

The Power of One (over N): Equal-Weighted Portfolios

Christopher Lewis | April 2023

Read Time: 20 mins

Key Takeaways

1. Equal-weighted portfolios have outperformed cap-weighted constructions across decades, factors, and countries. This alpha survives even after accounting for size, value, and momentum exposures.
2. For an active manager whose skill comes from picking winners and avoiding losers in their investment universe, equal weights generate the highest returns on that skill.
3. Equal weighting is not a free lunch. Implementing it in practice comes with liquidity and capacity challenges, tax consequences, and drawdown risks that managers must be willing and able to bear.

Summary

After an investment manager has decided which securities to hold in their portfolio, they face another important layer of decisions: what weights to assign those holdings. These assignments can be made piecemeal by discretion, collectively through a formula, or just about any hybrid in between. The optimal approach will vary across managers, based on their performance objectives and the type of investing skill they (might) have. For an investor focused on maximizing long-term returns, who is wary of assuming heroic amounts of manager skill, a surprisingly powerful approach is the simplest one: just assign the same weight to every holding. Equal-weighted portfolios have outperformed their cap-weighted counterparts in empirical analysis, they minimize the return impact of individual shock events, and they eliminate the need for precise estimates of future returns and correlations.

Background

While security selection is typically considered the cornerstone of successful active management, the weights assigned to holdings can also have a material impact on performance. The number of distinct weight combinations that could be applied to a portfolio is limitless, but there are a handful of systematic methods that stake out different extremes. Market cap-weighted portfolios generally minimize tracking error to a benchmark index, whereas equal-weighted portfolios minimize variation across the portfolio weights. Even more complex designs, which are often marketed as "optimization" tools are, by definition, attempting to locate some maximum or minimum along a metric of interest.

The first two schemas, cap weighting and equal weighting, are especially interesting because they are simple in design but otherwise diametrically opposed. The mathematics of compounding returns leads

to a natural skewness of market caps across firms¹. For that reason, it is not unusual to see a handful of "Mega Caps" carrying a large weight in the total market portfolio. The average investor always holds the total market portfolio (Sharpe, 1991), which makes cap weighting a rather obvious choice for benchmarks and passively minded products.

Equal weighting deviates from that passive inclination in the simplest manner. In fact, equal weighting is often referred to as the $1/N$ method – its complete formula for weight assignment (for N holdings in a portfolio). Yet despite its simplicity, equal-weighted portfolios consistently show higher returns in historical datasets. That advantage could be enhanced even further for an active manager whose skill comes from a classification ability to find outperformers and avoid underperformers in their universe. Sadly, these opportunities are not an arbitrage. The equal weight benefit that shows up on paper can be costly and prohibitive to implement, and as with traditional equity factors, equal weighting can suffer painful drawdowns along the way. Some might consider these issues to be flaws, but they are the very reasons that such a simple design feature could generate sustainable, long-term outperformance.

Higher Empirical Returns

The Russell 1000 Index

Different benchmark indices have varying levels of market cap skewness in their constituent weights. Large cap indices tend to be the most lopsided, since there is nowhere for their biggest firms to graduate up to. The Russell 1000 Index provides a compact example of this phenomenon, as it contains two sub-indices² that differ substantially in their weight concentration. The Russell Top 200 Index (stocks ranked 1-200 on market cap) carries a heavy weight in Mega Caps, whereas the Russell Mid Cap Index (stocks ranked 201-1000 on market cap) is more evenly distributed. These two components span the broader Russell 1000 Index of market cap ranks 1-1000.

An equal-weighted portfolio has no variation across its constituent weights, so an easy way to measure the degree of cap-weightiness for any portfolio is to calculate the standard deviation of its weights. Figure 1 shows this calculation for the three Russell 1000 pieces (right hand side) alongside their historical compound returns (left hand side). Across 37 years, the more equal-weighted Mid Cap Index has outperformed the Top 200 Index by nearly 1% per year. This spread is economically meaningful but somewhat unscientific, as there are plenty of other biases or characteristic differences which might explain the return gap. The size premium (Banz, 1981) or value premium (Fama and French, 1993) are


¹ If individual stock returns behave like Brownian Motion, then prices will follow right-skewed, Lognormal diffusion processes.

² Because of Russell's rebalancing rules there may not always be a perfect delineation at the 200th stock, but the segmentation of Large Cap = Mega Cap + Mid Cap is still a reasonable approximation.

two obvious candidates. However, even if those effects did account for the entire return difference, then leaning away from Mega Caps would still provide a simple way to capture some of those effects.

