Scientific Portfolio
An EDHEC Venture

A Scientific Portfolio Publication

Remember to Diversify Your Active Risk: Evidence from US Equity ETFs

March 2023

The authors can be contacted at contact@scientificportfolio.com.

About the Authors

Benjamin Herzog is CEO of Scientific Portfolio and has been with EDHEC since 2019. Previously, he started his career at Société Générale where he occupied several positions in quantitative research and structuring, and headed the equity QIS structuring team. He received Master's degrees in quantitative finance from Paris VI university and in computer science and artificial intelligence from Columbia University and holds a general engineering Master's degree from Ecole Centrale de Lyon.

Jenna Jones joined Scientific Portfolio in January 2022 as a Senior Researcher. Previously, she worked as an equity research analyst at Bloomberg and as a business analyst at a consultancy firm, and spent four years working as a researcher for an investment fund. She received a Bachelor's degree in Finance from the University of Florida and a Master's in Actuarial Science from HEC Lausanne, where she specialized in extreme event risks.

Shahyar Safaee joined Scientific Portfolio in September 2022 as Deputy CEO and Director of Business Development. Previously, he built a 20-year track record in both sell- and buy-side roles in London, Paris and NYC, primarily serving institutional clients at J.P. Morgan by taking and managing financial risks before joining EDHEC-Risk Institute to contribute to both research and business development efforts. He holds Master's degrees in engineering (Ecole des Mines de Saint-Etienne) and financial mathematics (Université Claude Bernard in Lyon).

Abstract

In this article, we estimate the level of risk diversification for a universe of US equity ETFs and observe the benefits of diversification for budgeting of active risk relative to a cap-weighted benchmark. We first introduce a risk model that only needs historical returns to break down relative risk into individual factor-related risk contributions, allowing for the construction of a measure of concentration directly inspired by the Equally weighted Risk Contribution (ERC) concept. We then use this concentration measure to highlight the impact of diversification on tracking error stability and find conclusive empirical evidence that US equity ETFs that have a diversified set of systematic active risk contributions have a more stable tracking error. These results suggest that investors seeking a stable tracking error for active risk budgeting purposes may benefit from selecting those ETFs that have a strong level of risk diversification.

Key Takeaways

- Investors only need historical returns to assess the level of risk diversification of a portfolio. The simple diversification metric that we propose relies on recent innovations in risk modelling that reconcile fundamental and statistical approaches.
- The diversification of factor-based active risk contributions is an effective way to manage active risk. It is associated with a more stable tracking error and in turn, facilitates budgeting of active risk.
- Investors in the US equity ETF market may benefit from these insights when selecting ETFs with strategies that are either actively managed or deviate from broad-based cap-weighted benchmarks.

Introduction

Introduction

Nearly fifty years ago, Jack Bogle created the first index mutual fund, now known as the Vanguard 500 (Arnott and Sherrerd (2022)). Its purpose was to provide retail investors with exposure to the returns of the Standard & Poor's 500 market-capitalisation weighted (CW) index. Specifically, Bogle created a passive investment product that was transparent and offered at a low cost for retail investors. Although some pension funds had already been able to invest in index strategies, ordinary retail investors did not have the practical means to do so (Malkiel (2022)). With the launch of the Vanguard 500 index mutual fund, the foundation was laid for the next step in the emergence of exchange-traded funds (ETFs) (Arnott and Sherrerd (2022)).

The next milestone was the creation of a vehicle that allowed a portfolio of stocks to be traded on stock exchanges as a single stock: the exchange-traded fund was born. In 1989, the world's first ETF, known as the Cash Index Participation Unit (CIP)[1], began trading on the Philadelphia Stock Exchange (PHLX). This was the first publicly available 'portfolio in a share' product in the US that could be traded on a stock exchange, making it available to both retail and institutional investors (see for example Ruggins (2017) for a detailed review on the history of CIPs). CIPs were designed in such a way as to provide more liquidity in the market and make trading much more efficient. The creation of CIPs was a seminal event that compelled the SEC to formally define an exchange-traded product and create legislation around it, setting the stage for the explosive growth of ETFs that later followed.

