Invesco

#01

1st issue 2020

Are you taking enough risk? How do you know?

Risk & Reward

Research and investment strategies

Asset management firms often stress the importance of risk management – and Invesco is no exception. But successful risk management is more than mitigating risk. Taking enough risk is equally important, especially when yields are as low as they are today.

My colleagues from our Multi-Asset team have developed a new risk measure, "Internal Portfolio Risk". They argue that it reflects portfolio risk much more accurately than traditional measures such as the standard deviation of past returns. This new indicator can help, we believe, construct portfolios with appropriate levels of risk: not too high but – equally important – not too low either.

Indeed, the low interest rate environment requires new ways of thinking. Sometimes, however, it helps to rethink traditional concepts as well. Convertibles have long been regarded as an interesting option for bond investors seeking extra yield. But with the convertibles market having shrunk by about half since its peak in 2008, there is a dearth of opportunities for yield-hungry investors. We have researched a flexible synthetic convertibles strategy that may respond to the current environment and investor needs.

We've also included an interview with my very experienced colleague Andrew Lo about investor attitudes towards mainland China, as well as covering four new studies on factor investing: we analyze a factor-based buy-and-hold strategy for bonds, discuss the merits of hierarchical clustering techniques for multi-asset multi-factor investing, investigate ways of integrating the low volatility factor in core equity portfolios and outline methods to combine forecasts in order to gain optimal factor exposures.

We hope you enjoy this latest issue of Risk & Reward.

Best regards,

Marty Flanagan
President and CEO of Invesco Ltd.

Contents

In focus

Are you taking enough risk?
How do you know?

By Michael Marshall

In brief

In a world of low volatility and low interest rates, it becomes increasingly important to understand the level of risk in an investment portfolio. Standard measures such as statistical volatility are often misleading. We propose using a different metric, Internal Portfolio Risk, which we think provides a far more accurate picture of the level of an investment portfolio's true risk.

A low volatility and low interest rate world has prompted a reassessment of the appropriate level of risk required to achieve a specified return outcome. As investors seek returns in this anaemic world, absolute return funds come under the spotlight because they are typically structured to deliver an attractive return with lower than equity-level risk. This dual focus on return and risk can also bring scrutiny, as many critics are concerned, about the promise of a "free lunch" – can one achieve a long-term, risk asset-like return and deliver it with a low standard deviation. In other words, are you taking enough risk to achieve a return target and, importantly, how do you know?

We believe that high information ratios – high returns for a given level of ex-post volatility – are possible if a manager can achieve a positive and persistent hit rate and a positive return skew, provided the risks interact in such a manner as to achieve low portfolio volatility.

Standard deviation, VaR, tracking error and other risk metrics have become synonyms, or even direct substitutes, for risk. But, fundamentally, they are not the same thing.

A volatility target put in place to limit downside exposure does not equate to the risk required to achieve a desired level of return. Low day-to-day volatility achieved through high levels of diversification allows portfolio managers to manage and limit short-term downside risk (drawdowns). However, it is the amount of internal risk – the risk associated with the individual positions or ideas – that allows the portfolio manager to achieve return targets.

Put another way, portfolio returns are not an outcome of the level of assumed volatility but rather of the skill of the manager in selecting positions and combining these in portfolios – volatility is merely an outcome of how the assumed risks behave.

Targets and targets
Any asset or portfolio generates returns in two ways: capital returns and yield (or carry).

Yields can be thought of as premia paid to investors bearing the risk of a position, e.g. dividends on stocks, bond coupons and the carry in currency, to name a few. If a market were to stand still, the yield would equal the total return.

The portfolio yield can be estimated ex-ante with some degree of certainty – for example, an investor can be confident of the yield to be received from a sovereign bond and, at an index level, we have some degree of certainty of the dividend yield.

Capital returns, on the other hand, require forward estimation: How far can the S&P 500 rally or the euro fall? Where will five-year US inflation be trading in a year's time? By how much will share A outperform share B?

How these capital returns play out, both idiosyncratically and combined, will, ex post, determine a portfolio's volatility. In other words: volatility is only concerned with the capital movements of positions.

Of course, nothing is certain, and some yields are more guaranteed than others – there is clearly a feedback loop between yield and capital risk. However, that risk is incorporated into thinking about what size position is appropriate for, say, a high yield credit market or low risk sovereign bond. Fortunately, in a world of derivatives and leverage, it is relatively easy to scale the volatility contribution of a particular asset, enabling flexibility in the asset's contribution to the portfolio of (post funding) yield and expected return.

Using this, we can define the capital target as the target return less expected portfolio yield. This gives us the amount of return the portfolio must generate from its capital moves – or the amount of return the capital risks must generate (figure 1).

It follows that, when trying to achieve a return target, the closer the portfolio yield is to the required return, the less capital risk is required to achieve or exceed that target.


Figure 1
The higher the yield, the less capital return is necessary
Effect of yield on return target

■ Yield   ■ Capital target
% of return target

Source: Invesco. For illustrative purposes only.

Figure 2
Adding risk widens the potential return distribution
Gaussian return distributions for different levels of yield and volatility

— 3.5% yield / 5% σ   — 1% yield / 10% σ   — Return target

Miss target   Exceed target

-40%   -30%   -20%   -10%   0%   10%   20%   30%   40%
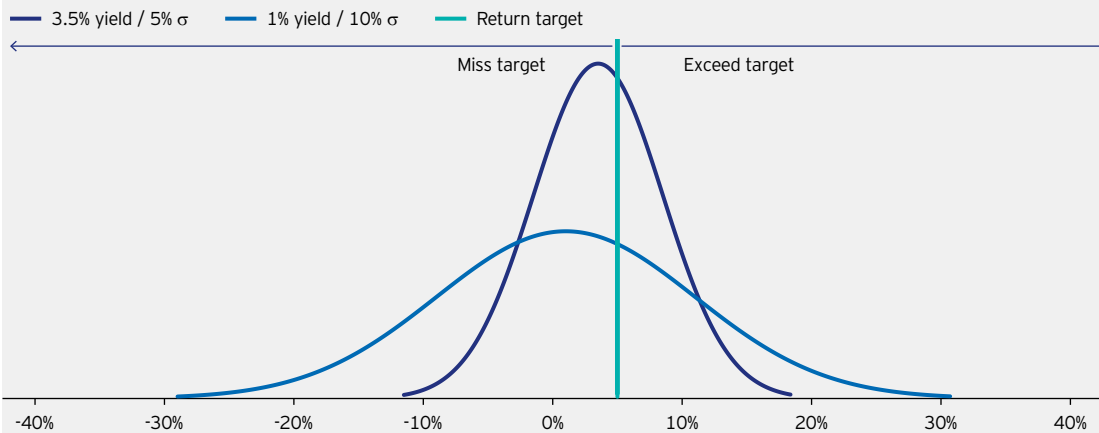
Source: Invesco. For illustrative purposes only.

To illustrate this in absolute terms: if risk-free cash rates are equivalent to the return target, then there is no need to take capital risk to achieve it. Adding risk will only increase the possible distribution of returns around the target and in doing so, increase the potential overshoot in both directions (figure 2).

Normally, however, the target return will exceed the expected portfolio yield, and some capital risk will have to be taken to achieve the target. Knowing this, we are still left with the problem of determining just how much risk is necessary to give us a good chance of generating capital returns at least equal to our capital target while minimising the chance that capital losses cause a target miss.

Focusing solely on standard deviation can be misleading.

Volatility ≠ Risk
Focusing solely on standard deviation can be misleading. The observable dynamic of a portfolio is driven by the level of volatility of each underlying investment combined with the correlation between those investments. However, these relationships do not remain static through time. Even using a simple asset allocation example of a stock/bond portfolio, the correlation between these two assets is not stable. This becomes increasingly true when you expand to more complex strategies, like those that look to isolate and reflect independent risks across idiosyncratic investment ideas.

Figure 3 shows that the correlation between equities and bonds has been both positive and negative at different times. In figure 4, we can also see that the standard deviation of both asset classes has evolved over time and, at times, has been at very similar levels.

This could lead to sub-optimal investment choices based on how much observable volatility is embedded

Figure 3
Correlations: sometimes positive, sometimes negative

— Rolling 1-year correlation between US Treasuries and global equities

%

80
60
40
20
0
-20
-40
-60
-80
2/91   2/94   2/97   2/00   2/03   2/06   2/09   2/12   2/15   2/18

Source: Bloomberg. Data as at 6 September 2019. For illustrative purposes only. Treasuries represented by 10-year US Treasuries. Global equities represented by MSCI World Index.

Figure 4
The standard deviation of returns is far from constant

— Global equities — US Treasuries
Annualized standard deviation

Source: Bloomberg. Data as at 6 September 2019. For illustrative purposes only.
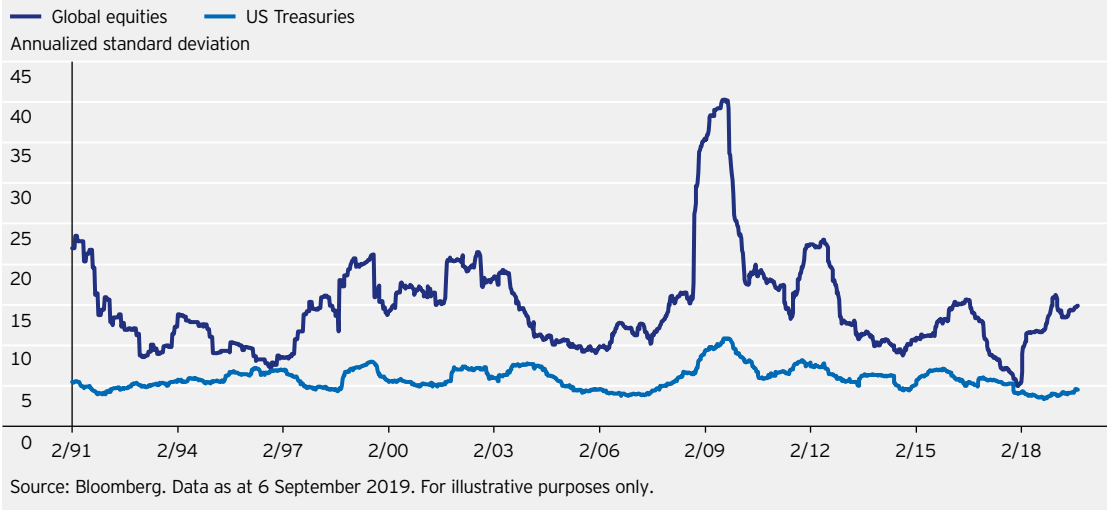
in the portfolio relative to the return outcome for each asset class.

Figure 5 highlights the lack of relationship between backward-looking observable risk metrics, such as ex-ante standard deviation, and subsequent returns. Comparing December 1995 with March 2009, both delivered subsequent annualized returns of around 12% over the next few years but, in 1995, the observable standard deviation was 5.5% and in March 2009 it was 9.3%.

We therefore need to look at a broader set of metrics to determine whether the underlying risk embedded within the portfolio is sufficient to achieve the portfolio return target.

The 'required information ratio'
In the textbook Gaussian world, where returns are normally distributed (clearly not the case in the real world), one can imply the probability of achieving a

given capital return by dividing the capital target by the expected portfolio standard deviation (the 'required information ratio') and inferring the outcome from the Gaussian distribution (figure 6). We use the after-yield or capital target for this because, as we have mentioned before, volatility is largely concerned with the capital moves of the assets – not the yield.

For example, if the required information ratio is 1, we can estimate that this portfolio has a 16% chance of reaching or exceeding its objectives, as 16% of the normal distribution lies to the right of one standard deviation. These are not great odds!

Likewise, if the expected yield equals the targeted return, the capital target and required information ratio is 0, so, regardless of the level of volatility (under this rationale), there is a 50% chance of exceeding the target, because 50% of the distribution lies to the right of zero.

Figure 5
Standard deviation is a poor predictor of returns

Subsequent 180-week annualized return
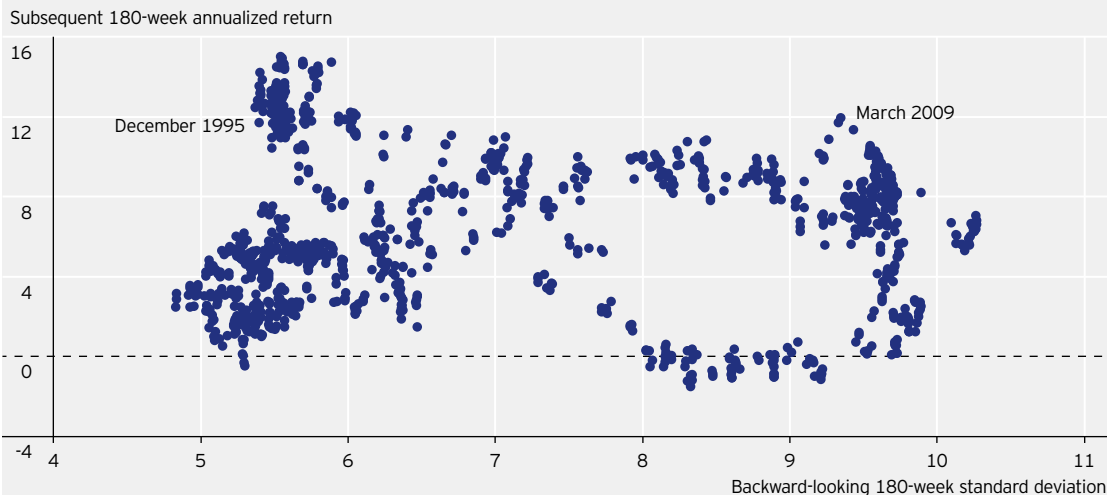
Source: Bloomberg, data as at 6 September 2019. The data points represent the standard deviation and returns of a 50-50 portfolio of the MSCI World and the S&P US Treasury bond 7-10 year total return index.

Figure 6
Risk and return in a Gaussian world
Probability of positive return outcome (normal distribution)

		50%
		16%
		2.5%

0σ　　　　　1σ　　　　　2σ　　　　　3σ

Source: Invesco. For illustrative purposes only.

Although this may be an easy heuristic to apply, it is extremely flawed in practice. Standard deviation is not necessarily a good estimator of risk as financial asset returns are not normally distributed and correlations and volatilities can evolve, as previously discussed. This thinking also completely ignores manager skill and, most importantly, it penalises diversification.

Through mathematical construct, more diversification equals less volatility and a higher required information ratio (ex-ante) for a given target. This implies that more manager skill would be required to achieve the desired return outcome.

Take the example of an illustrative multi-asset portfolio targeting a return of 3-month GBP LIBOR plus 5%, with a current ex-ante standard deviation of 3.5% and 2.5% expected yield. This implies a required information ratio of around 1. It seems fairly optimistic to expect the managers to be able to achieve that kind of ratio consistently over time but, as we have shown, this thinking is flawed.

When portfolios exhibit high levels of diversification – where the returns are driven by many truly independent factors – standard deviation can be extremely misleading and could prompt managers to understate the probability of achieving their targets, leading them to take on excessive levels of risk.

In other words, low volatility does not necessarily mean low risk.

Volatility – diversification – risk
Volatility
For any portfolio, there are many ways to calculate its variance, but the simplest is using parametric ex-post variance determined as:

$$\sigma_p^2 = W^T \cdot v \cdot W$$

where W is the vector of holdings' weights and v is the covariance matrix of the holdings' returns, itself being a function of the holdings' volatilities and their correlations.

It clearly follows that the parameters which increase or reduce portfolio variance are (a) the holding weights or mix, (b) the volatility of those holdings and (c) the correlation between the holdings.

Volatilities and correlations are notoriously unstable, which leads to unstable portfolio volatilities through time.

Ex-ante, the manager only controls the holding weights; volatilities and correlations can only be estimated. Both of these parameters are notoriously unstable, which leads to unstable portfolio volatilities through time. This can be seen in figure 7, which shows the rolling 1-year standard deviations for three static illustrative portfolios.

De-risking by allocating to lower volatility assets will reduce portfolio variance, as will diversification. But what does this mean for risk and our original problem of defining how much we are taking? And is it enough to achieve our targets? We first need to understand the effects of diversification.

Diversification
Principal component analysis (PCA) is a useful tool when investigating the level of portfolio diversification. We can expand PCA using a metric called the 'number of equally weighted independent factors' (NEWIF) to investigate the extent to which independent components have driven a portfolio's returns in the past.
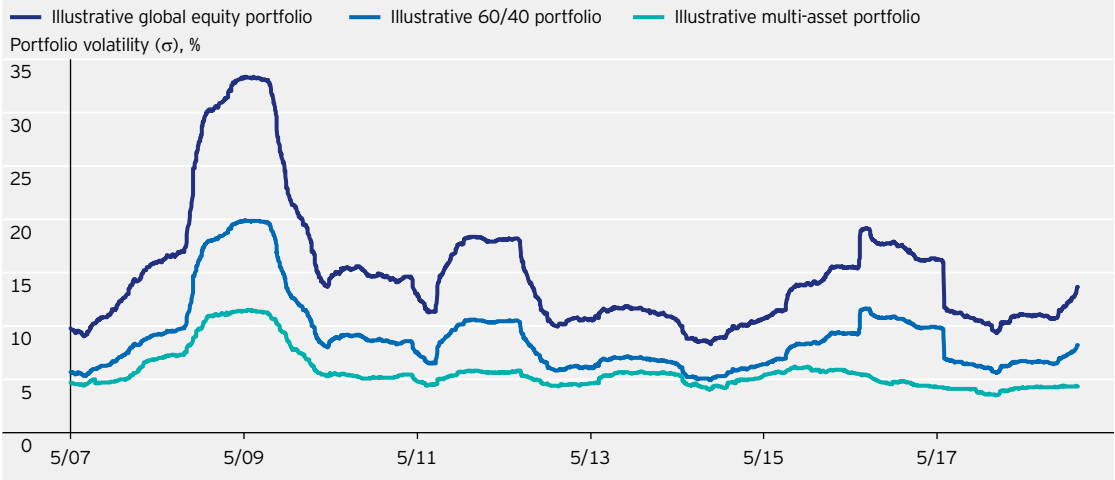
$$NEWIF = e \sum \left(p \cdot \ln(p) \right)$$

where

$$p = \frac{eigenvector(v_w)}{\sum eigenvector(v_w)}$$

Figure 7
Volatilities are notoriously unstable

— Illustrative global equity portfolio — Illustrative 60/40 portfolio — Illustrative multi-asset portfolio

Portfolio volatility (σ), %

Source: Invesco. For illustrative purposes only. Data as at 31 December 2018. Illustrative global equities portfolio represented by MSCI World Index, 60/40 = 60% MSCI World and 40% Barclays global bond indices, and the multi-asset portfolio is an illustrative portfolio targeting an annual return of 3-month GBP LIBOR plus 5%.

ν_w is the covariance matrix of the weighted holdings' returns, which can be calculated as $W_d \cdot \nu \cdot W_d$ with W_d being the diagonal matrix of the holding weights.

Simplistically, this number can be interpreted as how many purely uncorrelated factors were driving portfolio returns in a given period. We are effectively calculating the number of equally-weighted, statistically independent positions that the diversification prevalent in the portfolio during that period would have implied.

For the same three static illustrative portfolios in figure 7, we have calculated the NEWIF for the same rolling one-year periods (figure 8). Clearly, there is more going on in the illustrative multi-asset portfolio, as shown by the persistently greater number of independent factors at play. Also visible is the effect of unstable diversification, where markets can – and do – provide greater or lesser levels of diversification in different periods.