Figure 1: Large Cap Index Returns and Cap Weightiness

Compound annual returns from 1986 through January 2023 (left) and standard deviation of constituent weights at 1/31/23 (right). Both metrics in %s.



Source: Factset, FTSE Russell, and WEDGE Capital Management.


Equity Factor Returns

A more direct comparison of equal and cap weighting compares portfolio returns for the same constituents, weighted differently. Figure 2 graphs Sharpe Ratios at the factor level for four common equity factors: value, quality, investment, and momentum. The underlying factor returns themselves are long-short, meaning that in each case the top 30% of stocks ranked on that feature are "bought" while the bottom 30% of stocks are assumed to be sold short. With this construction the theoretical portfolio has zero net dollars invested, which makes any factor return that isn't 0% noteworthy. Because this long-short construction may be too hard or costly to recreate in practice, and a 0% return benchmark will not grow over time, Figure 2 focused on Sharpe Ratios rather than compound returns.

All of the factors show positive Sharpe Ratios in Figure 2 (which is always reassuring to see), regardless of whether their long end follows the traditional, cap-weighted formula or an equal-weighted construction. However, in every instance equal weighting the long side improves the performance.

Figure 2: Long-Short Factor Performance

Sharpe Ratios based on monthly, simple returns, shown annualized.



Source: Kenneth R. French data library and WEDGE Capital Management.

Note: Long side of factor trade represents the top 30% of stocks, and short side the bottom 30% of stocks, sorted on each metric annually. Short side of the factor trade is cap-weighted in all cases. Large and small size groups are formed separately and then averaged for each of the long and short side, just as in the traditional HML value factor. Factor returns start in July 1951 (Value), July 1963 (Quality and Investment), and July 1927 (Momentum), with all series running through January 2023.

That factor returns can be raised with $1/N$ weights is not some new epiphany. Fama and French surely knew this thirty years ago when they introduced their three-factor model (Fama and French, 1993). However, the purpose of academic factor models is to remove alpha – not create it. In that arena, a good factor model is one that can reattribute return patterns that might look like alpha as beta exposures to the factors. Finding the highest factor returns is not necessary to make that transference; what matters more is that the factors' fluctuations mimic common sources of return variation across assets. Since the average investor is always cap-weighted, mirroring that property for factor weights should cover more ground in explaining away potential alpha sources.

The comparison of equal versus cap weights can be isolated to just the long side of the factor trade too. Table I summarizes metrics for the return *difference* between equal-weighted and cap-weighted long factor portfolios (i.e. the stocks that the traditional factors "buy"). So once again, long-short portfolios have been formed, but this time the returns for each month t represent $R_{long,equal,t} - R_{long,cap,t}$. To help control for size effects, Table I is split into two panels, where the factors are formed separately within large and small size universes.

The Power of One (over N): Equal-Weighted Portfolios

Once again, the equal weight construction shows higher average returns, but here they appear across both large and small size classes. On paper the return advantage looks stronger for smaller caps, but in practice higher trading costs will eat into that gap. Small Cap quality is the black sheep of the group; it shows only a negligible benefit from equal weighting the higher quality firms.

Table I: Equal minus Cap-Weighted Return Spreads (Long Side Only)

Equal weight minus cap weight portfolio returns of the same, long-only constituents (top 30% of stocks sorted on each factor in each case). Monthly simple returns, shown in annualized %s.

	E/P (Value)	CF/P (Value)	OP/BV (Quality)	Asset Δ (Invest)	Prior Ret (Mom)
Years of Data	71.58	71.58	59.58	59.58	96.08
Panel A: Large Cap Stocks					
Avg Return	0.70	1.48	1.24	1.12	1.60
Std Deviation	5.51	5.80	6.32	5.99	5.42
Sharpe Ratio	0.13	0.26	0.20	0.19	0.30
t-stat	[1.07]	[2.16]	[1.51]	[1.45]	[2.90]
Panel B: Small Cap Stocks					
Avg Return	1.07	1.65	0.01	3.50	2.11
Std Deviation	4.58	4.98	5.10	8.71	5.31
Sharpe Ratio	0.23	0.33	0.00	0.40	0.40
t-stat	[1.98]	[2.81]	[0.02]	[3.11]	[3.89]

Source: Kenneth R. French data library and WEDGE Capital Management.