The first ETFs were designed as trading tools, an alternative to futures, that generated trading volume on stock exchanges for market participants that were unable or unwilling to trade on futures exchanges. Later, ETFs evolved to become suitable for long-term investors, and once structured as fund vehicles, were able to passively track CW indexes. However, this approach has since been acknowledged to have certain limitations. One such drawback of CW indexes is that a stock's weight is directly tied to its market capitalisation. This can cause overvalued stocks to be given a disproportionately high weight and undervalued stocks to be given a lower weight in the index (Colby (2007)), in direct contradiction to the intuitive value-driven investment strategy of buying low and selling high. Consequently, any overvaluation in a stock's price is magnified which can lead to a lag in performance compared to indexes with weightings based on other methods than valuation. Before 2005, there was no low-cost alternative weighting strategy (i.e., non-CW ETFs), that addressed the limitations of market capitalisation weighting; however, soon after, Rob Arnott, Jason Hsu and Philip Moore (2005) developed the Fundamental Index, a strategy that ignored stock price and market capitalisation, and instead weighted stocks using fundamental measures (Arnott and Sherrerd (2022)). This new strategy preserved a passive investment approach through combining an index-delivery structure with an alternative investment strategy. These straightforward, transparent, cost-effective investing solutions were later coined "smart beta" strategies, meaning valuation-indifferent strategies that follow a rules-based approach and often seek exposure to equity factors based on financial metrics such as value, quality, or momentum.

Today, there are two groups among the non-CW ETFs: those tracking a non-CW index (including sector ETFs and smart beta ETFs) and, more recently, those following a strategy devised by an active manager. Non-CW equity ETFs have in common their exposure to systematic equity risk factors other

1 - CIPs did not invest directly in stocks but rather utilised futures to mimic stock price movements.

Introduction

than the overall market factor (e.g., either fundamental risk factors that are deemed rewarded over the long term such as value, size and momentum, or industry risk factors). Actively managed ETFs presumably also include an additional layer of specific (non-systematic) risk linked to the manager's skill. Their recent growth (both in number of funds and in AUM) is often attributed to several features that differentiate them from actively managed mutual funds, such as tax efficiency (secondary market trading and in-kind redemptions/creations ensure that tax liabilities are not mutualised between investors), intraday pricing and liquidity and generally lower fees.

Equity investors now have access to a large variety of ETFs following passively managed smart beta strategies as well as actively managed strategies. At the end of June 2022, although ETFs that follow CW equity benchmarks still represented 77% of the global equity ETF AUM, 18% of the equity ETF *instruments* globally available to investors were following a passive smart beta strategy and 14% were actively managed[2].

It is reasonable for investors looking to allocate capital *away* from a CW ETF *to* a given non-CW ETF to assess the relative (incremental) risk they can expect from the latter. Note that for the remainder of the article and for ease of reading, we will interchangeably designate this relative risk as "active risk". More generally, the decision to invest in a non-CW ETF is always going to be implicitly compared to the optically safer choice of investing in a CW ETF because the latter is fully understood by investors and has acquired a *you-get-what-you-pay-for* status. The risk resulting from the purchase of a non-CW ETF is not only financial (the investment policy followed by the fund may indeed lead to relative underperformance compared to a CW ETF) but also somewhat reputational (for an institutional investor) or psychological (for a retail investor) because "deviating from the crowd" increases the chances of being singled out, for better or for worse. The metric traditionally provided by active managers or promoters of non-CW indices to inform an investor's assessment of active risk is the Tracking Error (TE), which can be estimated from the past returns of an ETF. Investors therefore translate their combined levels of risk appetite and trust with respect to a non-CW investment strategy into a TE risk budget they are willing to tolerate. Unfortunately, an empirically estimated TE is by design an ex-post (backward-looking) risk metric that may not always be a good estimator of future TE, especially if the investment strategy followed by the non-CW ETF generates large variations of TE. Therefore, investors need to carefully analyse the make-up of the observed TE after having identified its various components to ultimately assess whether they are comfortable relying on a backward-looking risk measurement for active risk budgeting purposes. Our intention in this article is to propose an ex-post risk analysis methodology that will help investors make such an assessment for non-CW ETFs.