Clearly there is more going on in the illustrative multi-asset portfolio, as shown by the persistently greater number of independent factors at play.

Risk
Our goal is to know whether we are holding enough risk in the portfolio to give us a good chance of reaching our return objectives. If we are right in our views, will we be able to achieve our capital targets?

Figure 8
How many equally weighted independent factors are in your portfolio?

— Illustrative global equity portfolio — Illustrative 60/40 portfolio — Illustrative multi-asset portfolio

Number of equally weighted independent factors (NEWIF)

Source: Invesco. For illustrative purposes only. Data as at 31 December 2018.

Knowing the portfolio variance and the NEWIF enables us to calculate the pre-diversification or 'internal risk' of the portfolio.

If all the holdings in a portfolio are uncorrelated the portfolio variance is merely the sum of the holdings' weighted variances. So, if the portfolio variance is known and all holdings are of equal weight and variance, we can solve for the volatility of each factor:

$$\sigma_f = \sqrt{\frac{\sigma_p^2}{NEWIF}}$$

The Internal Portfolio Risk can then be calculated as the sum of the independent factors' volatilities:

$$Internal\ Portfolio\ Risk = NEWIF \cdot \sigma_f$$

By normalising the diversification effect through PCA and calculating the Internal Portfolio Risk, we get an idea of how much risk we have in the portfolio – something diversified volatility might be hiding. We can then compare this to our capital target and model or assess whether we have enough risk to meet it.

The rolling calculation for our three illustrative portfolios is shown in figure 9, and figure 10 compares the volatility and Internal Portfolio Risk of all three sample portfolios as at 31 December 2018. The low levels of historic portfolio volatility of the illustrative multi-asset portfolio mask near equity-like levels of Internal Portfolio Risk.

If we hold only one asset in a portfolio, the Internal Portfolio Risk is equal to its volatility. But this risk has only one outcome driver. By holding the same amount of internal risk but split across multiple outcome drivers (the NEWIF), we can facilitate portfolio-level volatility reduction without compromising the level of risk.

Nevertheless, the manager still needs to be right about the positions to generate the returns, but

their opportunity set – the things he or she can be right about – is larger.

We have now effectively estimated the risk implicit in the portfolio. While volatility (or standard deviation) can give us an idea of how the portfolio may behave in the short term, the internal risk allows us to look through the effects of short-run diversification and assess how much risk there is to generate capital returns.

If managers are right in their position selection, it is the level of Internal Portfolio Risk that will determine the size of long-run portfolio outcomes.

If managers are right in their position selection, it is the level of Internal Portfolio Risk that will determine the size of long-run portfolio outcomes.

Conclusion
When trying to assess the level of risk embedded in a portfolio, we propose that considering a portfolio's statistical volatility metrics alone can be misleading when used to infer return potential.

We show that longer-term returns are not an outcome of the level of assumed volatility, but rather a by-product of the skill of the manager in selecting positions and the combination of these positions in the portfolio. Volatility (or standard deviation) is only an outcome of how the assumed risks behave.

As such, we believe the return potential of a portfolio is reflected most accurately through analyzing its embedded risk (Internal Portfolio Risk) - determining this requires an assessment of how diverse the drivers of a portfolio's returns truly are.

Figure 9
Internal Portfolio Risk: a better reflection of the portfolio's risk

— Illustrative global equity portfolio — Illustrative 60/40 portfolio — Illustrative multi-asset portfolio
Internal Portfolio Risk, %
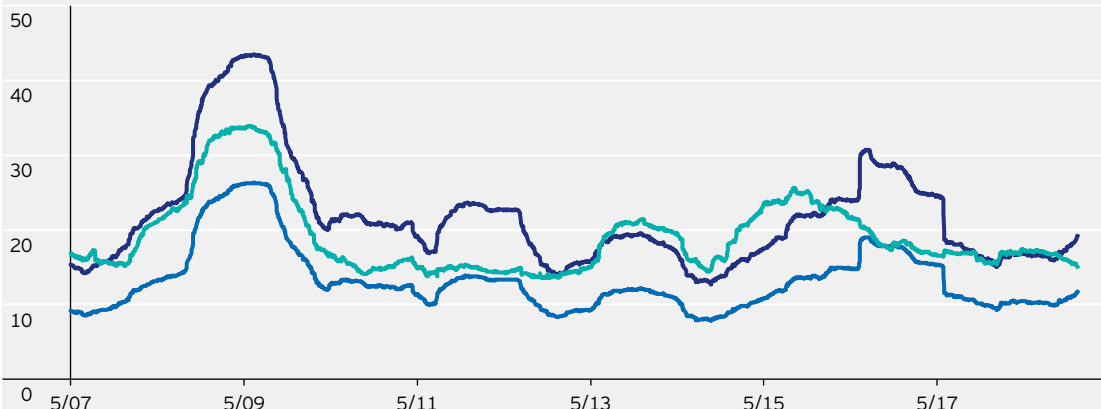
Source: Invesco. For illustrative purposes only. Data as at 31 December 2018.

Figure 10
Volatility and risk are not the same thing

■ Internal Portfolio Risk   ■ Volatility

Source: Invesco. For illustrative purposes only.

A manager should, therefore, be able to reduce portfolio-level volatility without compromising the level of actual risk if a given level of internal risk is split across many truly independent factors. Using this approach in a portfolio context, it is fully possible to achieve high information ratios – high returns for a given leve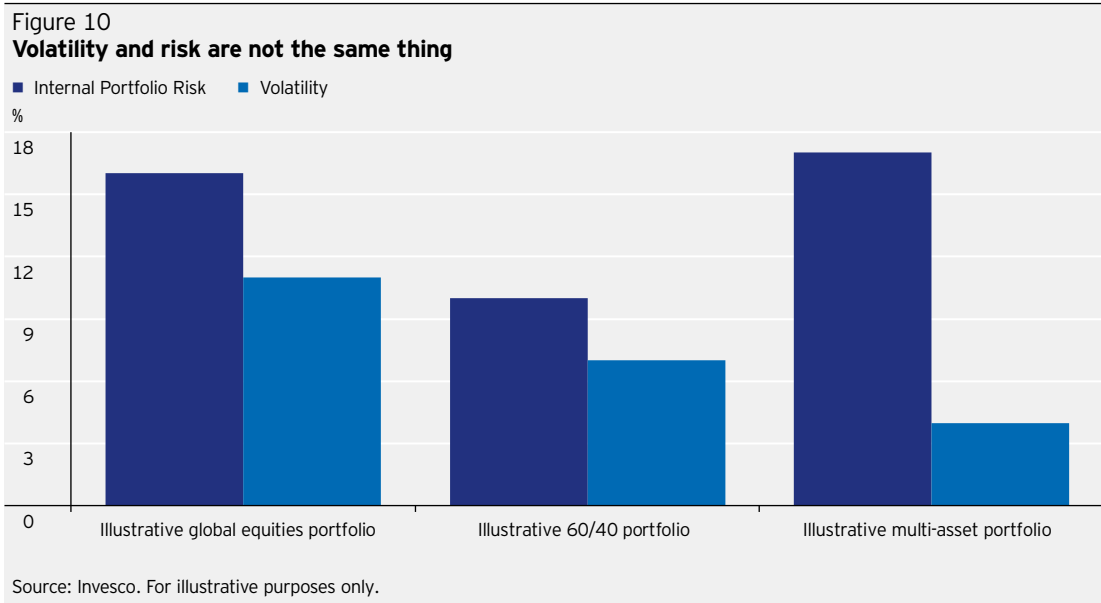l of ex-post volatility – provided the manager achieves a positive and persistent hit rate, positive return skew and if the portfolio risks interact so as to achieve low net day-to-day volatility.

A manager should be able to reduce portfolio-level volatility without compromising the level of actual risk if a given level of internal risk is split across many truly independent factors.

About the author

Michael Marshall, CFA, PRM
Risk Manager,
Invesco Multi Asset
Michael Marshall assists with the construction and risk management of our multi-asset portfolios.

Adding equity upside potential to fixed income portfolios

By Robert Young and Leyla Greengard

In brief

The low interest rate environment has compelled investors to search for investment strategies to improve portfolio performance while helping mitigate overall risk. A unique way to achieve this goal can be tactically pairing fixed income securities with equity options. Using this strategy, investors create convex equity return profiles that allow them to participate in the long-term potential of equity investments – with limited downside risk. We explain how this strategy works, analyze its performance and discuss its advantages over a traditional convertibles strategy. We also explore the benefits of including it in larger portfolios and ways to customize the strategy around investor objectives.

Low interest rates pose a significant challenge to fixed income investors. But measures to enhance returns often come with higher risks. We illustrate how enhancing a fixed income portfolio with equity options can lead to a more appealing risk/return profile.

Since the global financial crisis, strong equity markets, low interest rates and muted volatility have led to a prolonged period of market complacency. However, concern about the length of the equity bull market, coupled with interest rate uncertainty, has prompted many investors to search for ways to enhance returns without introducing meaningfully more risk to their portfolios.

A sophisticated way to attain this outcome can be to tactically combine fixed income securities with equity options. These two market exposures, in strategic combination, can offer investors a powerful, non-linear payout structure – similar to convertible bonds but far more flexible.

One critical element and advantage of this approach is that the fixed income and equity components are tailored independently to reflect compelling investment views. This distinguishes our concept from traditional convertible bonds, where the two components are usually from the same company.

Figure 1
Convex return profile of a portfolio containing of bonds and equity options
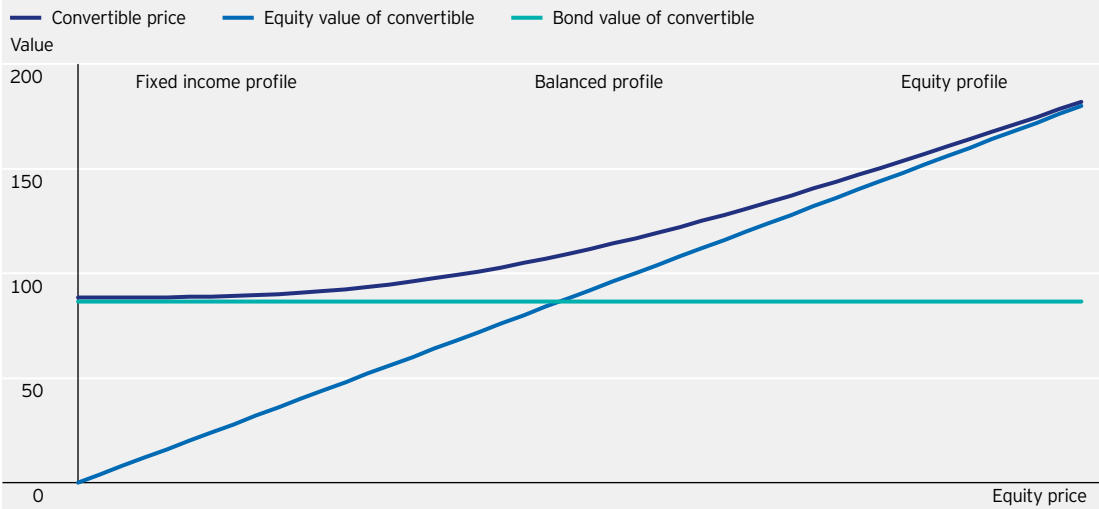
Source: Invesco and Bank of America Merrill Lynch. For illustrative purposes only.

The bond exposure can be sourced from broad fixed income universes and can therefore be tailored to reflect an investor's credit quality and duration requirements. The equity exposure can likewise be obtained from a broad equity universe (including equity indexes, exchange traded funds or individual stocks) and can be customized to introduce various degrees of equity sensitivity depending on investor goals. This flexibility creates the ability to optimize portfolios and provides a solution to help meet a broad range of investor objectives.

This strategy is designed around a very important principle: capital preservation. To minimize downside exposure, the equity option risk budget is set at the expected coupon from the bond portfolio. As a result, investors know in advance their approximate maximum loss potential.

Three reasons why this strategy should be considered as part of a well-balanced portfolio:
1. Attractive risk/return profile
Unlike traditional fixed income and equity return profiles, which are linear in nature, the return profile of the approach is convex. The curved return (figure 1) results in an increasing level of equity participation as

equities rise and a decreasing level of participation as equites fall. This profile is very appealing because it lies between a traditional fixed income portfolio and a traditional equity portfolio. The non-linear nature helps smooth the portfolio return and reduce portfolio risk.

To illustrate this, we constructed a portfolio consisting of the Bloomberg Barclays US Corporate Bond Total Return Index (90%) and call options on the S&P 500 Total Return Index (10%).[1] We then simulated the performance for the 10-year period from 30 June 2009 to 30 June 2019 and analyzed the monthly returns.

Table 1 shows the annualized returns and volatilities of this portfolio, call options on US equities and US corporate bonds over 1, 3, 5 and 10-years; as at 30 June 2019. The results show that, over the long term, the addition of a modest equity option allocation to a fixed income portfolio meaningfully increases overall portfolio return.

Although this strategy looks slightly more volatile than a pure corporate bond portfolio (figure 2), the limited equity option allocation can lead to a very

Table 1
Tactically pairing equity options and fixed income securities can improve overall risk/return profile

	Return p.a.:			Risk p.a.:			Return/risk p.a.:		
	Strategy	S&P 500 call options	US corporate bonds	Strategy	S&P 500 call options	US corporate bonds	Strategy	S&P 500 call options	US corporate bonds
1yr	10.4%	8.2%	10.7%	5.1%	17.5%	4.1%	2.02%	0.47%	2.58%
3yr	8.6%	11.9%	3.9%	6.3%	11.9%	3.9%	1.36%	1.00%	1.01%
5yr	7.5%	8.5%	4.1%	6.8%	11.8%	3.9%	1.11%	0.72%	1.04%
10yr	8.1%	12.3%	6.1%	6.7%	12.6%	4.3%	1.21%	0.98%	1.42%

Source: Invesco, Bloomberg as at 30 June 2019. Risk as measured by the standard deviation of returns. S&P 500 options; US corporate bonds as measured by the Bloomberg Barclays US Corporate Bond Total Return Index. **The figures refer to simulated past performance and past performance is not a reliable indicator of future performance.**

Figure 2
Volatility in comparison

— Strategy — S&P 500 call options — US corporate bonds

Rolling volatility

Source: Bloomberg, Invesco. Data as at 30 June 2019.
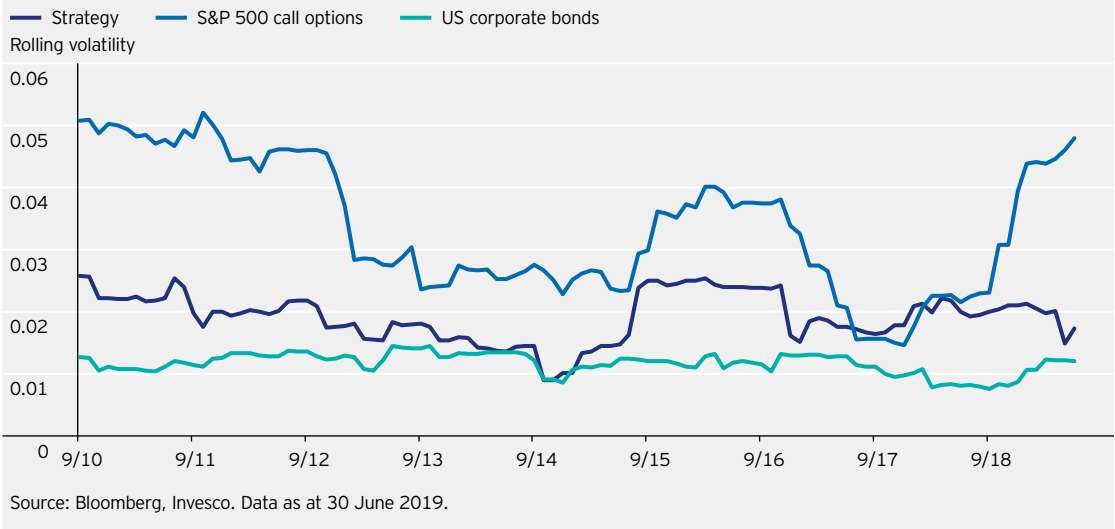
attractive risk/return ratio: the increase in risk for the equity exposure is minimal. Additionally, due to the intentional convexity of the option payoff, volatility in general remains significantly lower than the S&P 500.

More importantly, the strategy's downside during periods of S&P 500 drawdowns is limited by both the fixed income exposure and the convex option return profile (figure 3).

The data shows that this strategy would have added to the return of the Bloomberg Barclays US Corporate Bond Total Return Index while, in most cases, improving the risk/return profile. For investors who want to participate in the upside of equity markets but also desire some level of downside risk mitigation, this strategy may offer an attractive investment alternative.

2. Broader investment universe and greater liquidity than traditional convertibles

This approach sources market exposures from broad investment universes, eliminating the capacity constraints found in some asset classes. The larger universes also provide important liquidity advantages: the equity and fixed income components are bought and sold independently within their respective markets, rather than combined as a single security as with traditional convertible bonds. This provides a significant liquidity advantage, especially during challenging market environments.