Note: Long side of factor trade represents top 30% of stocks, and short side the bottom 30% of stocks, sorted on each metric annually.

Factor returns start in July 1951 (Value), July 1963 (Quality and Investment), and July 1927 (Momentum), with all series running through January 2023.

Although many of the factors show average equal-weighted return advantages above 1% per year, the volatility around those averages is much higher. The factor Sharpe Ratios capture this tradeoff well: equal weighting may be a useful return enhancement feature for a portfolio strategy, but it is not a silver bullet that can "solve" active management single-handedly.

A New Equal-minus-Cap "Factor"

Taking the next step forward, the Table I factor returns can be averaged into a single, composite factor of equal-minus-cap returns (called "EMC"). Formally, the EMC return for month t can be defined for any number of K different factors and the two size segmentations, s , within each factor:

$$R_{EMC,t} = \frac{1}{2K} \sum_{k=1}^K \sum_{s=1}^2 R_{long,equal,k,s,t} - R_{long,cap,k,s,t}$$

In the following analysis, two different Value factors are included to better reflect how a live, quantitative portfolio might look in practice. Value tends to be a leading focus even in multifactor models, in part because its lower turnover makes it more practical to implement.

Creating this EMC factor allows the return advantage of equal weighting to be properly tested against factor models like those described on page 4. Those academic models all put forth the same challenge: are a strategy's higher returns just a recycling of other, already known, factor exposures? For studying the $1/N$ approach, size, value, and momentum factors are important to consider because of the trading steps needed to maintain equal weights. Recent winners are always trimmed, and recent losers automatically topped up, which indicates that equal weighting will lean towards smaller, value stocks, and away from those with high price momentum.

Table II puts EMC through the gauntlet of academic factor models: the basic CAPM, plus the three, four, five, and six-factor models which appear with regularity in research papers³. Each row runs a time series regression of EMC's returns on those factor returns, using monthly data back to 1963. In each of the five tests EMC maintains a positive alpha, and in all but one case that alpha is statistically significant. The one instance where alpha falls short of significance⁴ (the three-factor model) is caused by momentum's absence, which makes sense given EMC's anti-momentum trading activity, which the four and six-factor models confirm. The broader conclusion, therefore, is that equal weighting's return advantage over cap weighting looks like a separate form of alpha which cannot be explained by existing equity factors – including the size factor!

The results are even stronger in international returns. Table V in the appendix replicates Table II but using a Developed ex-US dataset. The length of history to study is more limited there, although the three decades that are available provide further validation that there is more to $1/N$ weighting than just playing the size premium.

³ See Sharpe, 1964; Fama and French, 1993; Carhart, 1997; and Fama and French, 2015 for the original papers.

⁴ Using the rule of thumb that a t-statistic > 2 corresponds to a 5% likelihood of declaring nonzero alpha when in truth it is zero.

Table II: Factor Loadings for Equal Minus Cap Composite

Regressions of EMC factor on domestic, monthly factor returns from July 1963 through January 2023, with loadings (i.e. betas) shown in monthly %s.

Model	Intercept (Alpha)	RMRF (Market)	SMB (Size)	HML (Value)	UMD (Mom)	RMW (Quality)	CMA (Invest)	Adj R ²
CAPM	0.10	0.05	-	-	-	-	-	0.04
t-stat	[2.42]	[5.48]	-	-	-	-	-	-
3 Factor	0.05	0.02	0.23	0.08	-	-	-	0.42
t-stat	[1.53]	[2.01]	[20.95]	[7.53]	-	-	-	-
4 Factor	0.14	-0.01	0.23	0.05	-0.11	-	-	0.58
t-stat	[5.02]	[-1.10]	[24.30]	[4.68]	[-16.34]	-	-	-
5 Factor	0.09	0.00	0.21	0.08	-	-0.06	-0.06	0.43
t-stat	[2.67]	[0.55]	[18.61]	[5.16]	-	[-4.15]	[-2.62]	-
6 Factor	0.16	-0.01	0.22	0.02	-0.11	-0.04	-0.02	0.58
t-stat	[5.68]	[-1.85]	[22.01]	[1.60]	[-15.95]	[-3.29]	[-0.97]	-

Source: Kenneth R. French data library and WEDGE Capital Management.

Note: Composite EMC factor is constructed as the average across long-only, equal minus cap-weight formations used in Table I.

Why are Returns Higher?

