta of the portfolio.

CAPM fails and $\mu_Y > 0$: It is possible to construct a positive return zero-beta MVO portfolio, and we can add return by leveraging up this portfolio until we use up all of the volatility budget.

In this simple example, asset X is purely used to construct the beta portfolio for equity hedging, while asset Y is only used in the zero-beta MVO portfolio for diversification/return generation. In general, all assets could be used for both purposes.

This example, and more generally (A.14), show that there are two return sources when we deviate from the CAPM world. First, we pay a lower insurance premium to reach the beta target by using assets that have positive alphas. Second, because CAPM fails, we can construct a portfolio with zero-beta MVO and positive return to generate additional return using the remaining volatility budget.