It is well established that rigorous risk analysis ought to identify and separate systematic risk from non-systematic (specific) risk, because the former is more persistent, can be somewhat quantified, and cannot be fully diversified away. This approach is equally valid for investors trying to analyse active (relative) risk. For the equity asset class, the standard practice within academia and in the industry is to use a common set of factors to describe the systematic risk of equity portfolios (see for example Amenc and Goltz (2016) for an overview). Provided these systematic factors explain most of the risk

2 - https://etfgi.com/

Introduction

prevailing in the equity portfolios we wish to analyse, it may then be possible for a risk model to break down relative risk (and therefore TE) into individual factor-related *risk contributions*. The concept of risk contribution is commonly used among institutional investors to budget risk in accordance with the constraints of their mandate (see Qian (2006) for an intuitive interpretation of risk contributions as expected contributions to potential portfolio losses), and our study here aims at helping investors with their active risk budgeting process when reviewing a non-CW ETF: the more ex-post evidence one has about the stability of an ETF's TE, the easier it is to determine whether an appropriate amount of capital was allocated to this ETF as part of a risk budgeting review.

The question of how investors should use the relative risk contributions (obtained via the factor risk model) may be first addressed with a bit of intuition. A key insight of portfolio theory is indeed that diversifying the sources of uncertain returns leads (all else being equal) to lower portfolio risk, and more specifically, to smaller variations in portfolio returns. Going slightly "meta" and applying this insight to volatility itself tells us that diversifying the sources of (i.e., the contributions to) volatility should lead to smaller variations in portfolio volatility, all else equal. In other words, *a diversified set of active risk contributions should be associated with smaller variations in TE*. Maillard, Roncalli and Teïleteche (2010) have shown that diversifying risk contributions is an interesting portfolio construction heuristic (proved to be optimal under specific assumptions) that somewhat addresses the shortcomings of two other portfolios commonly used by practitioners: the naïve equal-dollar (1/N) portfolio and the minimum variance (MV) portfolio. Despite its out-of-sample robustness, the risk of the 1/N portfolio is often concentrated in the most volatile assets, while the MV portfolio is often concentrated (both in dollar and risk contributions) in the assets with the lowest ex-ante volatility, leading to out of sample instability. The Equal Risk Contribution (ERC) portfolio can be seen as a risk-aware 1/N portfolio, aiming for a balanced risk budget. It is shown to produce improved out-of-sample results compared to the two other portfolios, especially when the portfolio constituents are not strongly correlated. More specifically, the ERC portfolio is overall less volatile than the 1/N portfolio and less concentrated than the MV portfolio, therefore carrying smaller tail risk (VaR, drawdown). Roncalli and Weisang (2012) transpose the diversification problem to a factor-based framework (with fewer sources of risk that tend to be less correlated) and confirm that the factor-ERC portfolio has attractive properties in terms of tail risk (smaller kurtosis, smaller drawdowns) compared to the 1/N portfolio and even the asset-ERC portfolio. These properties are particularly useful for our (active/relative) risk budgeting exercise: a smaller tail risk is associated with limited variations in volatility and therefore, in line with our earlier intuition, the ERC portfolio appears to be a good reference point when evaluating the stability of a portfolio's risk.

To conclude, we now transpose the above insights to the universe of non-CW ETFs and apply the (factor) ERC concept to active risk: we expect non-CW ETFs whose TE is diversified in terms of (factor) active risk contributions to carry lower active tail risk and therefore have experienced a more stable TE. Conversely, non-CW ETFs whose TE is concentrated (in terms of factor active risk contributions) are expected to have had a more volatile TE, making the active risk budgeting exercise more challenging for investors.

Introduction

The objective of the article is to empirically test this insight for the US equity ETF market by proposing an ex-post methodology to assess active risk diversification based on the proximity with an ERC portfolio of active risks. Note that our analysis does not seek to uncover an ex-ante metric that would have attractive out-of-sample predictive power on the stability of future TE, but rather aims to verify that risk diversification and stable TE have historically gone hand in hand, thus providing evidence that one way of stabilising TE (for active risk budgeting purposes) is to maintain a good level of diversification of active risk contributions. A key feature of our methodology is that it does not require any knowledge of an ETF's holdings. The rest of the article is organised as follows. In the first section, we present our ETF dataset. In the second section, we briefly introduce the equity risk model we intend to use to assess factor risk contributions. In the third section, we introduce a measure of concentration in terms of factor active risk contributions directly inspired from the ERC concept and provide some intuitive interpretation of it. In the final section, we empirically observe the benefits of diversification for active risk budgeting in the US equity ETF universe.