Perhaps the next logical question to ask is why; what is it about equal weighting that would generate higher returns like this? Individual stocks do not sit on a spectrum of equal weightiness like they do for traditional factor metrics, so there cannot be a "this type of stock" behavioral bias or risk story at work.

Causality questions are always hard to answer, but there are two theoretical avenues that support equal weighting. First, equal weights squeeze the maximum benefit out of classification skill – the manager's (or factor's) ability to separate winners from losers in their investment universe. Secondly, equal weights minimize the return impacts of individual, unpredictable events that strike portfolio holdings.

Maximize Classification Skill

The first theoretical advantage is that $1/N$ avoids the need to distinguish securities beyond desirable and undesirable. Everything that makes the cut to be owned is implicitly assumed to have the same potential. If the manager's security selection abilities are limited to this type of classification decision, then an equal-weighted portfolio deploys all of her skill, completely, without reaching for extra sophistication where no skill exists. If the manager can identify "high conviction" ideas accurately, then different weights might make sense. However, as long as skill is limited to winner versus loser signals, any deviation from equal weights tacks on extra variance without an expected returns benefit.

Recent research has formalized this definition of winner/loser skill as "Omega" (Greek letter ω) and put it through a battery of tests (Bolshakov, Chincarini, and Lewis, 2022). One test simulates how a manager with Omega skill would perform against a cap-weighted benchmark, and whether they should assign equal or cap weights to their holdings. Table III summarizes the results from that paper's analysis. The scale of returns shown is arbitrary, but across all the different size leanings that the manager's benchmark could experience, the Omega manager achieves higher average returns by equal weighting. Higher average returns do come at a cost though: the tracking error across regimes is much higher too.

Table III: Simulated Return Outcomes for Omega Skill

Arbitrary units of $\pm 10\%$ individual security returns, assigned to actual S&P 500 constituent weights at 12/31/18.

	Cap-Weighted Benchmark Regime					Return Statistics	
	Favor Large Caps		Neutral to Size	Favor Small Caps		Across Regimes	
	Extreme	Mild		Mild	Extreme	Average	Std Dev
Benchmark Return	7.15	1.25	-	-1.25	-7.15	-	5.13
Panel A: Average Returns							
Equal-Weighted	0.20	0.20	0.20	0.20	0.20	0.20	0.00
Cap-Weighted	7.24	1.44	0.20	-1.06	-7.05	0.16	5.13
Panel B: Average Relative Returns							
Equal-Weighted	-6.96	-1.05	0.20	1.45	7.35	0.20	5.13
Cap-Weighted	0.09	0.19	0.20	0.20	0.11	0.16	0.05

Source: "Enhanced Indexing and Selectivity Theory" (2021; see Table 7) and WEDGE Capital Management.

Note: Each data point is an average across 100,000 Monte Carlo simulations of a manager with Omega skill of $\omega=1.1$ selecting 395 stocks out of a 500-stock benchmark, where half the constituents return +10% and half return -10%. As the source paper describes, these simplifying assumptions make for easy, one-period experiments, but can be relaxed without altering the conclusions. None of these "portfolios" are investable in practice.

Minimize Unpredictable Events

Related to the idea of Omega skill, $1/N$ minimizes the return impact of individual, unpredictable events. Even for a highly skilled manager, unpredictable things inevitably happen. In economics these occurrences are often called "shocks," and they can be good or bad. M&A announcements, natural disasters, fraud, etc. can all strike a portfolio holding without warning. Equal weighting minimizes the return impact that independent shocks like that have on portfolio returns. Otherwise, when the manager wakes up in the morning and learns that some event wiped out a portfolio holding overnight, they face a second layer of anxiety – was that one of their bigger positions or a smaller one? Over time, the volatility of whether shocks strike large or small positions puts extra variance into the portfolio's returns, which without an alpha benefit to compensate for it, lowers the portfolio's compound returns.

The reasoning above might make intuitive sense already, but there is a deeper framework at work. Entropy is a scientific measure of the amount of disorder in a system. Although "disorder" might sound like a bad quality to have, it helps protect against new shocks, as they will cause less *incremental* disorder to a system with high entropy already. This concept can be defined mathematically, and solved for optimal choices based on different circumstances. For an investment manager whose only constraint is long-only portfolio weights, *information* entropy is maximized by equal weighting their portfolio. Proof of this statement is left in the appendix for math lovers, but if more constraints or granular foresight were assumed, then the optimal weight formula would recalibrate accordingly.