Convertible bonds have traditionally been used to achieve this type of portfolio profile. However, the convertible asset class is facing significant headwinds as a result of persistent low interest rates. Today, the convertible universe is about half the size it was at its peak in 2008, resulting in a much smaller investment universe (figure 4). In addition, the lower level of

Figure 3
Performance during S&P 500 drawdowns in comparison
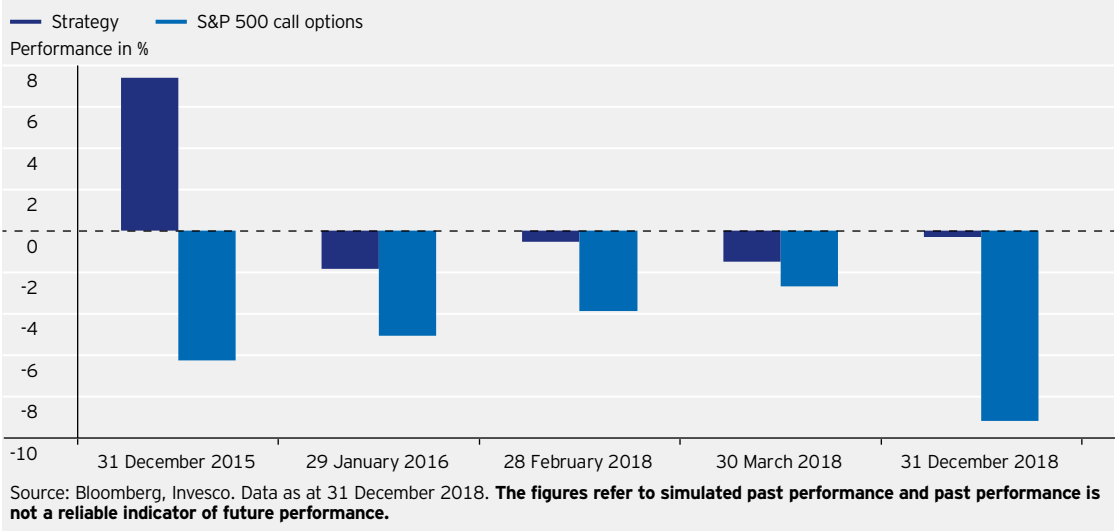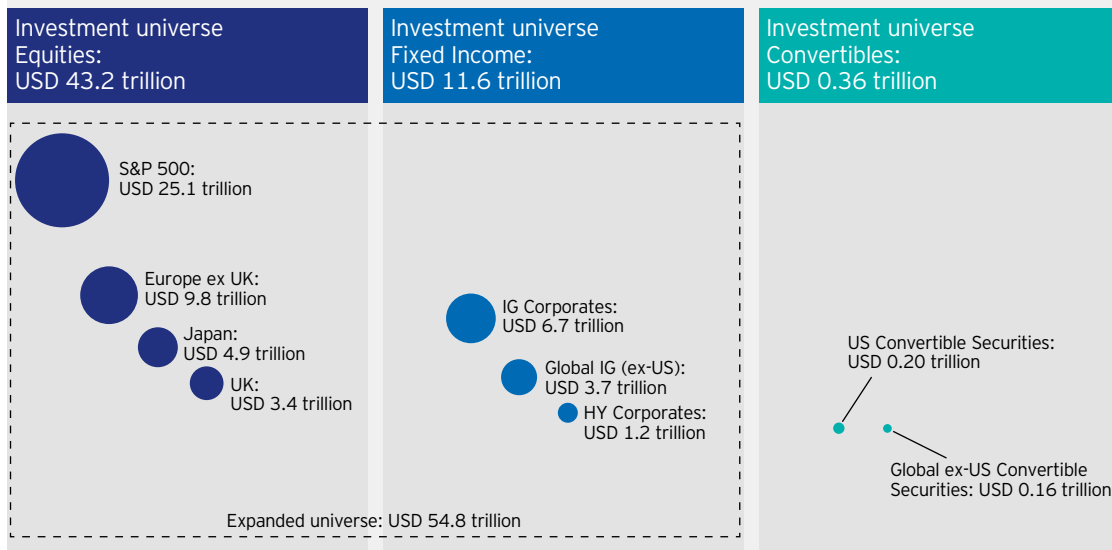
— Strategy — S&P 500 call options

Performance in %

Source: Bloomberg, Invesco. Data as at 31 December 2018. **The figures refer to simulated past performance and past performance is not a reliable indicator of future performance.**

Figure 4
Investment universes in comparison

Investment universe Equities: USD 43.2 trillion	Investment universe Fixed Income: USD 11.6 trillion	Investment universe Convertibles: USD 0.36 trillion

S&P 500:
USD 25.1 trillion

Europe ex UK:
USD 9.8 trillion

Japan:
USD 4.9 trillion

UK:
USD 3.4 trillion

IG Corporates:
USD 6.7 trillion

Global IG (ex-US):
USD 3.7 trillion

HY Corporates:
USD 1.2 trillion

US Convertible Securities:
USD 0.20 trillion

Global ex-US Convertible
Securities: USD 0.16 trillion

Expanded universe: USD 54.8 trillion

Source: Bloomberg. Data as at 30 June 2019.

liquidity in convertibles, along with the rigidity and inflexible nature of their structure, makes them less compelling in our view. The independent bond and option strategy may be more effective and could potentially provide enhanced opportunities to generate alpha.

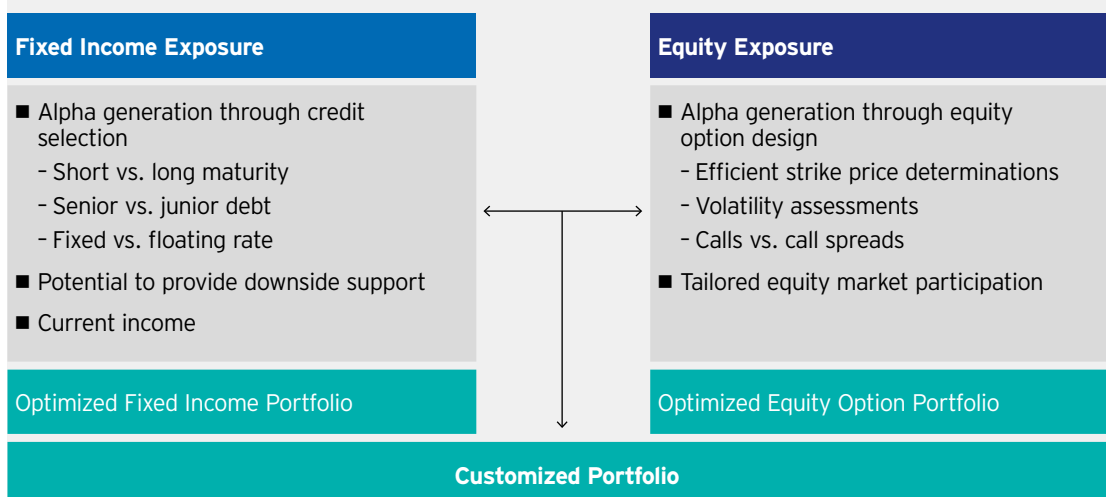**3. Customizable for different investor segments**
A critical advantage of this strategy is that both the fixed income and the equity option allocations are managed separately. A manager can tactically combine the two exposures to provide the optimal risk/return profile (figure 5).

As mentioned previously, for the fixed income exposure of the strategy, an investor can choose securities from a broad, liquid universe based on macroeconomic views, sector calls, individual credit determinations and specific investor requirements. This allows investors to tailor the fixed income exposure to their individual needs – e.g. credit quality, maturity, coupon type or level of capital structure.

For the equity option allocation, underlying securities can be chosen from a large, global investment universe. A manager can then identify option strike prices and expiries based upon current market conditions to reflect a client's desired equity sensitivity. The strategy can take a flexible approach to rebalancing, based either on the delta levels of the options (e.g. maintaining delta in a range to maximize convexity) or on a calendar basis. The manager can then pair these two separate exposures to create an optimal risk/return dynamic.

Figure 5
Combining two exposures with different purposes

Fixed Income Exposure	Equity Exposure
■ Alpha generation through credit selection - Short vs. long maturity - Senior vs. junior debt - Fixed vs. floating rate ■ Potential to provide downside support ■ Current income	■ Alpha generation through equity option design - Efficient strike price determinations - Volatility assessments - Calls vs. call spreads ■ Tailored equity market participation
Optimized Fixed Income Portfolio	Optimized Equity Option Portfolio

Customized Portfolio

Source: Invesco. For illustrative purposes only.

Important characteristics of equity options

We believe it is helpful to first remind investors about the role equity options can play in a portfolio and how they can be used to improve risk-adjusted returns. Compared to a traditional equity position in the S&P 500, which is linear in nature, call options on the S&P 500 offer two advantages. First, they allow an investor to participate in the upside of a stock's performance and, second, they provide this exposure in a convex, non-linear manner.

Some of the most important options characteristics are outlined below:

Sensitivity to the underlying price: *Delta is an estimate of how much an option's price will change given a move in its underlying asset. As the underlying price increases or decreases, the delta of the option also increases or decreases. Delta ranges from 0% (no equity sensitivity) to 100% (100% equity sensitivity). The rate of change in the delta, or the convexity of the option, is measured by gamma. This strategy is purposefully exposed to a high level of gamma to take advantage of the convexity of option returns.*

Option value and delta as a function of underlying price
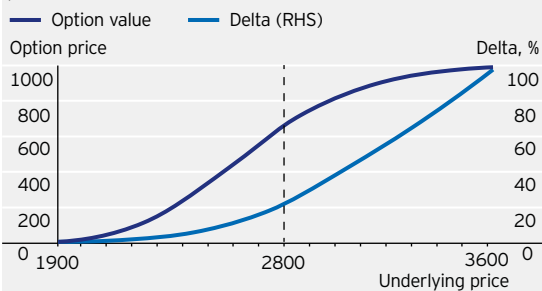
The graph illustrates the convexity (gamma) of the option value around the strike price, which is 2800. Note that the delta of the option is near zero when

the option is far out of the money, and it flattens at 100% when the option is far in the money.

Sensitivity to volatility: *The implied volatility of an option is the market's expectation of the fluctuation of a stock's price over a certain time period. In general, as market volatility increases, the implied volatility and the option value increase.*

Time sensitivity: *The value of options normally decreases as the option moves closer to expiry. The further away an option's expiry date, the slower the loss of value with the passage of time. Therefore, we focus on option expiries that are relatively long in nature to minimize time decay so we can spend time on identifying attractive underlying investments. The sensitivities of the option value to these different parameters are interconnected and each of them, in turn, depends on the others. However, in our analysis, of all the above sensitivities, the delta and the sensitivity to implied volatility are the most important factors contributing to changes in option prices.*

Derivatives risks: *The value of a derivative instrument depends largely on (and is derived from) the value of an underlying security, currency, commodity, interest rate, index or other asset (each referred to as an underlying asset). In addition to risks relating to the underlying assets, the use of derivatives may include other risks, including counterparty, leverage and liquidity risks. Derivatives create leverage risk because they do not require payment up front equal to the economic exposure created by owning the derivative. As a result, an adverse change in the value of the underlying asset could result in a loss that is substantially greater than the amount invested in the derivative, which may make the return more volatile and increase the risk of loss. Derivatives used for hedging or to gain or limit exposure to a particular market segment may not provide the expected benefits, particularly during adverse market conditions.*

Figure 6
This approach allows a manager to tailor portfolios based on an investors' equity sensitivity needs
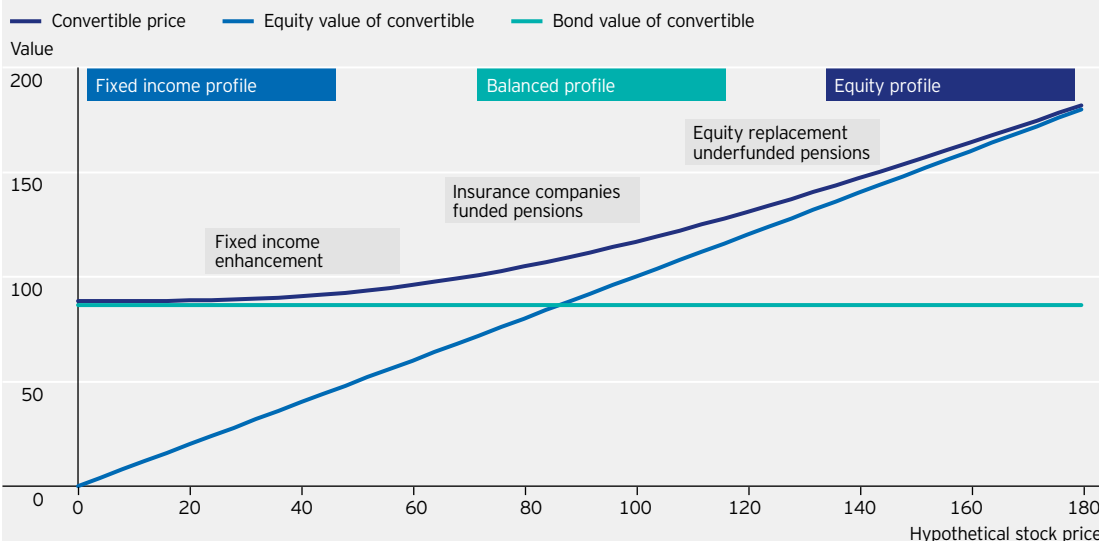
Sources: Invesco and Bank of America Merrill Lynch. For illustrative purposes only.

The fixed income security and equity option work together to form a convex risk/return profile. This pair can then be tilted in various ways to reflect unique investor risk and return needs by allocating differently across the fixed income and equity components (figure 6). For example, certain investors (e.g. underfunded pensions) may desire more equity sensitivity, and others less. A portfolio can be positioned at different points across the return spectrum to reflect those needs without changing the research or portfolio construction process.

Conclusion

A professionally managed portfolio can be constructed for a broad array of investor needs, introducing various levels of risk/return and asymmetric upside potential. In our view, the strategy can combine an attractive risk/return profile with broad market access and greater potential for alpha. The strategy can provide customized solutions without capacity constraints – and thus help institutional investors strengthen their overall asset allocation and navigate today's complex markets.

About the authors

Robert Young, CFA
Head of Institutional Convertibles,
Invesco Fixed Income
Robert Young is Head of Institutional Convertible Securities and a Senior Portfolio Manager for Invesco Fixed Income. Mr. Young manages and trades for institutional investment grade and non-investment grade convertible portfolios.

Leyla Greengard, PRM™
Head of Fixed Income Portfolio Risk,
Invesco Fixed Income
Leyla Greengard served as Head of Fixed Income Portfolio Risk from November 2011 – December 2019. In this role she analyzed portfolio risk and helped portfolio managers understand the sources of market risk in their portfolios.

Note

1   Options are SPX European calls, which give holders the right to buy the underlying on a given date (expiry) at a predetermined price (strike price). We assumed that options had 2 years to expiry and a strike price 2% higher than the underlying price at the time of purchase (out of the money). Option values were calculated using the Black-Scholes formula. The 2-year Treasury rate at the time of purchase was used as the discount rate and the VIX at the time of purchase as implied volatility. For ease of calculation, the portfolio was rebalanced monthly at month-end.

Forging ahead in China

A conversation with Andrew Lo on the China Position Study, a survey of global sentiment towards China-related investing.

Andrew Lo, Senior Managing Director und Head of Asia Pacific, Invesco Asia Pacific

Risk & Reward sat down with Andrew Lo, Senior Managing Director and Head of Asia Pacific, to talk about the recent China Position Study, what the survey results mean, his views on China and what the future of investment in China could look like. With almost four decades of experience in the investment management industry, Mr. Lo has seen the landscape evolve over time, ushering in a meaningful and fundamental shift with the emergence of China.

Risk & Reward
According to the study, how are investors approaching China-related investment?

Andrew Lo
Despite the ongoing uncertainty and softening global economic sentiment, survey respondents appear focused on exploring the long-term opportunities in China. For example, amid ongoing trade tensions, some 68.4% of respondents still reported that they are planning to increase investments in China over the next 12 months.

We know that China is a large and rapidly growing economy that presents countless opportunities, and that's stoking global investor interest. Indeed, 40.2% of respondents said that the growth potential of China's listed companies is driving their decision for a dedicated investment exposure to China.

We also note that market transparency and greater accessibility are providing additional impetus. Respondents cited better corporate reporting, legal protections for foreigners and the sophistication of financial intermediaries as further reasons behind committing investment resources to China.

Despite trade war uncertainty and a softening global economic outlook, **The China Position Study**, a global survey conducted by the Economist Intelligence Unit (EIU) and commissioned by Invesco[1] revealed bullish attitudes among professional investors and asset owners towards China-related investment. However, respondents in North America and Europe, the Middle East and Africa (EMEA) were more optimistic than those in Asia Pacific (APAC).

The study showed that over 80% of global investors surveyed plan to increase (either moderately or significantly) their organization's allocations to Chinese investments over the next 12 months, with only 4% planning to reduce exposure to China over the same period.

North American respondents are broadly bullish in their economic outlook across markets, with over 80% of respondents expecting better economic conditions both globally and in China over the next 12 months.

The study was conducted throughout August and September 2019 and received responses from 411 professional investors across North America, EMEA and APAC.

Overall, we think respondents are optimistic about China over the long run because it's becoming easier to understand and access its markets. And that denotes a shift in investor sentiment towards China – until recently, China was seen as a source of capital, not an investment destination in its own right. Now, organizations see a clear need to carve out a specific allocation to China. This is a key development in China's market evolution.

Risk & Reward
What have clients been telling you about investing in China?

Andrew Lo
There are two major themes I hear from our clients:

The first is that the index inclusion of Chinese equities and bonds is really making them take notice of the market, sometimes in a new light. Many have told us that, while they've been impressed by China's continued economic growth and reforms over the years, they were still skeptical about the attractiveness of China. They had questions like – "Yes, I can see the good economic performance, but does this translate into investment returns for us?" or "What about government and state intervention?"

So, the inclusion is a starting point for them to think differently about the market. It motivates our clients to speed up their learning about and understanding of China. I think this is where we can help them by sharing our knowledge about how China can fit into their portfolios.

Secondly, the survey showed that there are some segments of institutional investors willing to forge ahead in terms of their China exposure, be that through dedicated mandates or ETFs. These are the ones who want to explore opportunities, and whom we often partner with – we have the experience and the know-how to guide them.

Risk & Reward
Are the survey findings in line with your expectations, or did they surprise you?

Andrew Lo
I think what immediately jumped out at me and my team was the bullishness with regard to investing in China, particularly in the current climate. That was a bit surprising. But taking a step back, we see that the direction, the trajectory the results show is consistent with our expectations.

This bullishness is heartening. It shows a willingness and a level of sophistication in how some global investors view China, particularly how access to China's capital markets has improved over the years.

All this is very encouraging because it demonstrates a commitment to reforms on the part of the Chinese government, and already, through the survey results, we can see that investors appreciate these efforts.

Risk & Reward
How do the findings fit in with what we see happening around the world, in particular trade tensions?

Andrew Lo
The ongoing trade tensions cannot be ignored. You have the world's two largest economies in disagreement over trade, which is an important growth engine for the global economy. There are also some deeper structural issues, such as knowledge transfer mechanisms between the two countries, that need to be resolved. If handled incorrectly, there will be serious implications for everyone.

We see this uncertainty coming through in the survey results. Although the overall tone of the survey results were bullish, respondents were more

muted when asked about trade tensions – only 42% said that there would be a positive impact in the next 12 months, compared to 44% of respondents who have a negative view.

Personally, I'm more optimistic. I think the trade tensions are a chance for China to speed up market reforms. Historically, we've seen China push ahead with change whenever the country was at a major crossroads. This time, the Chinese government has been similarly measured and thoughtful in their policy responses. They are incentivizing innovation in the economy. They are keeping a tight lid on shadow banking. They are opening up financial markets even further – China recently announced that it would remove quota limits for offshore investors on two cross-border investment schemes: the Qualified Foreign Institutional Investors (QFII) and the Renminbi Qualified Foreign Institutional Investor (RQFII) schemes.

How will you use the study's findings to have better discussions with clients and share Invesco's China-related knowledge and experience?

Andrew Lo
For a start, we shared the results with Asia Pacific institutional clients at our recent conference in Beijing on 7 November 2019. EIU presented a session on the results, and we were glad to engage with the attendees on the many questions they had. Further afield, we've shared with the media, including flagship financial news outlets.

The survey also informs our go-to-market strategy in China. Insights gleaned from it help us gauge market sentiment towards investing in China and highlight where the market gaps are.

More importantly, the study benefits the entire industry. We understand that China, as a large, fast-changing and dynamic emerging market, may be hard to understand for foreign investors. The study helps our clients hear what their peers are doing about China. We believe it's an important piece that makes the Chinese market more comprehensible to all. This way, we can all play our part in improving the global and China's asset management industries.

How would you describe Invesco's commitment to the China market?

Andrew Lo
Simply put, we are in for the long term. We think external interest in China's market will continue to increase, especially given the rapid growth of the Chinese investment industry. If you look at forecasts, some are expecting that China will be the second-largest asset management market globally by the year 2025.[2] That's a rapid development considering that just six years ago, China was in eighth position.

As an independent asset management firm, we have established a strong background of experience and a respected market position in China – we intend to capitalize on the expected growth of the industry for our clients. Moreover, we have a strong institutional operation in China, and we are recognized for our market leadership and innovation. We have been on the ground in the Chinese Mainland for almost two decades, so you know the market well.

And that, I think, is where our greatest strength lies: we have a lot of experience and knowledge to share with international investors. Our many years of presence on the ground also means that we are trusted partners within the market. We have good, diverse, global platforms that offer a wealth of opportunity in China to onshore and offshore clients.

By deepening our engagement and showing that we can improve their understanding of how China relates to them – and to global markets in general – we endeavor to achieve better outcomes for potential investors. We've just successfully held our third Asia Pacific regional conference for institutional clients – one of the best buy-side conferences in the region. All three conferences have been held in China – which goes to show how important this market is to us.