Dataset

Dataset

Our study relies on a comprehensive database of ETFs kindly provided by Trackinsight[3], a leading provider of data and analytics on ETF markets. We use data reported as of September 2022. Out of a database of over 2,800 ETFs domiciled in the United States, our universe consists of 862 ETFs that are specifically pursuing a US "Delta One" equity strategy (i.e., we excluded ETFs exposed to other assets classes as well as equity ETFs involving leverage or non-linear/option-based strategies) and represents approximately USD3.8 trillion in assets under management (AUM). Our analysis is conducted on weekly total returns data covering the five-year period from October 2017 to September 2022, so our analysed dataset naturally only includes the subset of ETFs for which the full five years of returns with no material gaps are available. Additionally, for the purposes of our analysis, it is necessary that an ETF carries material relative risk with respect to a CW benchmark (determining the diversified or concentrated nature of an immaterial risk is indeed a moot point). Accordingly, ETFs which track or loosely track a broad, cap-weighted benchmark are excluded. More specifically, we exclude all ETFs that exhibit a TE less than 2% with respect to the SciBeta United States Cap-Weighted index (which is comprised of the 500 largest US stocks and whose returns are therefore close to those of the S&P 500 index)[4]. We determine the TE lower bound of 2% by identifying the types of ETFs we reasonably expect to fall within the scope of our analysis and determining a lower bound for each group. We consider that sector, actively managed and smart beta ETFs typically have strategies that deviate materially (based on their TE) from a CW benchmark and therefore wish to include them in the analysis. A review of both sector and actively managed ETFs indicates that they tend to have a higher level of TE than the smart beta group[5], and that the lower bound of TE will be driven by the smart beta segment. To obtain a lower bound that does not heavily depend on the universe of ETFs or the five-year sample period we have decided to analyse, we examine the five-year rolling tracking error for the SciBeta United States Maximum Deconcentration index3 (which is used to represent the smart beta segment of ETFs) versus the SciBeta United States Cap-Weighted index beginning in 2002 and ending in 2022. The SciBeta United States Maximum Deconcentration index is close to the equally weighted (EW) on the same universe as the SciBeta United States Cap-Weighted index, meaning that the same selection of stocks is made; however, the portfolio allocation within the two indices is different. As a result, we consider the SciBeta United States Maximum Deconcentration index as a good candidate to represent the subset of smart beta ETFs with the lowest TE. Figure 1 shows the resulting five-year rolling TEs by percentiles.

Figure 1: The five-year daily rolling tracking error of the EW index versus the CW index from 2002-2022, by percentiles

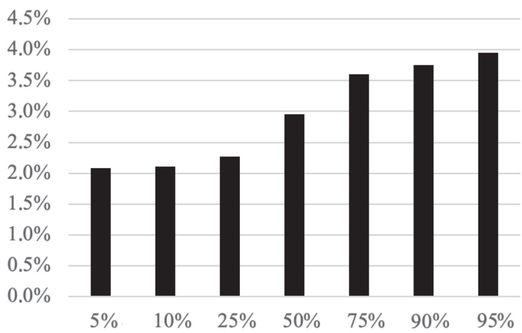

Note: The histogram shows the percentiles of the five-year daily rolling tracking error of the SciBeta United States Maximum Deconcentration index (EW) versus the SciBeta United States Cap-Weighted index (CW) from 2002 - 2022. The vertical axis shows the level of tracking error below which an ETF falls into the given percentile.

3 - https://www.trackinsight.com/
4 - https://www.scientificbeta.com/
5 - We rely on Trackinsight's database to identify sector ETFs (using the '*class_sector*' field), smart beta ETFs (using the '*class_factor*' field), and actively managed ETFs (using the '*ETFs$actively_managed*' field).

Dataset

The tracking error varies between approximately 2% and 4%. Therefore, 2% is selected as the lower bound of TE for the inclusion of ETFs for analysis. The resulting universe is comprised of 488 ETFs, representing AUM of USD2.3 trillion (approximately 61% of the initial US equity-focused universe AUM).