This last point is important: cap weighting is not the only alternative out there, despite being the main antagonist of this paper. All kinds of fancy "optimization" programs exist. However, they all require a considerable number of estimations or assumptions, which if wrong, can detract from the manager's underlying skill. In 2009 a group of researchers⁵ conducted a horse race test of some of these more complex weighting schemes, and reached the following conclusion:

Of the 14 models we evaluate across seven empirical datasets, none is consistently better than the 1/N rule in terms of Sharpe ratio, certainty-equivalent return, or turnover, which indicates that, out of sample, the gain from optimal diversification is more than offset by estimation error.

They go on to specify that it would take centuries of data to overcome the problem of estimation error:

...the estimation window needed for the sample-based mean-variance strategy and its extensions to outperform the 1/N benchmark is around 3000 months for a portfolio with 25 assets and about 6000 months for a portfolio with 50 assets.

How can this Persist?

A good sanity check for any alpha-seeking decision is to ask, "why does this opportunity exist for me?" If the case for equal weighting is so compelling and simple, then surely others would adopt it, which would impact market prices, and in turn dissolve the strategy's higher returns. Fortunately, there seem to be enough drawbacks and practical challenges to equal weighting that some investors will always be deterred from it. Ironically, it is these very drawbacks that could support sustained outperformance for 1/N into the future.

This type of circular logic can be exasperating: in daily life, how often is it declared that *fortunately* there are *drawbacks* to something? Yet in active investment management, where the rest of the market is competing for alpha and studying the past data too, identifying drawbacks can be cathartic for the

⁵ See DeMiguel, Garlappi, and Uppal: Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?

investor capable of bearing them. In the case of equal weighting, return drawdowns, trading costs, and marketing difficulties can all introduce challenges that may be insurmountable for some managers and investors, but acceptable for others.

Return Drawdowns

The boxer Mike Tyson once declared that "everyone has a plan until they get punched in the mouth." Just like with traditional factors, EMC's higher average returns can experience painful drawdowns along the way. Seeing these drawdowns on paper is one thing, but living through them in real time, when the recovery on the other side is not visible, can feel like a punch to the face. Even if the investment manager has the conviction to ride out these periods, their clients may not. Thus, for equal weighting to work in practice, the manager needs both a good implementation plan as well as the ability to defend the approach when its short-term track record is poor.

Table IV shows the five deepest and longest drawdowns that the EMC factor has experienced over the past six decades. Many of its worst stretches have overlapped with crisis periods in the market (e.g. the Great Financial Crisis, Savings and Loan Crisis, Black Monday crash, Covid-19 crash), which suggests there could be a risk aversion effect at work. For investors who value "downside protection" more than the same amount of "upside delivery," then cap weighting may provide them a higher utility, and a reason to (rationally) avoid the higher expected returns of equal weighting.

Table IV: Historical Drawdowns for Equal Minus Cap Composite Factor

Composite factor constructed as the average across equal minus cap-weight constructions of the long side of the typical quant trade, as shown in Table II. Returns below shown as cumulative %s.

Five Worst Drawdowns			Five Longest Drawdowns		
Start	Low	Maximum	High Water Mark		Months
Date	Point	D.D. Ret	Prior	Next	of D.D.
2006-05	2008-11	-17.97	1983-07	1994-02	126
1983-07	1990-12	-16.75	1969-01	1976-02	84
2014-02	2020-03	-15.54	2014-02	2021-02	83
1969-01	1973-06	-13.85	1996-05	2001-10	64
1996-05	1998-12	-12.60	2006-05	2009-09	39

Source: Kenneth R. French data library and WEDGE Capital Management.

Note: Composite EMC factor is constructed as the average across long-only, equal minus cap-weight formations used in Table I.

The length of drawdowns could also form a bifurcation across investors. EMC has experienced multiple drawdowns that took more than five years to dig out of. Even though five years is not that long from a quantitative perspective, it could certainly put patience to the test and cause some to abandon the

strategy. Thankfully, using equal weighting as a feature within a broader investment process should offer some diversification benefits, and help reduce the drawdowns described in Table IV.

Increased Costs

Pursuing an equal weight approach can be frustrating because the live holdings are always departing from the theoretical $1/N$ weights. Staying close to equal weighting incurs trading costs that do not exist for market cap weights. These costs can overwhelm the paper return advantage if the manager is not careful: even a monthly $1/N$ rebalancing of the S&P 500 would incur about 50% annual turnover, compared to under 5% turnover for the cap-weighted S&P 500 (Blitzer, Chincarini, and Dash, 2003).