Investing in any market takes knowledge and experience. The survey respondents recognize that – it is the second most-cited objective for why respondents maintained a dedicated China exposure, even higher than alpha generation. Our focus is to deliver the knowledge and experience that can result in opportunities to invest in China.

Thank you, Andrew, for your insight.

About risk
The value of investments and any income will fluctuate (this may partly be the result of exchange rate fluctuations) and investors may not get back the full amount invested. When investing in less developed countries, you should be prepared to accept significantly large fluctuations in value. Investment in certain securities listed in China can involve significant regulatory constraints that may affect liquidity and/or investment performance.

A factor-based buy-and-hold strategy for bonds

By Jay Raol, Ph.D., and Amritpal Sidhu

In brief

We show how factors can be used to build traditional buy-and-hold portfolios investing in bonds with fixed maturities that are held to maturity. Next, we combine several such fixed maturity buy-and-hold portfolios to form a "ladder", seeking to outperform a traditional fixed income benchmark with no turnover. We believe this option may be a solution to liability matching problems faced by insurance and pension funds.

A recent study by Invesco suggests that investors are increasingly turning toward factor strategies because they provide a well-researched and cost-effective way to potentially outperform traditional market value weighted indices.[1] While factor investing has historically been associated with equity markets, new research also points to its usefulness for corporate bond investing.[2] Given the relatively short track record of fixed income factor investing, however, some investors are concerned with implementation hurdles and the potential for excessive trading.[3] We aim to show that, even in portfolios with no trading - i.e. buy-and-hold portfolios - factor-based strategy can add value.

> While factor investing has historically been associated with equity markets, new research also points to its usefulness for corporate bond investing.

We start with a universe of US investment grade corporate bonds with maturities from four and a half to five and a half years. From this universe of bonds, two portfolios will be formed - a factor portfolio and a passive portfolio.

Figure 1
Total return comparison of five-year factor and market portfolios

■ Portfolio ■ Benchmark

Total return (%)

Vintage year

Source: Bloomberg L.P. and Invesco calculations. Data from 1 January 1990 to 31 December 2017. The figure shows the annualized returns of the five-year factor and market portfolios over the lives of the portfolios. The returns are dated by the initial investment year (vintage). **The figures refer to simulated past performance and past performance is not a reliable indicator of future performance.**

For the factor portfolio, the bonds in this universe are scored based on their exposure to the three factors quality, carry and value.[4] Each bond's overall score is a 10/40/50 percent blend of these factor scores. The portfolio is formed by taking half of the bonds in the universe (by market value weight) with the highest blended score and forming a market-value-weighted portfolio called the "factor five-year portfolio". The passive portfolio is formed by taking all of the bonds in the universe and market value weighting them. It is called the "market portfolio". Every year, at the beginning of January, this process is repeated, so that we obtain a series of market and factor portfolios with different vintages.

The bonds in these portfolios are held to maturity as long as they maintain a rating higher than CCC.

In other words, no change is made to the portfolios unless a bond approaches imminent default, at which time the bonds are sold and the cash proceeds kept in the portfolio. Otherwise, cash from coupons is reinvested pro rata into the portfolio. Proceeds from securities that are called early or mature earlier than the overall portfolio are also kept as cash in the portfolio. While, in practice, cash accumulated in the portfolio would be reinvested, simply accumulating it suffices for our simulation to illustrate the value of a factor approach.

Figure 1 shows the total returns of the factor five-year portfolios compared to those of the market five-year portfolios for different vintage years; figure 2 shows the associated active returns. The factor portfolios exhibit consistent outperformance against the market portfolios.

Figure 2
Active returns of the five-year factor portfolios

Excess return (bps)

Source: Bloomberg L.P. and Invesco calculations. Data from 1 January 1990 to 31 December 2017. The figure shows the annualized active (excess) returns of the five-year factor portfolio over the market portfolio for different vintage years. For example, the factor portfolio formed on 1 January 2009 outperformed the market value portfolio by 200 basis points annually over the subsequent five years. **The figures refer to simulated past performance and past performance is not a reliable indicator of future performance.**

Table 1
**Summary statistics of factor portfolios with different maturities**

(Portfolio maturity in years)	2	3	4	5	6	7	8	9
Total return	3.83	4.67	5.08	5.66	5.97	6.28	6.72	6.56
Excess return	0.15	0.21	0.20	0.33	0.17	0.22	0.34	0.08
Tracking error	0.65	0.64	0.40	0.47	0.44	0.47	0.46	0.43
Information ratio	0.24	0.33	0.51	0.71	0.39	0.47	0.74	0.19
Starting yield	5.32	5.66	5.91	6.41	6.80	7.04	7.26	7.18
Number of bonds	172	120	135	86	92	92	96	109
Default rate (%)	0.14	0.24	0.46	0.49	1.39	1.22	1.48	2.77
Active default rate (%)	0.08	0.06	0.07	-0.03	0.42	0.33	0.31	1.18
Yield loss from default	0.04	0.05	0.07	0.06	0.14	0.10	0.11	0.18
Active solvency charge	0.4	0.5	0.6	0.6	0.9	0.9	1.1	1.2
High yield share (%)	5	6	9	12	15	16	19	22
Active rating	0.6	0.6	0.6	0.5	0.6	0.6	0.6	0.6

Source: Bloomberg L.P. and Invesco calculations. Data from 1 January 1990 to 31 December 2017. The table shows the excess returns, tracking errors and information ratios of the factor portfolios versus the market portfolios. Tracking errors and information ratios of the portfolios are averaged over the back-test period. **The figures refer to simulated past performance and past performance is not a reliable indicator of future performance.**

Next, we repeat the construction of portfolios with maturities of two through nine years. Table 1 summarizes key statistics and results: the excess returns of the factor portfolios are all positive, which illustrates that the factors work regardless of maturity. The tracking errors are small, but the information ratios (IR) are consistent with those found in monthly rebalanced factor strategies.

The total return of any buy-and-hold strategy is a function of the starting yield less losses due to defaults, forced selling and any cash drag from the reinvestment of coupons, callability of bonds or recovery from default. To better understand the impact of defaults on portfolio returns, we determine the percentage of bonds that ended below a CCC rating during the life of each portfolio (see row labeled "Default rate"). Longer-maturity factor portfolios naturally have higher default rates since the cumulative default probabilities for any portfolio increase over time. The table also shows the active default rates (i.e. the factor portfolios' default rates in excess of the market portfolios' default rates) along with their yield impact. The advantage of factor-based portfolios is that their higher yields more than offset the negative return impact from additional defaults. The findings are consistent with Eisenthal-Berkovitz et al.[5]

Factor-based laddered portfolio construction
To extend the idea of utilizing factors in a zero-turnover portfolio, we use laddered portfolios to create factor-based solutions whose characteristics look similar to broad-based fixed income benchmarks. We construct a laddered portfolio by buying an equal share of fixed maturity buy-and-hold portfolios with durations similar to the chosen benchmark (figure 3). For example, to target a five-year duration portfolio, an equal-weighted portfolio is formed by investing in buy-and-hold portfolios with maturities from one to nine years (targeting a five-year duration). At the end of each year, the proceeds

Figure 3
How a laddered portfolio is constructed

January 2019 – weights

December 2019 – weights
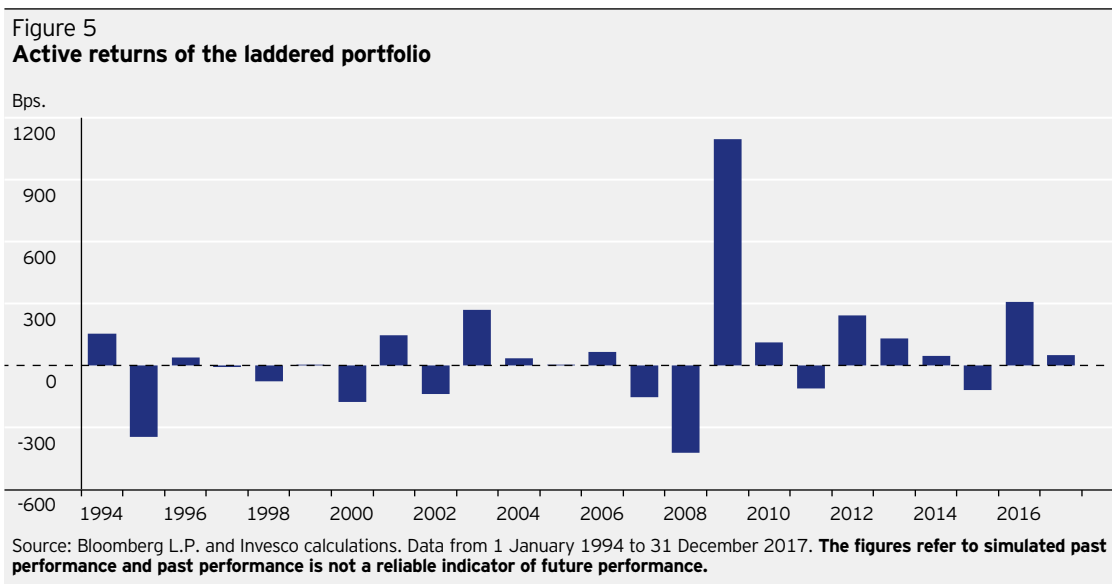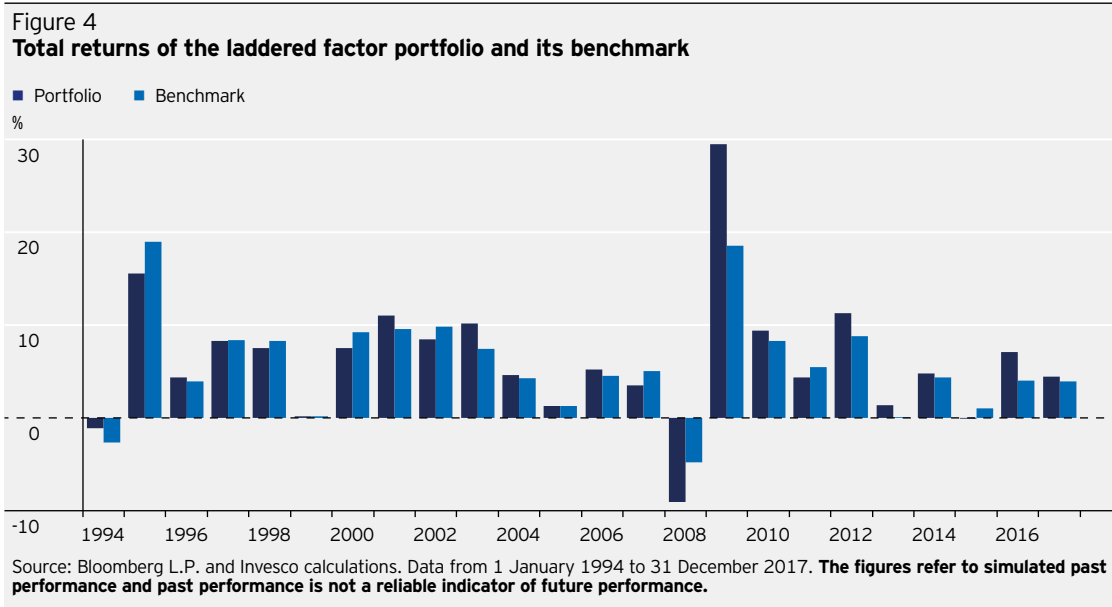
January 2020 – weights
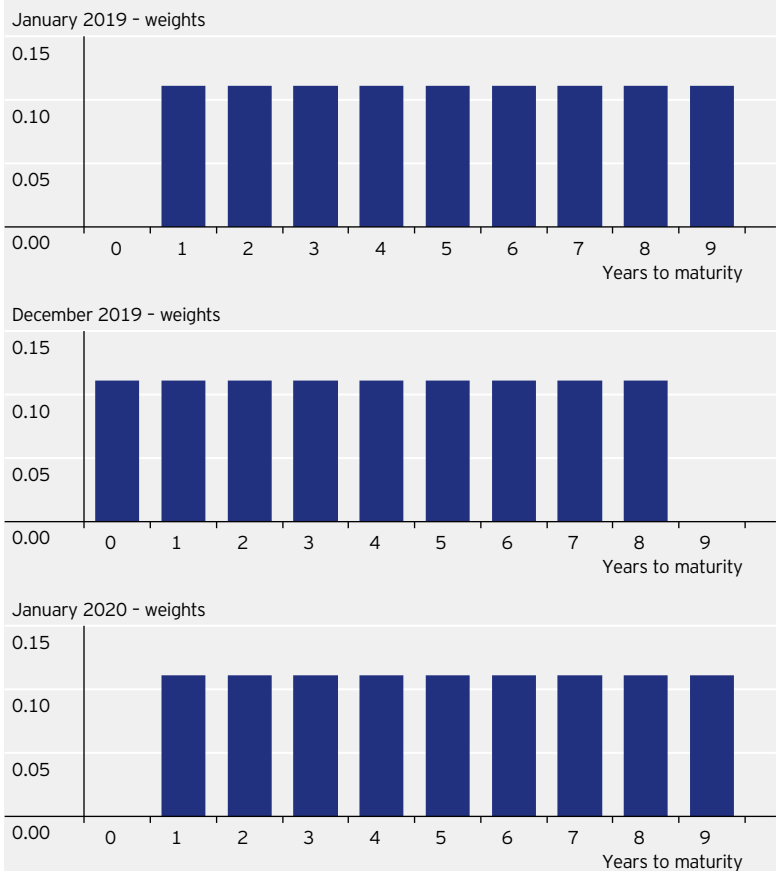
In the first panel, the portfolio is invested, in equal weights, in portfolios of maturities from one to nine years to target a five-year duration. The second panel shows how each portfolio has matured after one year. The cash generated from a maturing one-year portfolio is then used to buy a new nine-year portfolio, as shown in the last panel. In this way, the portfolio maintains a duration close to the desired five years, without incurring high trading costs.
Source: Bloomberg L.P. and Invesco calculations. Data from 1 January 1990 to 31 December 2017.

of the maturing portfolio are used to buy a new nine-year portfolio. This is repeated each year to keep the duration within 0.5 years of the desired portfolio duration.

The performance of the factor-based laddered approach relative to the Bloomberg Barclays Intermediate Corporate Index is shown in figures 4 and 5. Figure 4 shows the total returns for the factor-based portfolio and the market portfolio, figure 5 shows the active returns. The laddered factor portfolio outperforms in most calendar years. Table 2 compares the yield, OAS and duration of the laddered portfolio and the index benchmark. The factor-based approach results in the higher yield and higher return portfolio.

Figure 4
Total returns of the laddered factor portfolio and its benchmark

Source: Bloomberg L.P. and Invesco calculations. Data from 1 January 1994 to 31 December 2017. **The figures refer to simulated past performance and past performance is not a reliable indicator of future performance.**

Figure 5
Active returns of the laddered portfolio

Source: Bloomberg L.P. and Invesco calculations. Data from 1 January 1994 to 31 December 2017. **The figures refer to simulated past performance and past performance is not a reliable indicator of future performance.**

Table 2
Portfolio summary statistics

	Return p.a (bps)	Risk p.a. (bps)	Average yield (bps)	Average OAS (bps)	Average duration
Laddered portfolio	626	708	561	197	4.19
Benchmark	553	502	498	133	3.96
Active	73	311	64	63	0.23

Source: Bloomberg L.P. and Invesco calculations. Data from 1 January 1990 to 31 December 2017. **The figures refer to simulated past performance and past performance is not a reliable indicator of future performance.**

Conclusion

We believe factor-based buy-and-hold solutions can offer clients potentially compelling returns, especially relative to purely passive portfolios. Buy-and-hold portfolios help achieve desired yield outcomes with minimal turnover and transaction costs. Factors can further enhance returns when coupled with well-compensated risks. Indeed, as we demonstrated in this article, a simple factor-based laddered buy-and-hold portfolio beats its benchmark over most of the period covered in our analysis (1994 to 2017). We believe a duration-targeted laddered portfolio offers a simple way to harvest the benefits of a factor-enhanced buy-and-hold strategy.

We believe factor-based buy-and-hold solutions can offer clients potentially compelling returns, especially relative to purely passive portfolios.

About the authors

Jay Raol, Ph.D.
Director of Quantitative Research,
Invesco Fixed Income
Jay Raol is the Director of Quantitative Research for Invesco Fixed Income. His team leads the research in factor-based strategies across fixed income asset classes. He also leads the development of the quantitative tools that support the macro research process and factor-based portfolio construction process across the IFI platform.

Amritpal Sidhu
Quantitative Analyst,
Invesco Fixed Income
Amritpal Sidhu is a Quantitative Analyst with Invesco Fixed Income's Multi-Sector team. In this role, he is focused on researching alternative risk premia strategies in credit and rates.

Notes
1   Invesco Global Factor Investing Study (2019).
2   J. Raol and S. Pope (2018): Why Should Investors Consider Credit Factors in Fixed Income?, AIAR Vol. 7(3).
3   Invesco Global Factor Investing Study (2019).
4   The factor scores are the ranks of the bonds in US investment grade universe after ranking them according to certain characteristics. For carry, the ranking criterion is the bond OAS. For value, bonds are first arranged in similar groups based on sector and rating, and then ranked according to the OAS within each group. Quality score is the rank of the bond based on a blend of rating and maturity.
5   Y. Eisenthal-Berkovitz, J. Hyman, A. Maitra and S. Polbennikov (2014): Try and Hold Credit Investing, Cross Asset Research.

Economic versus statistical clustering in multi-asset multi-factor strategies

By Dr. Martin Kolrep, Dr. Harald Lohre, Erhard Radatz and Carsten Rother

In brief

Maximizing for diversification in the multi-asset multi-factor universe, the literature advances diversified risk parity strategies across economic clusters. For handling overly complex correlation matrices, hierarchical clustering techniques have recently been put forward to guide risk parity allocations. Indeed, such statistical clusters might be considered natural portfolio building blocks given that they automatically pick up the dependence structure and thus form meaningful ingredients to aid portfolio diversification. We explain the intuition and nature of hierarchical clustering techniques in the context of multi-asset multi-factor investing vis-à-vis the use of economic factors in diversified risk-based allocation paradigms such as 1/N, minimum-variance and diversified risk parity.

In an attempt to construct better, more efficient risk-managed portfolios, investors can diversify their portfolios through factors rather than traditional asset classes. Very often, however, this entails a correlation matrix so complex that it cannot be fully analyzed. In this study, we show how this problem can be addressed by using hierarchical clustering techniques and we investigate meaningful ways of generating a coherent multi-asset multi-factor allocation to harvest the associated asset and factor premia in a balanced fashion.

Standard portfolio theory suggests aiming for an optimal risk-return tradeoff by resorting to the seminal mean-variance paradigm of Markowitz (1952). Yet, given the notorious sensitivity of mean-variance portfolio optimization with regard to expected return inputs, one may disregard forecasting returns and focus on estimating risk instead. As a result, researchers have developed various risk-based allocation strategies in pursuit of portfolio diversification.

The literature has advanced diversified risk parity strategies designed to maximize diversification benefits across asset classes and style factors.