Risk Model

Risk Model

The remaining population of ETFs all contain a material amount of risk relative to the CW benchmark, measured by their TE. To precisely isolate, quantify and analyse the active risk of an ETF, we use a new risk model (Vaucher (2023)) that builds upon the IPCA model (Kelly, Pruitt and Su (2019)) in order to reconcile time-series/beta-based approaches with cross-sectional/fundamental approaches. Indeed, it is universally accepted that multivariate betas are the appropriate measures of systematic risk because they capture the covariances of a portfolio with respect to systematic risk factors. Put very simply, betas are covariances and covariances are risks. The academic debate revolving around characteristics and covariances is exclusively linked to asset pricing and the possible sources of equity expected returns (i.e., long term performance). For the purpose of risk analysis, we find no conceptual reason to resort to fundamental characteristics as is sometimes seen among industry practitioners. Indeed, one-dimensional characteristics are not good proxies for covariance-aware betas, and they are often subject to various layers of processing (e.g., winsorisation, creation of composite scores) that further increase the conceptual mismatch between risk exposures and characteristics (see for example Amenc et al. (2018) for a review of industry practices using equity characteristics).

This being said, multivariate betas create implementation challenges that should not be underestimated. They largely depend on the ad-hoc definition of a chosen list of factors (potentially creating model-dependency), the number of which is constrained to avoid collinearity issues that may affect the decomposition of systematic risks. A common workaround is the use of principal component analysis (PCA) and implicit factors, but this comes at the cost of losing intuition with respect to the economic meaning of each factor and ultimately reducing the actionability of the risk model; that is, the ability to easily amend unwanted risk exposures via asset allocation actions.

Our proposed methodology is to use the IPCA model while replacing fundamental equity descriptors by market-adjusted univariate exposures (estimated from historical returns) to explicit financial risk factors (regressors). More specifically, our approach makes use of an exhaustive list of systematic risk measures, including exposures to both fundamental (rewarded) risk factors and industry (non-rewarded) risk factors. The naturally overlapping nature of these exposures is then appropriately accounted for by the IPCA model via a reduced number of orthogonalised implicit factors, ultimately leading to a set of covariance-aware implicit betas. Systematic risk can then be decomposed into additive contributions that are reallocated back to the set of economically intuitive fundamental and industry factor exposures. The resulting model therefore leverages the IPCA methodology to combine the main benefits of both fundamental and statistical approaches in order to carry out portfolio risk analysis. Another major advantage of this risk model is that it only requires fund-level time-series data instead of needing to source and manage complex data at the stock-holdings level.

For a given ETF, the set of all (market-adjusted and univariate) exposures therefore becomes an informative risk identification card that fully determines, once processed by the model, the distribution of risk contributions across a group of 17 systematic sources of risk: seven fundamental equity factors (market, momentum, value, size, low volatility, investment and profitability) and 10 industries (Communication Services, Consumer Discretionary, Consumer Staples, Energy, Financials, Health Care,

Risk Model

Industrials, Information Technology, Materials, Utilities). The sum of the 17 risk contributions is equal to the total systematic volatility of the ETF.

We apply the risk model to our universe of ETFs and measure its performance by computing the R-squared for each ETF. If an ETF has a large proportion of its risk unexplained by the model, shown by an R-squared of less than 75%, it is excluded from the analysis. This exclusion process removes an additional 15 ETFs with a minimal impact on the AUM representation. Table 1 shows the model's precision at different levels for the remaining universe of ETFs.

Table 1: The R-squared pertaining to the universe of ETFs used for analysis

	R^2
25th percentile	93%
Mean	95%
Median	97%

The levels of R-squared reported in Table 1 confirm that the risk model is able to explain risk with a high level of precision, thus providing us with a satisfactory level of comfort as we plan to compute the active risk contributions for ETFs in our universe.

Measuring Active Risk Concentration in Non-CW US Equity ETFs

Measuring Active Risk Concentration in Non-CW US Equity ETFs

The risk model introduced in the previous section allows us to produce a decomposition of systematic active risk (i.e., the systematic volatility of the ETF's relative performance with respect to the CW benchmark) into individual factor-based active risk contributions, with a total of 17 systematic active risk contributions, which naturally add up to the systematic (i.e., explained by our model) portion of the TE. Note that we leave specific (unexplained) risk outside the scope of our analysis and do not opine on whether (and how) it contributes to the concentration or the diversification of an ETF. In this section, we propose to use this decomposition to construct a measure of concentration, which we call Concentration Risk Measure (CRM) that will help identify the non-CW ETFs whose TE is expected to be stable because of their proximity to the ERC portfolio. For a given ETF, the CRM measure we propose is defined as:

$$\text{CRM} = \frac{\sqrt{\sum_{i=1}^{17} \text{Active Risk Contribution}_i{}^2}}{\text{Systematic TE}}$$

where Systematic TE is the systematic portion of the TE of the ETF, equal to the sum of the 17 active risk contributions, i.e., $\sum_{i=1}^{17}$ Active Risk Contribution$_i$.