Individual liquidity circumstances can be problematic too. Even within large cap equities it is not uncommon to see a few individual stocks trading under \$5 million per day, especially around holidays. Equal weighting a position with that amount of liquidity could be problematic for even a well-diversified manager⁶. In smaller cap universes the liquidity challenges are even more daunting. Once again, the manager must have a system in place to either sidestep these potholes or adjust their weights accordingly.

Thankfully, there is evidence that these implementation costs can be manageable. In 2003 when Rydex Global Advisors (now Invesco) launched the Equal Weight S&P 500 Index, they also started a live, investable product⁷. Over the twenty years since its inception, that product has lagged the free, paper equal weight index by about 0.5% per year, but is still outperforming the cap-weighted S&P 500 Index by over 1% per year.

Marketing Difficulties

Finally, equal weighting can cause headaches within the investment manager's own operations. The turnover and liquidity challenges put an effective ceiling on how high product AUM can grow, whereas cap weighting has almost limitless growth potential. These capacity limits get even more restrictive in smaller asset classes. In addition, all the rebalance trading needed to maintain $1/N$ racks up short-term capital gains, so if taxable and tax-exempt clients are comingled together in a fund structure, equal weighting may be undesirable.

⁶ As an illustration, assume the manager wants to limit their daily participation rate to 10% of market volume, and holds about 100 stocks. In that scenario, each \$1 billion of AUM would consume 20 trading days to buy or sell a \$5 million ADV security.

⁷ This product is now an ETF trading under the symbol 'RSP.' As of 2/28/23, RSP's since inception NAV return is 10.92%; the Equal Weight S&P 500 Index return is 11.38%; and the S&P 500 Index return is 9.85% (all annualized).

Concluding Remarks

On paper, the performance benefits of using the $1/N$ method for portfolio construction look quite attractive. Empirical data shows higher long-term returns and Sharpe Ratios for weighting the same indices and factor constituents more equally. This performance advantage appears across factors, decades, and countries.

Equal weighting is not an arbitrage though. It behaves and feels a lot like a risk premium, akin to the traditional equity factors, and can introduce practical drawbacks for a manager. Ironically, drawbacks can be a good thing, as they build a rationale for why outperformance could persist into the future. In this type of iterative thinking, the manager (or client) must ask themselves if the pursuit of equal weighting is worth it and appropriate for them specifically. What are their product capacity and tax efficiency goals? How will they react if equal weighting underperforms during a market crisis?

The choice between equal and cap weighting doesn't have to be binary either. One could certainly lean into the concept by using a blend of equal and cap weights, imposing equal weights only within particular sector groups, or just about any other mixture for that matter. All these options are still active deviations, since the average investor always holds the cap-weighted market portfolio.

Disclaimer

This paper is being made available for educational purposes only and should not be used for any other purpose. The information contained herein does not constitute, and should not be construed as, an offering of advisory services or an offer to sell or solicitation to buy any securities or related financial instruments. Certain information contained herein is based on or derived from information provided by independent third-party sources and may be linked to third-party sources. WEDGE Capital Management L.L.P. "WEDGE" believes that the sources from which such information has been obtained are reliable; however, it cannot guarantee the accuracy of such information and has not independently verified the accuracy or completeness of such information.

This paper expresses the views of the author as of the date indicated and such views are subject to change without notice. WEDGE has no duty or obligation to update the information contained herein. Further, WEDGE makes no representation and it should not be assumed that past investment performance is an indication of future results. Further, wherever there is a potential for profit, there is also the probability of loss.

This paper, including the information contained herein, may not be copied, reproduced, republished, or posted in whole or in part, in any form without the prior written consent of WEDGE.