An innovative approach to managing diversification was introduced by Meucci (2009). Conducting a principal component analysis (PCA), he aims to identify the main risk drivers in a given set of assets. The ensuing principal components can be viewed as principal portfolios representing uncorrelated risk sources. A portfolio is considered well-diversified if the overall risk is distributed equally across these uncorrelated principal portfolios. Given the statistical nature of PCA, Meucci, Santangelo and Deguest (2015) propose a minimum-torsion transformation to derive uncorrelated risk sources that are economically more meaningful. Along these lines, the literature has advanced diversified risk parity strategies designed to maximize diversification benefits across asset classes and style factors; see Lohre, Opfer and Orszàg (2014), Bernardi, Leippold and Lohre (2018) and Dichtl, Drobetz, Lohre and Rother (2019).

Such an approach is dependent on designing an appropriate risk model, and it comes with several degrees of freedom. Accordingly, the recent literature presents risk parity allocation paradigms guided by hierarchical clustering techniques, prompting Lopez de Prado (2016) to label the technique 'hierarchical risk parity' (HRP). Given a set of asset class and style factor returns, the corresponding algorithm would cluster these according to some distance metric and then essentially allocate risk budgets equally along these clusters. Such clusters might be deemed more akin to natural building blocks than some aggregated factors in that they automatically pick up the dependence structure and are thus expected to form meaningful constituents to aid portfolio diversification.

In this article, we examine the mechanics and merits of hierarchical clustering techniques in the context of multi-asset multi-factor investing. We contrast these techniques with competing risk-based allocation paradigms, such as 1/N, minimum-variance and diversified risk parity. Hierarchical risk parity strategies generally build on two steps: first, hierarchical clustering algorithms uncover a hierarchical structure within the investment universe, represented in a tree-based map. Second, the portfolio weights are obtained by applying an allocation strategy along the hierarchy, which promises to deliver a meaningful degree of diversification.

As estimates of covariance or correlation matrices are subject to estimation errors, filtering correlation-based clusters and networks are meaningful for constructing diversified portfolios.

As estimates of covariance or correlation matrices are subject to estimation errors, filtering correlation-based clusters and networks are meaningful for constructing diversified portfolios resulting in more reliable outcomes, see Tumminello (2010). In this vein, Lopez de Prado (2016) argues that a correlation matrix is too complex to be fully analyzed and lacks a hierarchical structure. Instead of analyzing the full correlation matrix, he suggests considering the corresponding 'minimum spanning tree' (MST) with N nodes (one node for each asset) and only N-1 edges, i.e. focusing on the most relevant correlations. Deriving the MST requires the definition of a distance measure, often referred to as a 'dissimilarity measure'. The MST is naturally linked to the hierarchical clustering algorithm, called single linkage. In a direct way, the MST reflects the hierarchical organization of the investigated assets, and the optimal portfolio weights can be derived by applying an allocation scheme to the hierarchical structure.

Figure 1
Diversified risk parity across economic factors

Source: Invesco. For illustrative purposes only.

Economic factors in the multi-asset multi-factor universe

We consider a multi-asset multi-factor investment universe that combines the traditional asset classes equities, bonds (interest rates), commodities and credit, as well as different style factors. The monthly times series are available for the period from 31 January 2001 to 31 October 2018. The global equity and bond markets are represented by equity index futures for the S&P 500, Nikkei 225, FTSE 100, EuroStoxx 50, MSCI Emerging Markets and bond index futures for US 10Y Treasuries, Bund, 10Y Japanese Government Bonds (JGB) and UK gilts. The credit risk premium is captured by the Bloomberg Barclays US Corporate Investment Grade (Credit IG) and High Yield (Credit HY) Indices; both are interest rate duration hedged to synthesize pure credit risk. Gold, oil and copper indices are chosen to cover the commodity market.

In addition, we consider the four investment styles carry, value, momentum and quality (figure 1). Carry is based on the idea that high yield assets tend to outperform low yield assets, while momentum investors assume that recent winning assets outperform recent losing assets. Quality (or defensive) investing builds on the observation that high quality assets tend to have higher risk-adjusted returns than low quality assets. Value investing is based on the idea that cheap assets (according to a given valuation metric) tend to outperform expensive assets. We source the underlying return time series from Goldman Sachs (GS) and Invesco Quantitative Strategies (IQS). The factor definitions are given in the appendix.

To measure diversification, we consider an appropriate factor model encompassing suitable economic factors.

Figure 2
Minimum-spanning tree
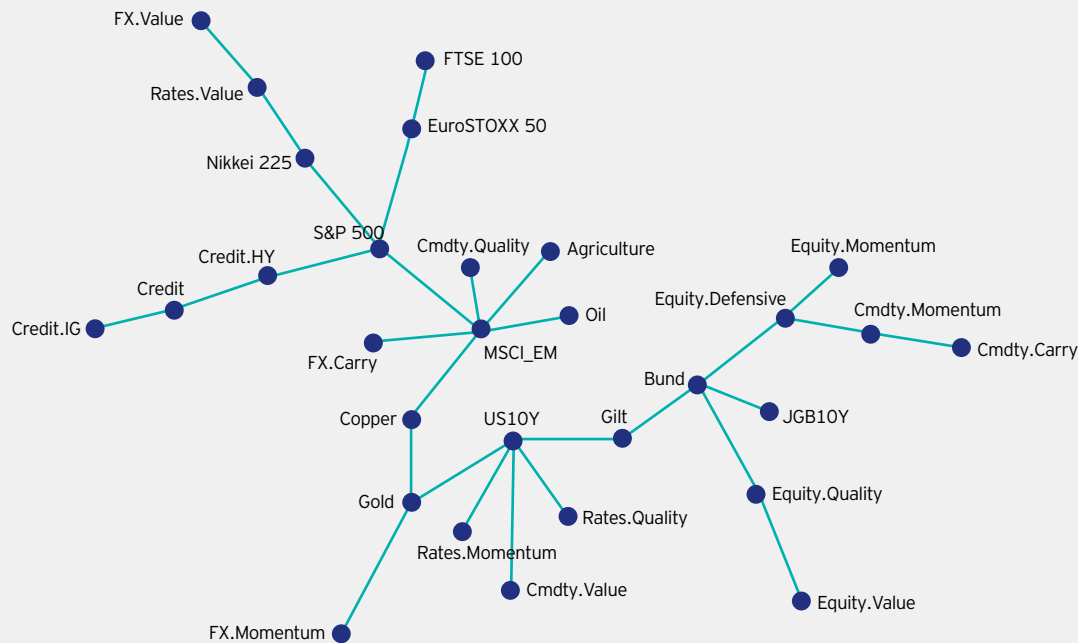
Notes: The figure depicts the correlation network as a minimum spanning tree for the multi-asset multi-factor universe building on the variance-covariance matrix using monthly data from the full sample period from 31 January 2001 to 31 October 2018.
Sources: Bloomberg, Invesco, Goldman Sachs.

To benchmark the statistical clusters vis-à-vis economic factors, we include a parsimonious set of market factors: equity + credit, duration and commodity (figure 1). Further, taking a pure style factor investing perspective, we build aggregate style factors across asset classes, i.e. the aggregate momentum style factor is based on equity momentum, FX momentum, rates momentum and commodity momentum. In the same vein, we construct aggregate carry, value and quality factors.

Diversified risk parity based on economic factors
Striving for a well-diversified portfolio, Meucci (2009) constructs uncorrelated risk sources embedded in the underlying portfolio assets. A well-diversified portfolio would follow a risk parity strategy applied to these uncorrelated risk sources; see Lohre, Opfer and Ország (2014). To construct uncorrelated risk sources, Meucci (2009) suggests using principal component analysis (PCA). Yet, follow-up research by Meucci, Santangelo and Deguest (2015) instead advocates a minimum-torsion transformation to derive the linear orthogonal transformation closest to the original assets (or a pre-specified factor model). Thus, we follow Dichtl, Drobetz, Lohre and Rother (2019) in using the minimum-torsions of the three market and four style factors introduced in figure 1.

The hierarchical structure of the multi-asset multi-factor universe
Portfolio optimization methods like the Markowitz mean-variance approach are sensitive to changes in input variables, and small estimation errors can lead to vast differences in optimal portfolio allocations. However, correlation and covariance matrices are quite complex, and they disregard

the hierarchical structure of asset interactions. To reduce complexity, one wants to focus on relevant correlations only. In this regard, a well-known approach from graph theory is the minimum spanning tree (MST) that connects all entities (here: assets and factors) without cycles but with the minimum total edge weight. An algorithm for obtaining the MST was introduced by Prim (1957). Before applying this algorithm, one has to define a distance measure, which is often based on the correlation coefficient. We will refer to this measure as the dissimilarity measure, since it aims to measure the dissimilarity of the assets (and factors).[1] Applying the dissimilarity measure to the correlation matrix leads to the so-called 'dissimilarity matrix' and allows us to derive the MST (figure 2). When the MST is based on correlations, it is also often referred to as a correlation network.

The MST reflects essential information contained in the correlation matrix and introduces a hierarchical structure. Looking at the branches, one particularly identifies an equity risk-like cluster and a more defensive cluster, among others.

Graph theory is linked to unsupervised machine learning. In particular, the MST is naturally related to the hierarchical clustering algorithm called single linkage. In a direct way, the MST conveys the hierarchical organization of the investigated assets and style factors, which results in a tree structure as represented by the dendrogram in figure 3.[2] Moving up the tree, objects that are similar to each other are combined into branches, i.e. the higher the height of the fusion, the less similar the objects are. Note that one has to define how to use this information for measuring the (dis)similarity among

Figure 3
Dendrogram based on Ward's method

Distance/(dis)similarity
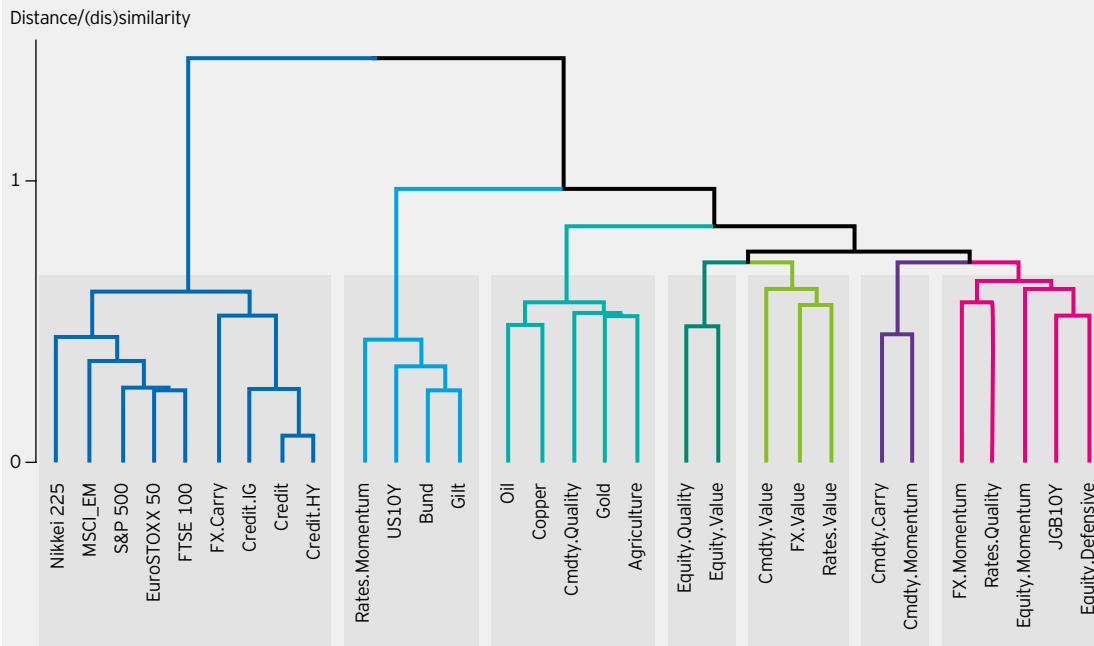
Notes: The figure depicts the dendrogram based on Ward's method for the multi-asset multi-factor universe building on the variance-covariance matrix using monthly data from the full sample period from 31 January 2001 to 31 October 2018.
Sources: Bloomberg, Invesco, Goldman Sachs.

clusters containing more than one element. This is done by the respective linkage criterion; we consider dendrograms based on Ward's method going forward.[3]

In the case of a large investment universe, it might make sense to consider the dendrogram only up to a certain level rather than taking the whole hierarchical structure into account. While this reduction leads to a loss of information, it makes finding the weight allocation faster. Cutting the dendrogram will partition the assets and style factors into clusters. There are different ways to determine an optimal number of clusters. One could simply choose a plausible number by looking at the dendrogram or apply a statistical criterion for determining the "optimal" number of clusters. An example is given in figure 3, where the number of clusters was deliberately chosen to be seven.

Portfolio allocation based on hierarchical clustering
Having determined the dendrogram, one has to decide how to allocate one's capital. Instead of using an algorithm based on recursive bisection as in Lopez de Prado (2016), Lohre, Rother and Schäfer (2020) propose investing along the nodes of the dendrogram to integrate the hierarchical information. Further, one has to choose an allocation technique within and across clusters – Lopez de Prado uses the inverse variance strategy in both cases, but there are various other alternatives. For instance, Papenbrock (2011) and Raffinot (2017) suggest a weighting scheme that allocates capital equally across cluster hierarchy and within clusters. In our study, we use a combination of risk parity based on equal risk contributions. The algorithm of Lohre, Rother and Schäfer is described

Box
Algorithm: Clustering-based weight allocation

1. *Perform hierarchical clustering and generate dendrogram*

2. *Assign all assets a unit weight* $\omega_i = 1 \ \forall \ i = 1,...,N$

3. *For each dendrogram node (beginning from the top):*
   a. *Determine the members of clusters C_1 and C_2 belonging to the two sub-branches of the according dendrogram node*

   b. *Calculate the within-cluster allocations $\widetilde{\omega}_1$ and $\widetilde{\omega}_2$ for C_1 and C_2 according to risk parity (equal risk contributions)*

   c. *Based on the within-cluster allocations $\widetilde{\omega}_1$ and $\widetilde{\omega}_2$ calculate the across-cluster allocation α (splitting factor) for C_1 and $1 - \alpha$ for C_2 according to to risk parity (equal risk contributions)*

   d. *For each asset in C_1 re-scale allocation ω by factor α*

   e. *For each asset in C_2 re-scale allocation ω by factor $1 - \alpha$*

4. *For each cluster containing more than one element:*
   a. *Determine the members of the cluster*

   b. *Calculate the within-cluster allocation*

   c. *For each asset in the cluster re-scale ω by the within-cluster allocation*

5. *End*

Table 1
Performance statistics of multi-asset multi-factor strategies

Performance statistics	1/N	MVP	DRP	HRP	HRP Smooth
Gross return p.a. (%)	4.23	3.95	4.20	3.95	4.43
Net return p.a. (%)	3.96	3.33	3.78	3.22	3.76
Volatility p.a. (%)	3.45	1.38	1.51	1.57	1.63
Sharpe ratio	0.74	1.37	1.54	1.14	1.42
Maximum drawdown (%)	-8.62	-1.20	-1.20	-1.83	-2.03
Calmar ratio	0.46	2.77	3.16	1.77	1.85
Number of bets	3.15	4.35	7.00	6.02	6.17
Turnover (%)	1.55	5.61	3.72	17.22	12.99

Notes: The table provides simulated performance figures for risk-based multi-asset multi-factor strategies from the perspective of a US-dollar investor. This model does not factor in all of the economic and market conditions that can impact results.
Sources: Bloomberg, Invesco, Goldman Sachs. Period: January 2006 to October 2018. **The figures refer to simulated past performance and past performance is not a reliable indicator of future performance.**

in the box. It starts at the top of the dendrogram and assigns weights by going from node to node. Note that step 4 in the algorithm needs only be executed if an optimal number of clusters is used, i.e. not all remaining clusters are singleton clusters.

Hierarchical risk parity for multi-asset multi-factor allocations

In this section, we focus on examining hierarchical risk parity strategies in the multi-asset multi-factor domain vis-à-vis the alternative risk-based allocation strategies 1/N, minimum-variance (MVP) and diversified risk parity (DRP). The traditional risk-based allocation strategies are directly applied to the seven aggregated factors resulting from the imposed aggregate factor model. These seven factors can be viewed as "economic" clusters, providing a benchmark for the "statistical" hierarchical clustering. As for HRP, the allocation strategies used either within or across clusters are risk parity, based on equal risk contributions. For hierarchical clustering, we use Ward's method and the dissimilarity matrices are derived from the correlation matrix.[4]

Portfolio rebalancing is conducted on a monthly basis. The strategies are assumed to be implemented using futures and swaps with associated transaction costs of 10 basis points (futures) and 35 basis points (swaps), respectively. Furthermore, 8 basis points per month are considered for holding a given swap. A 5-year rolling window of monthly returns is used for estimation of the covariance matrix, and the resulting correlation-based dendrograms are updated every month. We perform backtests of the investment strategies from January 2006 to October 2018.[5]

Table 1 shows performance and risk statistics as well as the average strategy turnover. First, we note that the 1/N strategy suffers from the highest volatility, as well as the highest maximum drawdown, rendering its risk-adjusted performance sub par. The underlying lack of diversification is discernible from only 3.15 bets averaged over time. Minimum-variance optimization enables an increase of this number to 4.35. Unsurprisingly, MVP exhibits the lowest portfolio volatilities in the sample period (1.38%). Maximum drawdown figures and risk-adjusted returns are also improved relative to equal weighting.

Next, we examine the middle-ground solution in between 1/N and minimum-variance: diversified risk parity, which is designed to have a maximum of seven bets over time. Its gross return is almost as high as that of 1/N (4.20%) while its turnover is in between the one of 1/N and MVP. As a result, it has the highest net Sharpe ratio.

Having investigated the risk-based strategies for economic factors, we are eager to learn how the approach based on statistical clusters fares. From a volatility perspective, we note that HRP is almost on par with MVP and DRP (1.57%). Also, despite being grounded in statistical clusters, we note that HRP captures 6.02 bets on average. However, the HRP allocation exhibits rather high turnover (17.22%), bringing the net Sharpe ratio down to 1.14. Moreover, the HRP is characterized by a maximum drawdown of -1.83%, which is more severe than the respective figures for MVP (-1.20%) and DRP (-1.20%).

Of course, one would have hoped to enable competing particularly in this statistic when diversifying by statistical clusters. Presumably, the statistical nature of the HRP renders the strategy too active following changes in the correlation structure. As a remedy, we have examined a smoothed HRP variant that is anchored in the optimal HRP portfolio but subject to a transaction cost penalty to smooth the overall allocation and implicitly reduce the associated transaction costs.[6] The last column of table 1 highlights the efficacy of the transaction cost penalty. We observe an increase in returns, yet risk characteristics are hardly affected, rendering it roughly on par with the outcome of the diversified risk parity strategy.

Conclusions

The main motivation to base an allocation strategy on hierarchical clustering is that the correlation matrix is too complex to be fully analyzed and lacks the notion of hierarchy. Hierarchical clustering reduces complexity by focusing on the correlations that really matter. Hierarchical risk parity is an intuitive investment approach, allowing for a high degree of flexibility. Though conceptually appealing, our empirical study suggests that a pure HRP allocation creates substantial turnover when seeking to follow the ensuing dynamic clusters and hierarchy.

We consider a transaction cost penalty to be an effective means to smooth the HRP allocation, rendering its return similar to a diversified risk parity strategy based on economic factors, but not its diversification and downside risk.

Hierarchical clustering reduces complexity by focusing on the correlations that really matter.

References

Bernardi, S., M. Leippold and H. Lohre (2018): Maximum diversification strategies along commodity risk factors, European Financial Management 24(1), pp. 53-78.