The CRM measure is effectively the 2-norm applied to the vector of the 17 active risk contributions (normalised so they add to 1). This follows the insight provided by DeMiguel et al. (2009) that controlling the 2-norm measure of the distance between the composition of a given portfolio and that of the 1/N portfolio can be directly achieved via a control on the 2-norm of the composition of the portfolio itself. CRM is therefore qualifying the distance (in terms of normalised active risk contributions) between an ETF and its corresponding hypothetical ERC portfolio, where the latter is designed to have a total risk equal to the systematic portion of the TE of the ETF, i.e., $\sum_{i=1}^{17}$ Active Risk Contribution$_i$. The higher the CRM value for an ETF, the more concentrated (that is, further away from its corresponding ERC portfolio) the ETF is and therefore, the more we expect to see large variations in its TE.

A risk-based interpretation of CRM is also possible, noting that $\frac{1}{CRM^2}$. is the *effective number of* (normalised) active risk contributions for the ETF (see Martellini and Milhau (2018) for practical examples of the use of the concept of "*effective number of*", directly related to the Herfindahl index) and remembering that a hypothetical portfolio of N uncorrelated assets sharing the same mean and volatility has a volatility of $\frac{\sigma}{\sqrt{N}}$ where σ is the volatility of each individual asset, indicating that diversification reduces portfolio volatility by a factor of $\frac{1}{\sqrt{N}}$. Using this analogy, we can see that CRM is homogeneous to a volatility and can be interpreted as a proxy for the diversification-related reduction factor created by a given portfolio.

We conclude this section by noting again that our risk decomposition and therefore our concentration risk measure does not require knowledge of holdings data.

Benefits of Diversification for Active Risk Budgeting in US Equity ETFs

Benefits of Diversification for Active Risk Budgeting in US Equity ETFs

We now intend to use the concentration risk measure CRM defined in the previous section to analyse the universe of non-CW US equity ETFs and highlight the benefits of diversification of active risk contributions.

First, we conduct an analysis where we identify whether the level of CRM for an ETF impacts its level of systematic TE. Figure 2 shows the systematic TE of ETFs in the universe conditional upon different levels of concentration as measured by CRM. Although the curve is not strictly monotonous, the trend clearly indicates that, as expected, systematic TE increases as CRM increases. For example, the median systematic TE in the lowest decile is approximately 6% while it is approximately 14% in the highest decile. This is consistent with the academic view (and findings) that the ERC portfolio is relatively low-risk and the more general idea that factor-based risk diversification effectively reduces portfolio risk. Note that we chose to represent median values in Figure 2 rather than averages in order to mitigate the impact of outliers and account for the fact that risk contribution concentration is not the only factor influencing the systematic TE of an ETF. Two ETFs may indeed be equally diversified across active risk contributions (i.e., share the same level of CRM because they are similarly close to their corresponding ERC portfolio) but be at the same time exposed to sources of risk of varying intensity, resulting in materially different levels of TE.

Figure 2: Median systematic TE conditional on a CRM decile

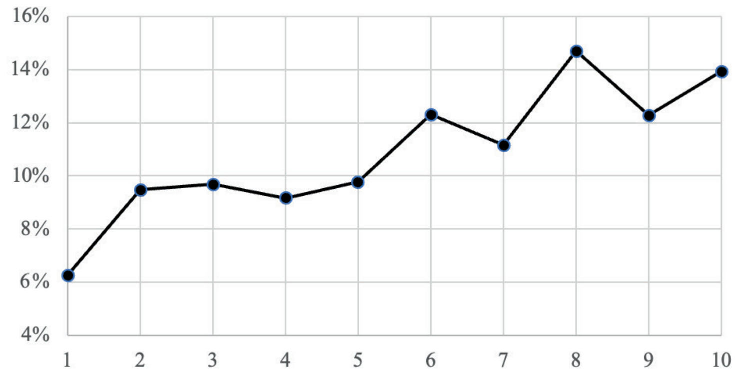

Note: The horizontal axis represents the deciles of ETFs measured by CRM. The first decile represents ETFs which are the least concentrated (i.e., most diversified) in terms of factor-based active risk diversification while the tenth decile represents ETFs which are the most concentrated in terms of factor-based active risk diversification. The number of ETFs in each decile is approximately 48. The vertical axis represents the median value, for each group of ETFs in a decile, of the systematic TE.

