CHAPTER 10

Benchmark-Relative and Absolute-Return Are the Same Thing: Conditions Apply

Robert Scott

10.1 Portfolio Objectives

Although they are both trying to maximise return for a given level of risk, benchmark-relative and absolute-return managers adopt different means for getting there. The benchmark-relative manager is maximising alpha, or return relative to the benchmark, subject to a tracking error limit, while the absolute-return manager is maximising total return subject to some risk limit such as a probability of loss or absolute volatility and so on. It might seem that the benchmark-plus-alpha that a benchmark-relative manager generates should be similar in magnitude to that which an absolute-return manager might deliver (at least during an up market); however, it turns out that this is not necessarily the case.

Our investigation is simulation-based, examining market views and optimal portfolios to test the impact of different investment objectives. The details of the simulation procedure can be found in the appendix and the specific objective functions and constraints for each strategy can be found in Table 10.1. Suffice it to say here that we have assumed that the

R. Scott (✉)
Official Institutions, Schroder Investment Management, London, UK
e-mail: robert.scott@schroders.com

N. Bulusu et al. (eds.), *Advances in the Practice of Public Investment Management*, https://doi.org/10.1007/978-3-319-90245-6_10

Table 10.1 Formal optimisation problems for absolute-return and benchmark-relative

Strategy	*Optimisation problem*
Absolute-return	Maximise $R_p = \beta_p F$ Subject to: $\sigma_{p,i} \leq$ Target max risk $i = 1 \ldots n$ where R_p is the return on the portfolio, β_p is the set of factor sensitivities in the portfolio, and F is the set of factor returns
Benchmark-relative	Maximise $\alpha_p = (\beta_p - \beta_{BM})F$ Subject to: $TE_p \leq$ Target max TE where α_p is the excess return of the portfolio over the benchmark and β_p and β_{BM} are the factor sensitivities in the portfolio and benchmark
Benchmark-relative, beta constrained	Maximise $\alpha_p = (\beta_p - \beta_{BM})$F Subject to: $TE_p \leq$ Target max TE $beta_p = 1$ where $beta_p$ is the overall beta of the portfolio relative to the benchmark (covariance divided by benchmark variance)
Benchmark-relative, risk capped	Maximise $\alpha_p = (\beta_p - \beta_{BM})F$ Subject to: $TE_p \leq$ Target max TE $\sigma_p \leq \sigma_{BM}$ where σ_p and σ_{BM} are the volatilities of the portfolio and benchmark

job of the benchmark-relative manager is to maximise their information ratio, and to apply this process to a multitude of similar portfolios, all with differing benchmarks and risk limits. For any one portfolio, the alpha is maximised subject to a tracking error (or other) limit. The alpha can be generated in the purest form by either market timing the beta (or betas) in the portfolio or through security selection.

In principle, the mandate can have a very low or a very high tracking error—there is nothing intrinsic to benchmark-relative investing that requires low tracking error. Absolute-return seeks to generate the highest return in the portfolio subject to a given risk limit, often captured as a measure of the probability of loss, or the likely frequency of losses over a particular horizon. By definition, for absolute-return investing, there is either no benchmark, or a margin over cash (or zero) is considered to be the benchmark. There is no requirement for return to be generated from either a single or multiple asset classes. We could therefore categorise absolute-return mandates into both single asset class (constrained) and multiple asset classes (unconstrained).