Dichtl, H., W. Drobetz, H. Lohre and C. Rother (2019): Active factor completion strategies, Working paper.

Lohre, H., H. Opfer and G. Ország (2014): Diversifying risk parity, Journal of Risk 16(5), pp. 53-79.

Lohre, H., C. Rother and K. Schäfer (2020): Hierarchical risk parity: Accounting for tail dependencies in multi-asset multi-factor allocations, Chapter 9 in: Machine Learning and Asset Management, Emmanuel Jurczenko (ed.), IESTE-Elsevier forthcoming.

Lopez de Prado, M. (2016): Building diversified portfolios that outperform out of sample, Journal of Portfolio Management 42(4), pp. 59-69.

Mantegna, R. (1999): Hierarchical structure in financial markets, European Physical Journal B - Condensed Matter and Complex Systems 11(1), pp. 193-197.

Markowitz, H. (1952): Portfolio selection, Journal of Finance 7(1), pp. 77-91.

Meucci, A. (2009): Managing diversification, Risk 22, pp. 35-40.

Meucci, A., A. Santangelo and R. Deguest (2015): Risk budgeting and diversification based on optimized uncorrelated factors, Risk 11(29), pp. 70-75.

Papenbrock, J. (2011): Asset Clusters and Asset Networks in Financial Risk Management and Portfolio Optimization, PhD thesis, KIT, Karlsruhe.

Prim, R. C. (1957): Shortest connection networks and some generalizations, Bell System Technical Journal 36(6), pp. 1389-1401.

Raffinot, T. (2017): Hierarchical clustering-based asset allocation, Journal of Portfolio Management 44(2), pp. 89-99.

Tumminello, M., F. Lilio and R. N. Mantegna (2010): Correlation, hierarchies, and networks in financial markets, Journal of Economic Behavior and Organization 75(1), pp. 40-58.

Appendix

Here, we briefly describe the single asset and style factor indices underlying the article's empirical analyses. The global equity and bond markets are represented by equity index futures for S&P 500, Nikkei 225, FTSE 100, EuroSTOXX 50, MSCI Emerging Markets and bond index futures for 10-year US Treasuries, German Bunds, 10-year JGBs and Gilts. The credit risk premium is captured by the Bloomberg Barclays US Corporate Investment Grade (Credit IG) and High Yield (Credit HY) indices (both duration-hedged to synthesize pure credit risk). To capture commodity markets, we consider total return indices of S&P GSCI for crude oil and gold as well as total return indices from Bloomberg for copper and agriculture.

All style factors are constructed in a long-short fashion and all non-equity style factors are sourced from Goldman Sachs (GS); see table 2 for the style factor indices used. For equity style factors, we utilize the Invesco Quantitative Strategies definitions as laid out in "Investing in a multi-asset multi-factor world", Risk & Reward, #3/2017. In particular, equity value, momentum and quality each follow a multi-factor approach that combines several metrics proxying for the respective style dimension. For equity defensive, we build on a long-short approach that is long a minimum-volatility portfolio while shorting a beta-adjusted market portfolio.

Table 2: **Overview of style factor series**

Style factor	Equity	Fixed Income	Commodity	FX
Carry	–	GS Interest Rates Carry 05	GS Macro Carry Index RP14	GS FX Carry C0115
Value	IQS Value	GS Interest Rates Value 05	GS Commodity COT Strategy COT3	GS FX Value C0114
Momentum	IQS Momentum	GS Interest Rates Trend	GS Macro Momentum Index RP15	GS FX Trend C0038
Quality	IQS Quality & IQS Defensive	GS Interest Rates Curve C0210	GS Commodity Curve RP09	–

About the authors

Dr. Martin Kolrep
Senior Portfolio Manager,
Invesco Quantitative Strategies
Dr. Martin Kolrep is involved in the development and
management of factor-based multi-asset strategies
mainly for institutional clients.

Dr. Harald Lohre
Senior Research Analyst,
Invesco Quantitative Strategies and Visiting Research
Fellow EMP/Lancaster University
Dr. Harald Lohre develops quantitative risk and
return forecast models used in the management of
multi-asset multi-factor strategies.

Erhard Radatz
Portfolio Manager,
Invesco Quantitative Strategies
Erhard Radatz manages multi-asset portfolios that
include the elements: factor-based investing, active
asset allocation and downside risk management.

Carsten Rother
Research Analyst,
Invesco Quantitative Strategies
Carsten Rother develops quantitative multi-factor
models in equities and across asset classes while
pursuing his doctoral studies at the University of
Hamburg.

Integrating low volatility style exposure into core equity factor investments

By Michael Fraikin, Xavier Gerard, Ph.D., and André Roberts

In brief

Low volatility style investing is popular, but not easy to implement in a core equity strategy. We show some possible routes to this goal in the light of realistic investment constraints. The aim is to create an investable portfolio that harvests the low volatility factor alongside quality, value and momentum.

Firms with historically low stock return volatility have been found to offer a higher risk-adjusted return than their high risk counterparts. The empirical evidence for this effect has survived decades of scrutiny, but some authors have recently cast doubt on the robustness of the empirical results presented in some studies, suggesting that they may vanish in the light of realistic investment constraints. While we cannot speak for all conceivable low volatility strategies, we challenge the generality of this position and show how to create an investable portfolio designed to capture the low volatility anomaly that offers an important contribution to the factor toolkit.

One of the most enduring observations in the financial economics literature is that, on a risk adjusted basis, low risk assets tend to outperform their high risk counterparts (Haugen and Heins, 1975). Earlier explanations for this effect have focused on limits to arbitrage, such as investors' leverage constraints (Black, 1972), which imply a lower risk-adjusted return for high-beta stocks than for low-beta ones. Somewhat related explanations are provided by delegated-agency models (Baker, Bradley and Wurgler, 2011), where most investment managers are benchmark-constrained and try to achieve higher information ratios by holding higher volatility stocks.

Behavioural explanations have been advocated as well: some investors' preference for stocks with lottery-like payoffs (Kahneman and Tversky, 1979); the representativeness heuristic (Kahneman and Tversky, 1983), which could explain why people

would overpay for speculative investments; and overconfidence (Cornell, 2009), which is particularly problematic when outcomes are uncertain (Diether, Malloy and Scherbina, 2002).

Novy-Marx and Velikov (2018) have cast serious doubts on the robustness of at least one low volatility strategy.

Recently, however, Novy-Marx and Velikov (2018) have cast serious doubts on the robustness of at least one low volatility strategy that is particularly popular in academia: the Betting-Against-Beta (BAB) strategy of Frazzini and Pedersen (2014). In short, the authors argue that the staggering performance of BAB is the result of ad-hoc methodological choices that boil down to a huge overweight of very small stocks. After accounting for transaction costs, they show that the strategy is explained by common factors. The authors are also highly critical of the estimation of stock betas, which strikes at the core of the strategy. These concerns are shared by Welch (2019).

In light of this, we aim to show that it would be inappropriate to generalize and infer that these specific criticisms apply to all low volatility strategies. Using a global universe of developed market stocks from January 1997 to March 2019, we show that an investable low volatility strategy can be devised to successfully capture the low volatility anomaly and that its performance compares remarkably well to that of a set of equally investable common factors.

Constructing an investable low volatility strategy
To capture the low volatility anomaly, we define two investable portfolios: a lower risk and a higher risk portfolio. Both portfolios are adjusted monthly.

Our higher risk portfolio is simply the market portfolio. While a portfolio of more volatile securities could theoretically be used, the benefit of shorting them does not appear to offset the costs involved. In contrast, futures on the market portfolio are widely available and can be shorted at little cost.

Our low risk portfolio is a long-only minimum variance portfolio with stringent investment constraints. The position sizes depend on the estimated daily trading volume of the individual stocks, so that we effectively forbid investments in the 30% of stocks expected to be the least liquid in a given month. The one-way turnover of the portfolio is also kept to a minimum, at 30% per year. Finally,

Betting-Against-Beta (BAB) strategy
The BAB strategy dynamically leverages up a long portfolio of low risk stocks and deleverages a short portfolio of high risk stocks such that each side achieves the same beta of 1.

we limit country, region, industry and sector exposures to +/-10 percentage points relative to the market portfolio and enforce diversification by capping holdings at 1%.

Key to capturing the low volatility anomaly is the understanding that low risk assets are expected to outperform on a risk-adjusted basis. In other words, in the following regression:

$$r_{MinVar} - r_f = \alpha + \beta \cdot (r_{Mkt} - r_f) + \varepsilon$$

where r_{MinVar} is the return of the investable minimum variance portfolio; r_{Mkt} is the return of the market portfolio; r_f the risk-free rate and α the abnormal return earned by the low volatility anomaly, we expect α to be positive and significant.

Crucially, this is not the same as saying that

$$r_{MinVar} > r_{Mkt} .$$

To earn the low volatility premium, we need to risk-adjust the performance of the market portfolio. Our low volatility strategy does so by going 100% long in the minimum variance portfolio and β short in the market portfolio, with the remaining $1-\beta$ borrowed at the risk-free rate. In mathematical terms, we have:

$$r_{Low\text{-}Volatility} = (r_{MinVar} - \beta \cdot r_{Mkt}) - (1-\beta) \cdot r_f$$

For each month, we compute β using a covariance matrix of stock returns (estimated with our proprietary risk model) and the portfolio holdings of the minimum variance and market portfolios, namely:

$$\beta = (h_{Mkt}' \cdot \Omega \cdot h_{Mkt})^{-1} \cdot (h_{Mkt}' \cdot \Omega \cdot h_{MinVar})$$

This strategy is reminiscent of the BAB strategy (Frazzini and Pedersen, 2014), which measures risk with the beta of individual stocks and captures the following return:

$$r_{BAB} = \frac{1}{\beta_{LowVol}} \cdot r_{LowVol} - \frac{1}{\beta_{HighVol}} \cdot r_{HighVol} - \left(\frac{1}{\beta_{LowVol}} - \frac{1}{\beta_{HighVol}}\right) \cdot r_f$$

The key difference is that our strategy addresses the investability concerns raised by Novy-Marx and Velikov (2018).

In addition, our beta-neutral strategy is a mean-variance efficient allocation between the minimum variance portfolio and the market portfolio. In case this is not immediately obvious, we compare our proposed weights to those obtained by solving the following mean-variance optimization problem:

$$\max_{w} w' \cdot r - \frac{\lambda}{2} \cdot w' \cdot \Omega \cdot w$$

$$s.t. \ w' \cdot \begin{pmatrix} \beta \\ 1 \end{pmatrix} = 0$$

where r is a 2×1 column-vector of expected excess returns:

$$r = \begin{pmatrix} r_{MinVar} - r_f \\ r_{Mkt} - r_f \end{pmatrix};$$

Ω is the covariance matrix of excess returns:

$$\Omega = \begin{pmatrix} \sigma^2_{MinVar-rf} & cov \\ cov & \sigma^2_{Mkt-rf} \end{pmatrix}$$

λ is the risk aversion parameter; and β is the market beta of the minimum variance portfolio.

The solution to this problem is well-known:

$$w = \frac{1}{\lambda} \cdot \Omega^{-1} \cdot \left\{ I - \begin{pmatrix} \beta \\ 1 \end{pmatrix} \cdot \left[\begin{pmatrix} \beta \\ 1 \end{pmatrix}' \cdot \Omega^{-1} \cdot \begin{pmatrix} \beta \\ 1 \end{pmatrix} \right]^{-1} \cdot \begin{pmatrix} \beta \\ 1 \end{pmatrix}' \cdot \Omega^{-1} \right\} \cdot r$$

where I is the identity matrix.

After some tedious, albeit straightforward, algebra we find:

$$w = \frac{1}{\lambda \cdot \left(\sigma^2_{MV-rf} \cdot \sigma^2_{Mkt-rf} - cov^2 \right)} \cdot$$

$$\left\{ \begin{matrix} \sigma^2_{Mkt-rf} \cdot \left[\left(r_{MinVar} - r_f \right) - \beta \cdot \left(r_{Mkt} - r_f \right) \right] \\ -cov \cdot \left[\left(r_{MinVar} - r_f \right) - \beta \cdot \left(r_{Mkt} - r_f \right) \right] \end{matrix} \right\}$$

Recognising that

$$\alpha = \left(r_{MinVar} - r_f \right) - \beta \cdot \left(r_{Mkt} - r_f \right),$$

we can restate as follows:

$$w = \frac{1}{\lambda \cdot \left(\sigma^2_{MV-rf} \cdot \sigma^2_{Mkt-rf} - cov^2 \right)} \cdot \left\{ \begin{matrix} \sigma^2_{Mkt-rf} \cdot \alpha \\ -cov \cdot \alpha \end{matrix} \right\}$$

We choose λ to ensure that the weight of the minimum variance portfolio is equal to 1. Replacing λ (in the equation for the market weight) with this value, we obtain:

$$w_{Mkt} = \frac{\left(\sigma^2_{MV-rf} \cdot \sigma^2_{Mkt-rf} - cov^2 \right)}{\left(\sigma^2_{MV-rf} \cdot \sigma^2_{Mkt-rf} - cov^2 \right)} \cdot \frac{-cov \cdot \alpha}{\sigma^2_{Mkt-rf} \cdot \alpha} = -\beta$$

Accordingly, in line with the original formula, the solution to the optimization problem holds that our beta-neutral and mean-variance efficient allocation between the minimum variance portfolio and the market portfolio can be written as:

$$r_{Low\text{-}Volatility} = \left(r_{MinVar} - \beta \cdot r_{Mkt} \right) - \left(1 - \beta \right) \cdot r_f$$

Finally, in keeping with our objective of constructing a strategy that provides a realistic picture of the premium that can be earned by exploiting the low volatility anomaly, we adjust the returns of the minimum variance and market portfolios using a conservative one-way transaction cost of 75bp.

Constructing a set of investable common factors
In addition to testing the stand-alone performance of the low volatility strategy, we investigate whether it adds alpha over and above a strategy using the factors: quality, value and momentum. This set of factors is often used in the asset pricing literature and all three have withstood proper scrutiny. By that we mean that the key papers that established these factors have been widely replicated and their motivations highly debated so that, even when

several explanations are advocated for these effects, the consensus is that they all seem credible enough. For our purposes, we define the three factors as follows:

Quality – good quality firms with more conservative accounting practices, which are able to generate a healthy amount of cash from operations and that distribute this cash back to investors, have been shown to outperform over the long term. We therefore compute a quality factor that combines 'cash-based-operating profitability' – replaced by 'operating income for financials' (Ball, Gerakos, Linnainmaa and Nikolaev, 2016), 'change in net operating assets' (Richardson, Sloan, Soliman and Tuna, 2005) and the extent of 'share issuances and buybacks' (Pontiff and Woodgate, 2008).

Momentum – firms that have experienced a recent increase in their earnings and stock returns also tend to outperform over the medium term. This is known as the 'momentum anomaly', which we harvest with a simple 'price momentum' metric (Jegadeesh and Titman, 1993) and the 'revision of analysts' earnings forecasts' (Chan, Jegadeesh and Lakonishok, 1996).

Value – there is substantial empirical evidence that value firms with higher ratios of fundamentals-to-price earn higher returns in the long run. Our value factor comprises 'book-to-market', 'operating cashflow yield' (replaced by historical 'earnings yield for financials') and 'forward-earnings yield' (among others: Basu, 1983; Fama and French, 1992).

For the sake of comparability, each signal in a factor family is normalized every month. The signals are then equally weighted, and we form a 100% long and 100% short factor portfolio by holding the top and bottom 20% of stocks in proportion to their average score values.

It is common practice in academia to abstract from real-world investment constraints. However, we are chiefly concerned with findings that can be acted upon. In other words, when assessing the performance of the investable low volatility strategy, we ultimately want to know whether it adds value over and above investable versions of the raw factors described above.

To create investable monthly quality, value and momentum portfolios, we minimize the tracking error to the returns of the above-described targeted factors while incorporating real-world investment constraints.

First, we impose a number of risk controls when constructing our factors each month. We remove the influence of industry and country effects. The returns of value and momentum being often negatively correlated, we adjust the value portfolio so that its returns have zero expected sensitivity to those of momentum. The reciprocal adjustment is performed on the momentum portfolio. The returns of quality can also be sensitive to those of book-to-market, and so we adjust the quality portfolio in a similar fashion (Novy-Marx, 2013).

Second, we force rather stringent investment constraints in the optimization process. As with the construction of the minimum variance portfolio,

holdings are scaled by estimates of each stock's daily trading volume and cannot be larger than +/- 2%. The liquidity constraint implies that these portfolios do not invest in the 30% least liquid stocks in the universe. Still further, turnover is severely restricted with realized values on average equal to 20% of their unconstrained levels.

Finally, all return series are adjusted for costs beyond merely price impact. In international markets, exchange fees and taxes can be steep. Borrowing costs for shorting are frequently high. Structures involving swaps and leverage are costly. In essence, this implies a conservative one-way transaction cost of 1%.

Empirical analysis
We carry our analysis over a global developed country universe with approx. 4500 stocks each month and investigate performance from January 1997 to March 2019.

Table 1 shows the performance of the market portfolio and each individual strategy, as well as that of a multi-factor portfolio (QMV) defined as an equally weighted combination of quality, value and momentum, rescaled to be 100% long and 100% short. The first important observation is that the low volatility strategy performs better than most common factors. The only factor with a higher t-statistic is quality. But many of the component signals in this factor were discovered over the period of our analysis, so that its performance is largely in-sample and should be looked at with a healthy degree of skepticism.

Turning to figure 1, where we plot the cumulative net returns of each strategy, we find that the performance of low volatility went through two periods of almost unabated positive returns. The first one lasted from the tech bubble burst until the financial crisis and was only derailed by a brief yet sharp episode of underperformance at the end of

Table 1
Performance diagnostics

	Return (annualized, %)	Standard deviation (annualized, %)	t-Stat	Maximum drawdown (%)	Turnover one-way (%)
MSCI World Hedged	5.39	14.09	1.81	55	10
Quality	2.72	3.58	3.59	8	89
Momentum	2.29	5.21	2.08	21	176
Value	2.49	3.97	2.96	13	116
QMV	3.22	4.29	3.54	8	112
Low Volatility	3.60	4.77	3.56	16	28

Source: Invesco calculations. MSCI World Hedge is actual performance. The selected factor performance is back-tested data. **Past performance (actual or back-tested) is not a guide to future results.** Actual performance may vary significantly from any hypothetical or historical performance shown. Back-tested performance is not actual performance, but is hypothetical and based on criteria applied retroactively with the benefit of hindsight and knowledge of factors that may have positively affected the performance, and cannot account for all financial risks that may affect the actual performance.

Figure 1
Cumulative net return

Source: Invesco calculations. Data as at February 2019. **Back-tested past performance is not a guide to future results.** Actual performance may vary significantly from any hypothetical or historical performance shown. Back-tested performance is not actual performance, but is hypothetical and based on criteria applied retroactively with the benefit of hindsight and knowledge of factors that may have positively affected the performance, and cannot account for all financial risks that may affect the actual performance.

Table 2
Regression analysis

	Regression 1	Regression 2	Regression 3	Regression 4	Regression 5
Annualized alpha (%)	3.82	3.70	3.39	4.16	3.51
t-Stat	3.77	3.53	3.27	4.07	3.45
Market-rf	-0.04	-0.04		-0.07	
t-Stat	-1.92	-1.69		-3.14	
QMV		0.03	0.07		
t-Stat		0.44	0.99		
Quality				-0.32	-0.22
t-Stat				-3.53	-2.55
Momentum				-0.04	0.00
t-Stat				-0.62	0.07
Value				0.31	0.27
t-Stat				3.77	3.27
R-squared	0.014	0.014	0.004	0.091	0.046

Source: Invesco calculations.