We now turn to an analysis of the variations of the TE over time to determine whether the stability of TE is dependent on the level of risk diversification of an ETF. We use daily returns data to compute TE for six-month daily rolling windows across the five-year period. This generates a time series of short-term TEs.

To determine the impact of CRM on the level of TE stability of non-CW ETFs, the median value of sample standard deviations of the six-month daily rolling TE's time series is calculated and plotted against CRM deciles. Results are displayed in Figure 3. We see, as with Figure 2, an upward sloping curve that indicates that concentration of active risk contributions is associated with more variability in TE. For

Benefits of Diversification for Active Risk Budgeting in US Equity ETFs

example, the median value of the TE standard deviation for the most concentrated decile is more than 1.5x greater than that of the most diversified decile. This is again consistent with the expected behaviour of the ERC portfolio and its reduced kurtosis compared to other heuristic approaches (see Maillard, Roncalli and Teïleteche (2010)).

Figure 3: Median standard deviation of six-month daily rolling TE conditional on a CRM decile

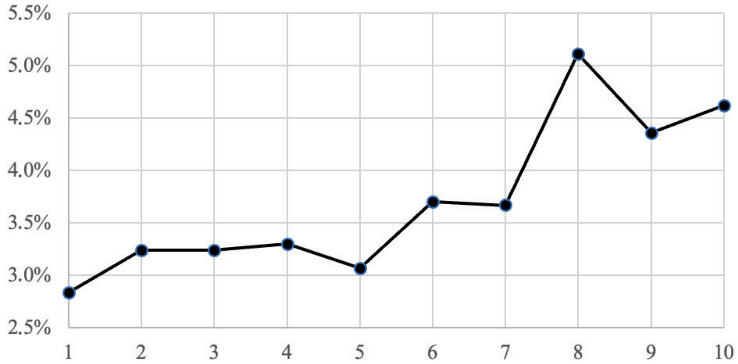

Note: The horizontal axis represents the deciles of ETFs measured by CRM. The first decile represents ETFs which are the least concentrated (i.e., most diversified) in terms of factor-based active risk diversification while the tenth decile represents ETFs which are the most concentrated in terms of factor-based active risk diversification. The number of ETFs in each decile is approximately 48. The vertical axis represents the median value, for each group of ETFs in a decile, of the time-series standard deviation of six-month daily rolling TEs.

We conclude this section by examining the relationship between active risk concentration and extreme moves in TE. The (historically) highest 5% of the six-month rolling TEs represents the tail of the time-series distribution of the TEs and therefore acts as a measure for extreme active risk. Figure 4 shows the median of the difference between the 5% quantile of TEs and the average of TEs, plotted against CRM deciles. We observe the same upward slopping pattern as in the two previous exhibits, confirming the consistency of our empirical results with the expected properties of portfolios diversified in terms of active risk contributions, namely their reduced active tail risk which translates naturally into a more stable TE.

Figure 4: Median value of 5% historical quantile of six-month daily rolling TE minus average of six-month daily rolling TE, conditional on a CRM decile

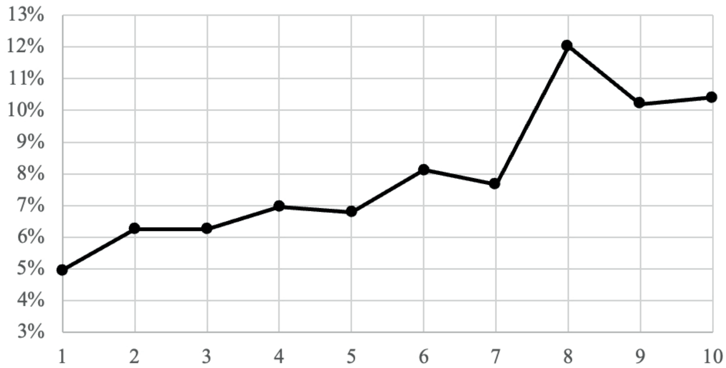

Note: The horizontal axis represents the deciles of ETFs measured by CRM. The first decile represents ETFs which are the least concentrated (i.e., most diversified) in terms of factor-based active risk diversification while the tenth decile represents ETFs which are the most concentrated in terms of factor-based active risk diversification. The number of ETFs in each decile is approximately 48. The vertical axis represents the median value, for each group of ETFs in a decile, of the difference between the 5% historical quantile of six-month daily rolling TE and the time-series average of six-month daily rolling TEs.