References

- Banz, R. (1981). The Relationship between Return and Market Value of Common Stocks. *Journal of Financial Economics*, 9, 3-18.
- Blitzer, D., Chincarini, L., and Dash, S. (2003). The S&P Equal Weight Index: Uses, Properties, and Historical Simulations. Available at: <https://ludwigbc.com/pubs/SPEWIWhitePaper010703.pdf>.
- Bolshakov, A., and Chincarini, L. (2020). Manager skill and portfolio size with respect to a benchmark. *European Financial Management*, 26 (1), 176-197.
- Bolshakov, A., Chincarini, L., and Lewis, C. (2022). Enhanced indexing and selectivity theory. *European Financial Management*, 28 (4), 964-998.
- Carhart, M. (1997). On Persistence in Mutual Fund Performance. *Journal of Finance*, 52 (1), 57-82.
- DeMiguel, V., Garlappi, L., Uppal, R. (2009). Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? *The Review of Financial Studies*, 22 (5), 1915-1953.
- Fama, E., and French, K. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33, 3-56.
- Fama, E., and French, K. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116, 1-22.
- Fama, E., and French, K. (2018). Choosing Factors. *Journal of Financial Economics*, 128 (2), 234-252.
- French, K. (2023). Kenneth R. French – Data Library. Retrieved from: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.
- Gibbons, M., Ross, S., and Shanken, J. (1989). A Test of the Efficiency of a Given Portfolio. *Econometrica*, 57 (5), 1121-1152.
- Khang, K., Picca, A., Zhang, S., Zhu, M. (2021). How Inefficient is the 1/N Strategy for a Factor Investor? *Journal of Investment Management* (forthcoming). Available at SSRN: <https://ssrn.com/abstract=3972442> or <http://dx.doi.org/10.2139/ssrn.3972442>.
- Shannon, C. (1948). A Mathematical Theory of Communication. *The Bell System Technical Journal*, 27, 379-423.
- Sharpe, W. (1964). Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. *Journal of Finance*, 19 (3), 425-442.
- Sharpe, W. (1991). The Arithmetic of Active Management. *Financial Analysts Journal*, 47 (1), 7-9.

Appendix

International Evidence

Table V replicates the regressions of Table II using a global data set of developed, ex-US returns. The length of data history is more limited here, but the alpha results are even stronger.

Table V: International Factor Loadings for Equal Minus Cap Composite

Regressions of developed ex-US EMC factor on monthly factor returns from November 1990 through January 2023. Loadings shown in monthly %s.

Model	Intercept (Alpha)	RMRF (Market)	SMB (Size)	HML (Value)	UMD (Mom)	RMW (Quality)	CMA (Invest)	Adj R ²
CAPM	0.27	-0.00	-	-	-	-	-	-0.00
t-stat	[6.36]	[-0.32]	-	-	-	-	-	-
3 Factor	0.24	0.02	0.26	0.05	-	-	-	0.36
t-stat	[7.09]	[2.19]	[14.73]	[3.58]	-	-	-	-
4 Factor	0.32	-0.01	0.27	0.01	-0.09	-	-	0.49
t-stat	[10.25]	[-0.80]	[17.15]	[0.80]	[-9.85]	-	-	-
5 Factor	0.27	-0.00	0.24	0.07	-	-0.08	-0.14	0.40
t-stat	[7.64]	[-0.52]	[14.17]	[3.83]	-	[-2.92]	[-5.42]	-
6 Factor	0.32	-0.02	0.26	0.03	-0.09	-0.03	-0.11	0.51
t-stat	[9.86]	[-2.06]	[16.59]	[1.77]	[-9.17]	[-1.00]	[-4.59]	-

Source: Kenneth R. French data library and WEDGE Capital Management.

Note: Composite EMC factor is constructed as the average across long-only, equal minus cap-weight formations as described for Table II. However, the developed ex-US dataset contains only one value factor – Book-to-Market – so that value factor alone is used in this international construction of EMC.

Entropy Maximization

To prove that equal weighting maximizes information entropy for an unconstrained long-only portfolio, consider a simple optimization problem: design a function $p(x)$ that maximizes aggregate information entropy (itself defined by the log-likelihood function), such that the cumulative density sums to 1.

Translating this description to investment portfolios, $p(x)$ would be the formula that assigns a weight to security x , and for a long-only portfolio, the total weight across holdings must add up to 100%.

$$\begin{aligned} \max_{p(x)} & - \int_a^b p(x) \ln p(x) dx \\ \text{s. t.} & \int_a^b p(x) dx = 1 \end{aligned}$$

The lack of additional constraints, such as specified return or covariance requirements, reflects the limited forecasting power available in active management. If investors could know precisely all

expected returns and their massive covariance matrix, then more complexity could be solved for. Sadly, it is hard enough to find alpha sources that are positive, let alone being able to precisely pinpoint those alphas and the thousands of cross-correlations between securities.