Table 3
Optimal weighting schemes and information ratios

	Vector 1	Vector 2
Quality (%)	48	
Momentum (%)	0	
Value (%)	15	
QMV (%)		53
Low Volatility (%)	37	47
Information ratio	1.15	1.03

Source: Invesco calculations.

2002. The second period of sustained outperformance started in mid-2009 and lasted until mid-2016. Out of all the other periods, it is during the tech bubble that low volatility witnessed its worst drawdown. The same is true for value, and it is interesting to note that the two strategies share similar dynamics, not just over this period but also in its immediate aftermath: the burst of the dot.com bubble and the ensuing recession.

Nevertheless, despite at times sharing the dynamics of the common factors, figure 1 shows that the dynamics of the low volatility strategy display unique characteristics that should provide a valuable enhancement to existing factors. We explore this formally in table 2, where the net returns of the low volatility strategy are regressed against the market and the three common factors. Regression 1, where we only control for the excess return over the market, shows that the ex-ante beta neutrality of the low volatility strategy is mostly preserved ex-post. While the strategy inherits a negative ex-post beta, the latter is small and hardly significant. In fact, the proportion of the strategy variance explained by its exposure to the market is just over one percent.

Irrespective of whether we add all common factors separately (regression 4) or in combination within QMV (regression 2), we always find that the low volatility strategy earns a significant annualized abnormal return of approximately 4%. This indicates that the strategy is not spanned by common factors and the market, so that an optimal combination of these with low volatility would improve mean-variance efficiency. While these findings are interesting, the results in regressions 3 and 5, where we exclude the excess market return, are arguably of greater practical importance. This is because one would typically entertain combining the low volatility strategy with quality, value and momentum. Again, the large and statistically significant alphas in these regressions suggest that an optimal combination of

Table 4
Performance diagnostics

	Return (annualized, %)	Standard deviation (annualized, %)	Information ratio	t-Stat	Maximum drawdown (%)	Turnover one-way (%)
QMV	3.22	4.29	0.75	3.54	8	112
QMVL	3.32	3.50	0.95	4.47	7	91

Source: Invesco calculations. **Back-tested past performance is not a guide to future results.** Actual performance may vary significantly from any hypothetical or historical performance shown. Back-tested performance is not actual performance, but is hypothetical and based on criteria applied retroactively with the benefit of hindsight and knowledge of factors that may have positively affected the performance, and cannot account for all financial risks that may affect the actual performance.

low volatility with the set of common factors would improve mean-variance efficiency.

Next, table 3 displays the values of these optimal weights. For ease of interpretation, we have scaled them such that their sum is equal to one. Relying solely on historical data to select the optimal combination of quality, value, momentum and low volatility (vector 1), one would place a very large weight of 48% on quality, no weight on momentum and a moderate weight on low volatility (37%), which is more than twice as large as value (15%). Vector 2 returns the optimal weighting scheme where quality, value and momentum are equally weighted within QMV. Over our study period, it would have been optimal to assign almost equal weight to QMV (53%) and low volatility (47%). It should come as no surprise that the information ratio (IR) of vector 1 is larger than that of vector 2 since the latter imposes equal weights on quality, value and momentum. However, there is only a ten percent difference between the resulting portfolios' IRs, and their returns display a correlation of 0.9, suggesting that these two factor-combinations lead to portfolios that are not too dissimilar.

Nevertheless, vectors 1 and 2 are both relatively unbalanced, so that we evaluate the performance (table 4) of a portfolio (QMVL) with a weight of 75% on QMV and 25% on the low volatility strategy. Even though this is not the weighting scheme that would have maximized mean-variance efficiency of the combined portfolio, we do find a meaningful improvement over the performance of QMV. The volatility of the strategy and (to a lesser extent) its maximum drawdown both decrease while its return increases – albeit only marginally – leading to a markedly higher IR for QMVL than QMV.

Integrating the low volatility anomaly in a long-only portfolio
As indicated from the outset, the low volatility anomaly refers to the fact that, on a risk-adjusted basis, lower risk stocks dominate their higher risk counterparts. This implies that some leverage is needed to capture the low volatility anomaly, i.e. the benefit of incorporating it into a long-only strategy may not seem immediately obvious.

However, things become fairly straightforward once we recognize that the choice of implementing a low volatility bet within a long-only strategy involves two distinct active decisions: (1) an asset allocation decision, which refers to the choice of reducing risk by deviating from a beta of one to the market and

> Things become fairly straightforward once we recognize that the choice of implementing a low volatility bet within a long-only strategy involves two distinct active decisions.

(2) a decision to enhance the return of the targeted asset allocation by harvesting the low volatility anomaly.

Taking an investment in a minimum variance portfolio as an example, we first define the low volatility anomaly as the difference between the minimum variance portfolio and the risk-adjusted market portfolio, with the net investment value borrowed at the risk-free rate:

$$r_{Low\text{-}Volatility} = \left(r_{MinVar} - \beta \cdot r_{Mkt}\right) - \left(1 - \beta\right) \cdot r_f$$

We then interpret the decision to deviate from a beta of one to the market as an asset allocation decision:

$$r_{Asset\ Allocation} = \beta \cdot r_{Mkt} + \left(1 - \beta\right) \cdot r_f$$

In turn, the performance of the investable long-only minimum variance portfolio can be decomposed as follows:

$$r_{MinVar} = r_{Low\text{-}Volatility} + r_{Asset\ Allocation}$$

In practice, many investors are unwilling to deviate entirely from the market portfolio towards the minimum variance portfolio. They often prefer to target some moderate risk reduction. This can be readily done by combining the market portfolio with an investable long-only minimum variance portfolio and choosing the weights between them so as to achieve the desired risk target.

Let us define the following composite benchmark with weights (w) chosen to achieve a desired risk reduction:

$$r_{Composite} = w \cdot r_{Mkt} + \left(1 - w\right) \cdot r_{MinVar}$$

Replacing in this equation the minimum variance portfolio with its expression derived above and re-arranging terms, we have:

$$r_{Composite} = \left(1 - w\right) \cdot r_{Low\text{-}Volatility} + \left\{\left(1 - w\right) \cdot \beta + w\right\} \cdot r_{Mkt} + \left\{\left(1 - w\right) \cdot \left(1 - \beta\right)\right\} \cdot r_f$$

Introducing $\omega = \left\{\left(1 - w\right) \cdot \beta + w\right\}$, the asset allocation decision becomes:

$$r_{Asset\ Allocation} = \omega \cdot r_{Mkt} + \left(1 - \omega\right) \cdot r_f$$

In turn, we have the following decomposition for the performance of the composite benchmark:

$$r_{Composite} = \left(1 - w\right) \cdot r_{Low\text{-}Volatility} + r_{Asset\ Allocation}$$

One could then entertain deviating further from the low volatility benchmark by adding exposures to other rewarded factors, including value, momentum and quality. We argue that this must be done in a risk-controlled manner. The intuition for this is simple and can be best described using the long-only minimum variance portfolio as an example. Specifically, the main difficulty when adding factor exposures to a long-only minimum variance portfolio is avoiding a large exposure to higher volatility stocks, which could otherwise significantly detract from the risk reduction objective.

Table 5
Performance decomposition

Panel A: Decomposition of net returns

	Enhanced Low Volatility (ELV)	Market	Active (ex market) ELV	Total return asset allocation	Active return (ex market) asset allocation	Low volatility anomaly	QMV
Return (annualized, %)	10.09	7.42	2.66	6.63	-0.80	1.81	1.65
Standard deviation (annualized, %)	12.80	15.10	3.92	12.62	2.49	2.51	1.61
Sharpe ratio	0.59	0.32	–	0.32	–	–	–
Information ratio	–	–	0.68	–	-0.32	0.72	1.03
t-Stat	3.73	2.33	3.21	2.49	-1.52	3.41	4.87

Panel B: Decomposition of active risk

	Full covariance matrix	Diagonal covariance matrix
Asset allocation (%)	43	41
Low volatility anomaly (%)	40	42
QMV (%)	18	17

Source: Invesco calculations. **Back-tested past performance is not a guide to future results.** Actual performance may vary significantly from any hypothetical or historical performance shown. Back-tested performance is not actual performance, but is hypothetical and based on criteria applied retroactively with the benefit of hindsight and knowledge of factors that may have positively affected the performance, and cannot account for all financial risks that may affect the actual performance.

To put it simply, when adding some factor exposures to the long-only minimum variance portfolio:

- While one can get positive exposures to stocks with good factor scores across the universe;

- One can only underweight stocks with negative factor scores that also have low risk.

- In turn, if the trade-off between the expected active return and the active risk from these deviations is not adequately controlled for, the active bets tend to come with large exposures to higher volatility stocks – potentially defeating the original objective of reducing risk.

Therefore, to add a multi-factor equity strategy that captures quality, value and momentum (QMV), we create an optimized portfolio relative to the low risk benchmark, where the active component is subject to restrictions on liquidity, turnover, single stocks, sector, industry and country active bets; and crucially where tracking error is explicitly controlled for.

To illustrate our approach, we have tested, in the same global universe of developed countries, a monthly rebalanced strategy that uses a composite benchmark with a weight of 45% on the investable long-only minimum variance portfolio and 55% on the market portfolio and where the target tracking error for QMV relative to this composite benchmark is 1.5%. Some key results for this backtest are presented in table 5, where we assume a conservative one-way transaction cost of 75bp.

Our structured approach to portfolio construction allows us to precisely decompose performance according to each active decision. For instance, our asset allocation decision reduces the risk of the portfolio relative to that of the market by just over 15%. While our exposures to the low volatility anomaly and QMV have little impact on the overall risk reduction objective, they add 3.46% of net active return, which more than compensates for the lowered market exposure (-80bp). In turn, our enhanced low volatility strategy achieves a much higher information ratio than the market (0.59 vs. 0.32) thanks to significant improvements from both the return and risk sides of the portfolio.

Conclusion
In light of recent concerns that the low volatility anomaly may not be robust to the imposition of realistic investment constraints, we developed an investable low volatility strategy and showed that it would have significantly improved the mean-variance efficiency of a portfolio that only combines investable value, momentum and quality factors. Importantly, these results hold not only when looking at long/short strategies but also when imposing a long-only constraint. All in all, our findings largely comport with our view that the low volatility anomaly should be part of a balanced factor allocation alongside other traditional "quant" factors.

References
Baker, M., B. Bradley and J. Wurgler (2011): Benchmarks as limits to arbitrage: understanding the Low volatility anomaly, Financial Analyst Journal 67 (January), pp. 1-15.

Ball, R., J. Gerakos, J. T. Linnainmaa and V. Nikolaev (2016): Accruals, cash flows, and operating profitability in the cross section of stock returns, Journal of Financial Economics 121 (July), pp. 28-45.

Basu, S. (1983): The relationship between earnings yield, market value, and return for NYSE common stocks: Further evidence, Journal of Financial Economics 12 (June), pp. 129-156.

Black, F. (1972): Capital market equilibrium with restricted borrowing, Journal of Business 4 (July), pp. 444–455.

Chan, L. K. C., N. Jegadeesh, and J. Lakonishok (1996): Momentum strategies, Journal of Finance 51 (December), pp. 1681–1713.

Cornell, B. (2009): The Pricing of volatility and skewness: A new interpretation, Journal of Investing 18 (Fall), pp. 27–30.

Diether, K. B., C. J. Malloy, and A. Scherbina (2002): Differences of opinion and the cross section of stock returns, Journal of Finance 57 (October), pp. 2113–2141.

Fama, E. F., and K. R. French (1992): The cross-section of expected stock returns, Journal of Finance 47 (June), pp. 427–465.

Frazzini, A., and L. Pedersen (2014): Betting against beta, Journal of Financial Economics 111 (January), pp. 1–25.

Haugen, R. A., and A.J. Heins (1975): Risk and the rate of return on financial assets: Some old wine in new bottles, Journal of Financial and Quantitative Analysis 10 (December), pp. 775-784.

Jegadeesh N., and S. Titman (1993): Returns to buying winners and selling losers: Implications for stock market efficiency, Journal of Finance 48 (March), pp. 65-91.

Kahneman, D., and A. Tversky (1979): Prospect theory: An analysis of decision under risk, Econometrica 47 (March), pp. 263–291.

Kahneman, D., and A. Tversky (1983): Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment, Psychological Review 90 (October), pp. 293–315.

Novy-Marx, R. (2013): The other side of value: The gross profitability premium, Journal of Financial Economics 108 (April), pp. 1-28.

Novy-Marx, R., and M. Velikov. 2018): Betting Against Betting Against Beta, Working paper.

Pontiff J., and A. Woodgate (2008): Share issuance and cross-sectional returns, Journal of Finance 63 (April), pp. 921-945.

Richardson, S. A., R. G. Sloan, M. T. Soliman, and I. Tuna (2005): Accrual reliability, earnings persistence and stock prices, Journal of Accounting and Economics 39 (September), pp. 437-485.

Welch, I. (2019): Simple Better Market Betas, Working paper.

About the authors

Michael Fraikin
Global Head of Research,
Invesco Quantitative Strategies
Michael Fraikin and his team are responsible for the maintenance and further development of the quantitative models that drive the decisions within the investment products.

Xavier Gerard, Ph.D., CFA
Senior Research Analyst,
Invesco Quantitative Strategies
Xavier Gerard develops quantitative models to forecast risk and return used in the management of equity strategies.

André Roberts, CFA
Senior Portfolio Manager,
Invesco Quantitative Strategies
André Roberts is responsible for managing Australian equity portfolios.

A forecast combination approach to equity factor timing

By Michael Fraikin, Edward Leung, Ph.D., and Dr. Harald Lohre

In brief
We investigate the benefits of forecast combination for timing equity factors based on predictive regressions using macro predictors. Relative to standard predictive regression models, forecast combination reduces the noise of forecasts and hence improves their out-of-sample predictive accuracy. Given the nature of macro predictors, the ensuing dynamic model reacts when major macro events happen. Before transaction costs, portfolio simulation results show considerable outperformance of the factor timing model over a static factor allocation. But much of this performance wedge is eroded when transaction costs are taken into account, rendering this article a cautionary tale about the benefits of factor timing.

Factor investing has become popular among academics and asset managers since the global financial crisis in 2007-2009. And though its value is now widely accepted, it is still a matter of debate whether factor timing can add value over a diversified static factor allocation. In this article, we investigate the benefits of forecast combination for timing equity factors based on predictive regressions and using time series macro predictors as inputs.

In testing for the potential benefits of equity factor timing through forecast combination, we follow five steps: (1) running predictive regressions for each of four equity factors, (2) combining the individual factor forecasts, (3) translating the combined equity factor forecasts into dynamic factor weights, (4) constructing a dynamic multi-factor model from the dynamic factor weights, (5) examining the performance of the dynamic model over a static factor allocation.

Running predictive regressions for each of the four equity factors
Our starting point lies in predictive regressions based on monthly US large cap data from April 1991 to December 2017.[1] The dependent variable is one-month forward looking return (proxied by decile spreads) of broad equity factors, such as price trend, relative value, earnings momentum and quality. Price trend (PT) considers various price momentum factors, such as specific and risk-adjusted momentum. Earnings momentum (EM) subsumes various earnings-related metrics, such as earnings revision, sales revision and earnings surprise. Relative value (RV) is based on value factors constructed using data from various portions of the financial statements, such as book yield, earnings yield and gross profit

yield. Finally, quality (Q) is based on various quality and efficiency factors, such as external financing and return on equity.

The independent variables used to forecast the above four equity factors are 30 standard macro predictors taken from the Fred-MD database by McCracken and Ng (2016).[2] This set of macro predictors derives from various categories, such as output and income, labour market, housing, consumption, orders and inventories, money and credit, interest and exchange rates, etc. We lag all predictors by two months to adjust for reporting lags.[3]

Combining the individual factor forecasts through forecast combination

In order to improve return forecasts of predictive regressions, we apply forecast combination. First, we consider averaging the forecasts of many predictive regressions, striving to improve the reliability of the dynamic factor weights.[4] We also test machine learning tools, such as Adaptive LASSO and Dynamic Model Averaging, as more advanced ways of averaging the forecasts.

A good forecast is "close" to the actual outcome. To quantify this closeness, we require a metric that defines forecasting mistakes. One common metric is the mean squared forecast error (MSFE), defined as the average sum of squared differences between the forecasts and the actual outcomes. Based on the MSFE, Campbell and Thompson (2008) define an out-of-sample R^2 as follows:

$$R^2_{OOS} = 1 - \frac{MSFE_{predictors}}{MSFE_{long\text{-}term\ average}}$$

If the out-of-sample R^2 is > 0, the forecast based on predictors is more accurate than a naïve forecast that assumes the outcome to be the long-term average.

Based on the 30 predictors, we now calculate the out-of-sample R^2 for different methods of aggregating the forecasts. Table 1 summarizes the results for each of our four equity factors.

The first two columns are included for information and show the average in- and out-of-sample R^2 of the individual predictive regressions. Whereas all in-sample R² are positive, most out-of-sample R^2 values are negative and statistically insignificant. The third column shows the in-sample R^2, the fourth the out-of-sample R^2 of a simple average of the 30 individual forecasts. For three of our four factors,

the out-of-sample R^2 is positive and statistically significant, at least at the 10% level.[5]

In the fifth and sixth columns, we use machine learning tools like Adaptive LASSO (ALASSO)[6] and Dynamic Model Averaging (DMA)[7] as more advanced ways of averaging the forecasts. However, the out-of-sample R² values of both approaches are statistically insignificant, even though they are similar in magnitude to (or even higher than) those derived from simple averaging. This observation can be explained by the nature of the two machine learning techniques: at every point in time, ALASSO will pick a subset of the macro predictors as independent variables, and DMA will average the forecasts of a subset of state space models. These variable and model selection processes generate a higher variance relative to the average forecast across all the macro predictors. Hence, the out-of-sample R² values of DMA and ALASSO are statistically insignificant. Note that none of the approaches works for predicting the quality factor.

> Taking the simple average of all 30 forecasts for each factor delivers the most consistent out-of-sample results.

In summary, taking the simple average of all 30 forecasts for each factor delivers the most consistent out-of-sample results. Given these findings, we will focus on simple averaging for the remainder of the article.

Translating the combined equity factor forecasts into dynamic factor weights

The next step is to construct the dynamic factor weights based on the combination forecast of each of the four equity factors. Let $F(X,t)$ be the forecast of factor X at time t. Then, we set the weight of factor X at time t to:

$$W(X,t) = \frac{F(X,t)+c}{\sum\limits_{X \in \{RV, PT, EM, Q\}} \left(F(X,t)+c\right)} = \frac{F(X,t)+c}{\sum\limits_{X \in \{RV, PT, EM, Q\}} F(X,t) + 4c}$$

Table 1

In-sample and out-of-sample R² for different aggregation methods

	Individual regressions (average): In-sample	Out-of-sample	Simple averaging: In-sample	Out-of-sample	DMA out-of-sample	ALASSO out-of-sample
Price Trend	0.008	0.004	0.12	0.020**	0.029	0.034
Earnings Momentum	0.012	-0.00021	0.23	0.024*	0.012	-0.07
Relative Value	0.009	-0.002	0.10	0.025**	0.067	-0.04
Quality	0.003	-0.020	-0.06	-0.007	-0.009	-0.16

* means out-of-sample R² is statistically significant at the 10% level. ** means out-of-sample R² is statistically significant at the 5% level.
Source: Invesco. Based on US large cap data from April 1991 to December 2017.