Conclusion

Conclusion

Using only historical returns, our proposed equity risk model combines the benefits of fundamental and statistical approaches and explains a large portion of the risk of US equity ETFs (mean R-squared of 95% and median R-squared of 97% for the set of funds we have analysed), allowing us to assess the concentration of systematic active risk contributions in US equity ETFs whose investment strategy deviates from a cap-weighted benchmark. Financial intuition and the insights of the literature on the Equally Weighted Risk Contribution portfolio tell us that non-cap-weighted portfolios that are exposed to a diversified set of systematic active risk contributions should also have a more stable tracking error. The US equity ETF market provides conclusive empirical evidence of this effect, indicating that investors who wish to budget their active risk more accurately may do so by primarily selecting those non-cap-weighted ETFs that maintain a strong level of diversification. In the end, a diversified TE budget is not only a promise of robustness, but also the sign of prudent risk-taking by some managers of non-CW ETFs. These managers have likely made more balanced bets compared with managers with concentrated exposures via a small number of risk factors (therefore possibly generating large variations in TE, and in turn, large variations in relative performance). The finance literature reminds us that, on average, winning managers do not repeat their outperformance, so opting for a concentrated TE may, in the long-term, lead to disappointing risks and risk-adjusted returns.

References

References

- Amenc, N. and F. Goltz (2016). Long-Term Rewarded Equity Factors: What Can Investors Learn from Academic Research?" *Journal of Index Investing* 7(2): 39-56.
- Amenc, N., F. Goltz, A. Lodh and B. Luyten (2018). Measuring Factor Exposure Better to Manage Factor Allocation Better. Scientific Beta white paper.
- Arnott, R., J. C. Hsu, and P. Moore (2005). Fundamental Indexation. *Financial Analysts Journal* 61(2): 83–99.
- Arnott, R. and K. Sherrerd (2022). Jack Bogle and Smart Beta: Disruption in the Investment Management Industry. *Journal of Beta Investment Strategies* 13(1): 28-37.
- Colby, R. (2007). Evolution of Enhanced Analytical Indices. *ETFs and Indexing* 2007(1): 187-199.
- DeMiguel, V., L. Garlappi, F. Nogales and R. Uppal (2009). A Generalized Approach to Portfolio Optimization: Improving Performance by Constraining Portfolio Norms. *Management Science* 55(5): 798-812.
- Kelly, B.T., S. Pruitt and Y. Su (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics* 134(3): 501-524
- Maillard, S., T. Roncalli and J. Teïleteche (2010). The Properties of Equally Weighted Risk Contribution Portfolios. *Journal of Portfolio Management* 36 (4): 60-70.
- Malkiel, B. (2022). Jack Bogle: American Business Hero. *Journal of Beta Investment Strategies* 13(1): 14-18.
- Martellini, L. and V. Milhau (2018). Proverbial Baskets Are Uncorrelated Risk Factors! A Factor-Based Framework for Measuring and Managing Diversification in Multi-Asset Investment Solutions. *Journal of Portfolio Management* 44 (2): 8-22
- Qian, E. (2006). On the financial interpretation of risk contributions: Risk budgets do add up. *Journal of Investment Management* 4(4): 41-51.
- Roncalli, T. and G. Weisang (2012). Risk Parity Portfolios with Risk Factors. *SSRN Electronic Journal*. Working paper. http://dx.doi.org/10.2139/ssrn.2155159
- Ruggins, S. (2017). Building blocks: A historical sociology of the innovation and regulation of exchange traded funds in the United States, 1927-2000. University of Edinburgh doctoral thesis.
- Vaucher, B. (2023). Reconciling beta and fundamental approaches – an application of the IPCA equity model. Working paper (forthcoming).

Scientific Portfolio Publications

Scientific Portfolio Publications

2023 Publications

- Herzog, B., Jones, J., Safaee, S. Remember to Diversify Your Active Risk: Evidence from US Equity ETFs. (March).

2022 Publications

- Bouchet, V., Vaucher, B., Herzog, B. Look up! A Market-Measure of the Long-Term Transition Risks in Equity Portfolios. (December).

Disclaimer

Scientific Portfolio
An EDHEC Venture