Thus, for the problem at hand with just the total weight constraint, a Lagrangian function only needs a single multiplier λ to set up a solvable system of equations across the partial derivatives:

$$\mathcal{L}(p(x), \lambda) = - \int_a^b p(x) \ln p(x) dx + \lambda \left(\int_a^b p(x) dx - 1 \right)$$

First, differentiate with respect to $p(x)$ and set that equal to zero. Although $p(x)$ is an unknown function, it can be thought of as many discrete points with their own – potentially unique – derivatives which must all equal zero. However, it quickly becomes apparent that those derivatives are the same for any x . Use the Chain Rule and the fact that the derivative of $\ln(u)$ is $1/u$ to get the following:

$$\frac{\partial \mathcal{L}}{\partial p(x)} = -\ln p(x) - 1 + \lambda = 0$$

$$\ln p(x) = \lambda - 1$$

$$p(x) = e^{\lambda-1}$$

Already the conclusion is coming into focus. $p(x)$ depends only on λ , which is unknown but static across x values. Therefore all $p(x)$ outputs must be the same. Nevertheless, one can complete the exercise by differentiating with respect to λ , and arriving at the uniform distribution as the optimal function.

$$\frac{\partial \mathcal{L}}{\partial \lambda} = \int_a^b p(x) dx - 1 = 0$$

$$1 = \int_a^b e^{\lambda-1} dx$$

$$1 = (b - a)e^{\lambda-1}$$

$$\frac{1}{b - a} = e^{\lambda-1}$$

$$p(x) = \frac{1}{b - a}$$

The uniform distribution can be confirmed as the maximum since the second derivative of the objective function is negative whenever $p(x)$ is positive, which is a requirement of any logarithmic function.

$$\frac{\partial^2 \mathcal{L}}{\partial p(x)^2} = -\frac{1}{p(x)}$$

Spanning Regressions

Table VI explores spanning tests of seven domestic factors: the six used as right-hand-side (RHS) variables in Table II, plus the EMC factor. In this setup seven different regressions are run along the horizontal rows, in each case with that particular factor shifted to the left-hand-side (LHS) of the equation, while the other six remain as RHS variables. The bottom row of Table VI, with EMC on the LHS, is the same regression as the bottom row of Table II. A nonzero alpha indicates that the LHS factor cannot be replicated by the other factors, and therefore seems useful to retain on the RHS for future tests on other asset returns. An insignificant alpha suggests that the LHS factor could be redundant.

For brevity, only coefficients on alpha and EMC are shown in Table VI. The $a^2/s^2(e)$ column divides the squared alpha by that regression's variance of residuals. This metric measures how impactful including the intercept is compared to the unexplained variance left behind, which has implications for other model-selection metrics (namely the GRS test) that are not pursued here.

Table VI: Spanning Regressions of EMC plus Six Traditional Factors

Regressions of monthly factor returns from July 1963 through January 2023, with loadings (i.e. betas) shown in monthly %s.

LHS Variable	Intercept (Alpha)			EMC Loading		Adjusted R ²		
	Value	t-stat	$a^2 / s^2(e)$	Value	t-stat	w/ EMC	w/o EMC	Change
RMRF	0.93	[5.93]	0.06	-0.38	[-1.85]	0.23	0.23	0.00
SMB	-0.14	[-1.64]	0.00	1.88	[22.01]	0.51	0.17	0.33
HML	-0.02	[-0.19]	0.00	0.17	[1.60]	0.52	0.52	0.00
UMD	0.91	[6.70]	0.07	-2.47	[-15.95]	0.34	0.10	0.24
RMW	0.41	[5.13]	0.04	-0.34	[-3.29]	0.18	0.17	0.01
CMA	0.24	[4.31]	0.03	-0.07	[-0.97]	0.54	0.54	-0.00
EMC	0.16	[5.68]	0.05	-	-	-	0.58	-

Source: Kenneth R. French data library and WEDGE Capital Management.

Note: Composite EMC factor is constructed as the average across long-only, equal minus cap-weight formations used in Table I. Adjusted R-Squared w/o EMC re-runs each regression without the EMC factor and reports that value.

The HML value factor stands out as the least additive, which aligns with existing research from Fama and French (2015 and 2018). More interestingly, adding EMC into the candidate pool reduces the size factor's uniqueness too, whereas it maintained a strong prominence in those prior studies. Particularly if the purpose of the factor model is to explain live manager returns, where AUM levels might influence weighting decisions, replacing the traditional size factor with some version of an EMC factor might make sense. For testing anomalies in the cross-section of stock returns though, EMC is nonsensical and should be avoided, given that it has no interpretability as a characteristic that varies across firms.