This formula implies a higher weighting of factor *X* at time t when the one-month forecast for the factor is more optimistic. In addition, the use of a constant *c* prevents the denominator from being zero or negative. Note that increasing *c* will render the weights less extreme. In our setting, *c* is 1.5.[8]

Constructing a dynamic multi-factor model from the dynamic factor weights
We now use this formula to construct a dynamic multi-factor model. Figure 1 shows the dynamic factor weights based on the above weight function *W(X,t)*.

Since our first out-of-sample forecast is for October 2004, this is where the time series of factor weights begins. During the financial crisis, the model put significantly more weight on relative value, switching over after 2014 to price trend and earnings momentum, which the market was rewarding at

the time. Indeed, in our sample, we see significant weight shifts only during the financial crisis (decreasing momentum, increasing relative value) and after 2014, when the reverse happened.

To enable investigation of a longer time series of factor weights, we reconstruct dynamic factor weights using generic data taken from the Fama and French US data set[9] for academic versions of the momentum, value and quality factors.

This extension of history covers several major episodes, such as the 1987 Black Monday, the 1990s recession, the Asian financial crisis in 1997, the Russian debt crisis in 1998, the inflation and collapse of the tech bubble between 1999 and 2002, and the global financial crisis from 2007 to 2009. Figure 2 shows the dynamic factor weights using the generic data. With this data set, the first out-of-sample forecast is for August 1978. Similar

Figure 1
The evolution of factor weights derived from a forecast combination based on 30 forecasts focusing on the out-of-sample period from Oct 2004 to December 2017
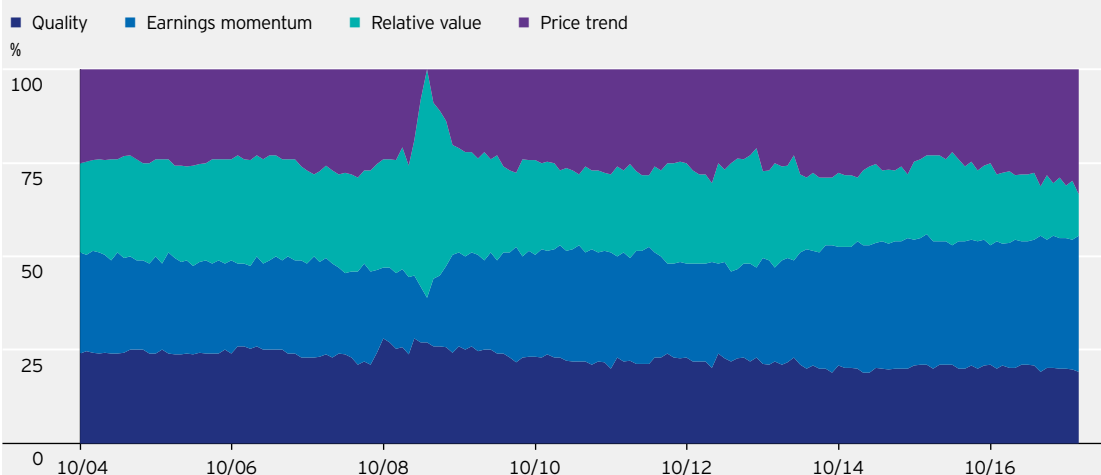
Source: Invesco. Based on US large cap data from October 2004 to December 2017.

Figure 2
The evolution of factor weights derived from a forecast combination based on 30 forecasts
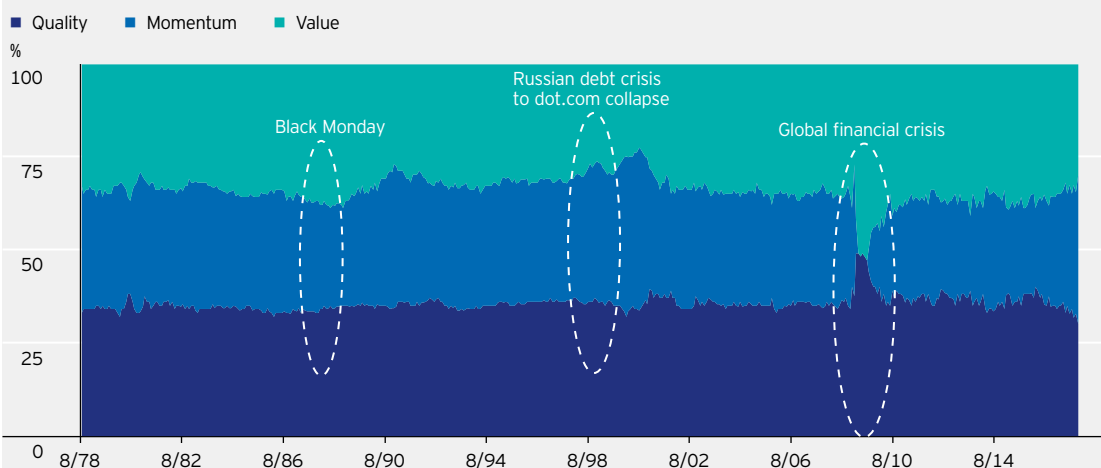
Source: Invesco. Based on the Fama and French US data set from August 1978 to December 2017.

Comparing static and dynamic market-neutral multi-factor models

	Static model	Dynamic model
Benchmark performance		
Return (%)	1.71	1.71
Risk (%)	0.54	0.54
Gross performance		
Return (%)	3.95	5.51
Active return (%)	2.24	3.80
Information ratio	0.64	1.09
Turnover and transaction costs		
Annual turnover (%)	46.79	163.18
Annual transaction costs (bps)	27.87	127.89
Net performance		
Return (%)	3.66	4.17
Active return (%)	1.95	2.46
Active risk (%)	3.47	3.50
Information ratio	0.56	0.70
Maximum drawdown (%)	-8.09	-7.07

Source: Invesco. Based on US large cap data from October 2004 to December 2017. Estimated transaction cost is based on a multi-sloped transaction cost model. **The figures refer to simulated past performance and past performance is not a reliable indicator of future performance.**

to figure 1, this model put significantly more weight on value and quality during the financial crisis and more weight on momentum after 2014. It also put more weight on momentum and less on value and quality around the dot.com era. After Black Monday in 1987, it put more weight on value and quality but less on momentum.

Table 3
Calendar-year net returns of the static and dynamic market-neutral multi-factor model

Holding period return net (active)	Static model	Dynamic model
2004	-0.15	-0.13
2005	1.52	-0.65
2006	2.59	3.26
2007	6.40	0.90
2008	1.40	-0.89
2009	-5.84	-4.09
2010	-0.07	3.72
2011	3.51	8.39
2012	0.21	2.36
2013	5.59	4.95
2014	2.17	7.10
2015	7.54	4.89
2016	-5.23	-2.42
2017	7.26	5.59

Source: Invesco. Based on US large cap data. **The figures refer to simulated past performance and past performance is not a reliable indicator of future performance.**

Examining the performance of the dynamic model over a static factor allocation
Recall that we modify the multi-factor score for each stock to take factor timing into account by multiplying the single-stock broad equity factor score by its respective dynamic factor weights at each point in time. By the same token, we can construct a static factor allocation model i.e. an equal weight factor allocation model, by multiplying the single-stock factor score by a constant, which is 0.25 in the case of four equity factors. In addition, we need to make sure that the factor exposure of this static factor allocation model is also the same for each of the four factors because factor scores change over time.

In short, we use portfolio simulations to compare performance gross and net of transaction costs of a multi-factor model with dynamically weighted factors versus a static weighting scheme. This comparison enables us to evaluate the value added of factor timing in a realistic setting.

Tables 2 and 3 contain the simulation results in a market-neutral portfolio implementation of the two multi-factor models.[10] First, we note that, in terms of gross performance, the dynamic multi-factor portfolio outperforms the static one by 1.56 percentage points p.a. (5.51% vs. 3.95%). Given similar risk for both strategies, this outperformance translates into a considerable difference in gross information ratios (1.09 vs. 0.64).

However, the outperformance comes at the cost of more than three times the turnover of the static model. As a consequence, some 100 bps of the relative excess returns of the dynamic model are lost when transaction costs are taken into account. Furthermore, the information ratio of the dynamic

model is reduced to 0.70, whereas the information ratio of the static model falls to 0.56. The difference between the two net information ratios, however, is statistically insignificant based on the paired test for equality of Sharpe ratios (p-value of 0.69).[11] In other words, a significant portion of the benefit of factor timing is offset by transaction costs.

A significant portion of the benefit of factor timing is offset by transaction costs.

Table 3 shows the calendar year net returns of the two models. The dynamic model outperformed the static model most noticeably during the recovery from the financial crisis in 2010 and 2011, when it reined in the exposure to value. These two years account for more than half of all the apparent performance advantage of the dynamically weighted model. The period from 2014 onwards, with a further reduction in exposure to value and an increased exposure to momentum, is also positive whereas the period ending in the early part of the financial crisis would have been better handled by the static model.

This evidence suggests treating overly positive factor timing claims with caution.

Conclusion
We have investigated the benefits of forecast combination for timing equity factors based on predictive regressions using time series macro predictors as inputs and testing a dynamic factor allocation in real world portfolio simulations. While the gross information ratio of the factor timing model clearly exceeds that of the static model,[12] the net information ratio exhibits a more modest difference. Obviously, this evidence suggests treating overly positive factor timing claims with caution. In future work, more research on robustness is needed to validate the genuine value of a dynamic factor allocation. Furthermore, the speed of factor timing can be examined by incorporating faster-moving factor efficacy metrics, such as factor momentum.[13]

References
Aiolfi, M., C. Capistrán and A. Timmermann (2011): Forecast Combination, In M. P. Clements and D. F. Hendry (eds.), Oxford Handbook of Economic Forecasting, Chapter 12, Oxford University Press.

Campbell, J. Y. and S. B. Thompson (2008): Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?, Review of Financial Studies 21(4), pp. 1509-1531.

Diebold, F. X. (2015): Comparing Predictive Accuracy, Twenty Years Later: A Personal Perspective on the Use and Abuse of Diebold-Mariano Tests, Journal of Business and Economic Statistics 33(1) pp. 1-9.

Dichtl, H., W. Drobetz, H. Lohre, C. Rother and P. Vosskamp (2019): Optimal Timing and Tilting of Equity Factors, Financial Analysts Journal 75(4), pp. 84-102.

Lee, J. H., Z. Shi and Z. Gao (2018): On LASSO for Predictive Regression, arXiv Working Paper.

Leung, P. L. and W. K. Wong (2008): On testing the equality of multiple Sharpe ratios, with application on the evaluation of iShares, Journal of Risk 10(3), pp. 15-30.

Kelly, B. and T. Gupta (2019): Factor Momentum Everywhere, Journal of Portfolio Management 45(3), pp. 13-36.

Kong, A., D. Rapach, J. K. Strauss and G. Zhou (2011): Predicting Market Components Out of Sample: Asset Allocation Implications, Journal of Portfolio Management 37(4), pp. 29-41.

Koop, G. and D. Korobilis (2012): Forecasting Inflation Using Dynamic Model Averaging, International Economic Review 53(3), pp. 867-886.

McCracken, M. W. and S. Ng (2016): FRED-MD: A Monthly Database for Macroeconomic Research, Journal of Business and Economic Statistics 34(4), pp. 574-589.

Moody, J., L. Wu, Y. Liao and M. Saffell (1998): Performance Functions and Reinforcement Learning for Trading Systems and Portfolios, Journal of Forecasting 17(5-6), pp. 441-470.

Rapach, D. E. and G. Zhou (2019): Sparse Macro Factors, Working Paper.

Wright, J. A., S. C. P. Yam and S. P. Yung (2012): A note on the test for the equality of multiple Sharpe ratios and its application on the evaluation of iShares, Technical Report.

Zhu, H. (2006): The Adaptive Lasso and its Oracle Properties, Journal of American Statistical Association 101(476), pp. 1418-1429.

Note that simulated performance is based on historical data.

Technical appendix

Adaptive LASSO

In short, LASSO[14] is Ordinary Least Squares (OLS) regressions with a penalty term to minimize over-fitting. Adaptive LASSO is chosen because this approach is more appropriate when some predictors are non-stationary, and it has better asymptotic properties compared to LASSO on its own. The objective function of Adaptive LASSO is:

$$\beta^* = \arg\min_\beta \left(y - \sum_{j=1}^{p} x_j \beta_j \right)^2 + \lambda \sum_{j=1}^{p} w_j \left| \beta_j \right|$$

where $w_j = \dfrac{1}{\left| \beta_j^{ini} \right|^\gamma}$

Unlike LASSO, the weight w adjusts the penalty differently for each coefficient.

Dynamic model averaging (DMA)

DMA is a state space framework that improves forecasts by minimizing not only the impact of parameter instability (regression coefficients changing over time) but also model uncertainty (which model is the best in terms of predictability).

Suppose we have M models that are characterized by having different subsets of predictors. Unlike LASSO, DMA averages the forecasts across models. Those containing important combinations of predictors receive high weights in the averaging process:

$$r_{t+1} = z_t^{'(m)} \beta_t^{(m)} + \varepsilon_{t+1}^{(m)} \qquad \varepsilon_t^{(m)} \sim N\left(0, V_t^{(m)}\right)$$

$$\beta_t^{(m)} = \beta_{t-1}^{(m)} + \omega_t^{(m)} \qquad \omega_t^{(m)} \sim N\left(0, W_t^{(m)}\right)$$

Recall that, at every point in time, ALASSO will pick a subset of the macro predictors as independent variables, and DMA will average the forecasts of a subset of state space models. So, both approaches restrict the set of indicators, but in a different manner.

Given dynamic factor weights, we can construct a dynamic multi-factor model by multiplying the single-stock factor score by its respective dynamic factor weights (i.e. dynamic factor exposure at each point in time) and examine the performance gross and net of transaction costs via portfolio simulations relative to a static multi-factor model (i.e. a multi-factor model with each factor exposure equally weighted). This comparison enables us to evaluate the value added of factor timing in a realistic setting.

Softmax function

The softmax function enables us to compute a probability of choosing an action based on its estimated value. See Moody, Wu, Liao and Saffell (1998):

$$W^{soft\,max}(X,t) = \frac{F(X,t)}{\sum_{X \in \{RV, PT, EM, Q\}}^{4} e^{F(X,t)}}$$

About the authors

Michael Fraikin
Global Head of Research,
Invesco Quantitative Strategies
Michael Fraikin and his team are responsible for the maintenance and further development of the quantitative models that drive decisions within the investment products.

Edward Leung, Ph.D.
Quantitative Research Analyst,
Invesco Quantitative Strategies
Edward Leung conducts quantitative research into factors, portfolio construction and risk modelling, currently focusing on machine learning and alternative data.

Dr. Harald Lohre
Senior Research Analyst,
Invesco Quantitative Strategies and Visiting Research Fellow EMP/Lancaster University
Dr. Harald Lohre develops quantitative risk and return forecast models used in the management of multi-asset multi-factor strategies.

Notes
1 The US large cap universe builds on the Russell 1000 constituents.
2 https://research.stlouisfed.org/econ/mccracken/fred-databases/.
3 See Rapach and Zhou (2019).
4 See Kong, Rapach, Strauss, and Zhou (2011) and Aiolfi, Capistrán and Timmermann (2011).
5 The statistical significance is determined by the DM test. See Diebold (2015) for details.
6 See Lee, Shi and Gao (2018), Zhu (2006) and the technical appendix.
7 See Koop and Korobilis (2012) and the technical appendix.
8 We have also tried solving the issue of the denominator being zero or negative by applying the softmax function from ML, see Moody, Wu, Liao and Saffell (1998) and the technical appendix. Yet, the dynamic weights are similar.
9 http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html#Research.
10 Our market-neutral portfolio consists of two portfolio legs, one with long and one with short names which broad market exposures are expected to neutralize in the aggregate. The optimal portfolio implementation derives from a mean-variance optimization of the dynamic alpha subject to a minimum level of risk and beta exposure using cash as benchmark.
11 See Wright, Yam and Yung (2012) and Leung and Wong (2008).
12 Dichtl, Drobetz, Lohre, Rother and Vosskamp (2019) arrive at a similar conclusion.
13 See Kelly and Gupta (2019).
14 Least Absolute Shrinkage and Selection Operator (LASSO).

About risk

The value of investments and any income will fluctuate (this may partly be the result of exchange rate fluctuations) and investors may not get back the full amount invested.

Important information

www.invesco.com

当資料ご利用上のご注意

当資料は情報提供を目的として、インベスコ・アセット・マネジメント株式会社(以下、「当社」といいます。)が当社グループの各運用拠点に在籍する運用プロフェッショナル（以下、「作成者」）が作成した英文資料を当社グループから入手してご提供するものです。当資料は信頼できる情報に基づいて作成されたものですが、その情報の確実性あるいは完結性を表明するものではありません。また過去の運用実績は、将来の運用成果を保証するものではありません。当資料に記載された一般的な経済、市場に関する情報およびそれらの見解や予測は、いかなる金融商品への投資の助言や推奨の提供を意図するものでもなく、また将来の動向を保証あるいは示唆するものではありません。また、当資料に示す見解は、インベスコの他の運用チームの見解と異なる場合があります。本文で詳述した当資料の分析は、一定の仮定に基づくものであり、その結果の確実性を表明するものではありません。分析の際の仮定は変更されることもあり、それに伴い当初の分析の結果と重要な差異が生じる可能性もあります。当資料について事前の許可なく複製、引用、転載、転送を行うことを禁じます。

受託資産の運用に係るリスクについて

受託資産の運用にはリスクが伴い、場合によっては元本に損失が生じる可能性があります。各受託資産へご投資された場合、各受託資産は価格変動を伴う有価証券に投資するため、投資リスク（株価の変動リスク、株価指数先物の価格変動リスク、公社債にかかるリスク、債券先物の価格変動リスク、コモディティにかかるリスク、信用リスク、デフォルト・リスク、流動性リスク、カントリー・リスク、為替変動リスク、中小型株式への投資リスク、デリバティブ《金融派生商品》に関するリスク等）による損失が生じるおそれがあります。ご投資の際には、各受託資産の契約締結前書面、信託約款、商品説明書、目論見書等を必ずご確認下さい。

受託資産の運用に係る費用等について

投資一任契約に関しては、次の事項にご留意ください。【投資一任契約に係る報酬】直接投資の場合の投資一任契約に係る報酬は契約資産額に対して年率0.88％（税込）を上限とする料率を乗じた金額、投資先ファンドを組み入れる場合の投資一任契約に係る報酬は契約資産額に対して年率0.605％（税込）を上限とする料率を乗じた金額が契約期間に応じてそれぞれかかります。また、投資先外国籍ファンドの運用報酬については契約資産額に対して年率1.30％を上限とする料率を乗じた金額が契約期間に応じてかかります。一部の受託資産では投資一任契約に加えて成功報酬がかかる場合があります。成功報酬については、運用戦略および運用状況などによって変動するものであり、事前に料率、上限額などを表示することができません。【特定（金銭）信託の管理報酬】当該信託口座の受託銀行である信託銀行に管理報酬をお支払いいただく必要があります。具体的料率については信託銀行にご確認下さい。【組入有価証券の売買時に発生する売買委託手数料等】当該費用については、運用状況や取引量等により変動するものであり、事前に具体的な料率、金額、上限または計算方法等を示すことができません。【費用合計額】上記の費用の合計額については、運用状況などによって変動するものであり、事前に料率、上限額などを表示することができません。