While maximising the Sharpe ratio or the information ratio might sound like very similar things, in fact, the process of maximising the information ratio does not deliver the highest possible Sharpe ratio for the end investor (see Roll 1992). For this reason, the opportunity set of possible returns for active investors are better under an absolute-return mandate than for a typical benchmark-relative strategy. This is true, so long as the benchmark is not mean-variance efficient: in other words, if the benchmark is not constructed by maximising returns as a function of risk. There is theoretical and empirical evidence in support of capitalisation-weighted benchmarks being inefficient (see Haugen and Baker 1991, 2010). The process of achieving the highest information ratio incentivises the portfolio manager to create portfolios that effectively "leverage" the beta in the benchmark to some degree as we show later. The end result is a higher information ratio, but a sub-optimal Sharpe ratio. Figure 10.1 shows the possible portfolios available to the investor for a given set of expected returns and risk tolerances: either total risk for absolute-return or tracking error for benchmark-relative. These possible portfolios are based on hypothetical risky assets with characteristics described in greater detail in the appendix. The portfolios constructed are based on maximising the expected return for absolute-return and expected alpha for benchmark-relative for the same set of expected asset returns. The only constraints for these initial portfolios are the risk limits (either total volatility or tracking error). If the benchmark-relative investor maximises alpha for a given level of tracking error, their resulting portfolio lies below an absolute-return portfolio with similar risks. Put another way, the Sharpe ratio is lower. Table 10.2 details some of the characteristics of the portfolios used to create the previous charts. In column one, we show that the benchmark is designed to have factor sensitivities, perhaps beta and duration derived from stocks and bonds. The next column makes clear that the absolute-return portfolio with similar risk levels to the benchmark has a higher expected return, this because it is constructed so as to be mean-variance efficient. The following two columns show some sample benchmark-return portfolios with different levels of tracking error. They are constructed to maximise the expected alpha subject only to the tracking error limit. Note the betas are fairly high, and the correlation between alpha and beta is also quite high. The information ratios, however, are the highest of all sample portfolios whereas the Sharpe ratios are among the lowest.

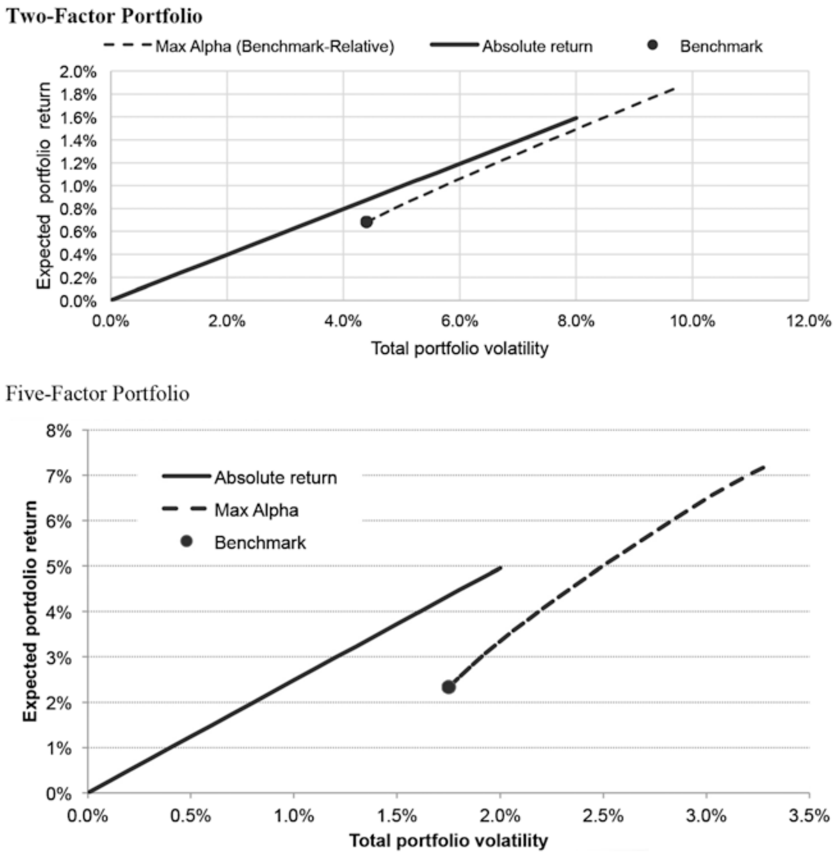

Fig. 10.1 Traditional benchmark-relative approaches lag absolute-returns for two- and five-factor portfolios. Two-Factor Portfolio

To improve the Sharpe ratio, an additional incentive is needed to induce the benchmark-relative active manager to improve the end investor's overall return for a given level of overall risk. One very effective method, we will argue, is for the investor to actually increase the constraints in the mandate.

Table 10.2 Sample portfolios under various constraints—two-factor model

	Benchmark	*AR risk = 4.4%*	*Max alpha TE = 90 bps*	*Max alpha TE = 3%*	*Beta = 1 TE = 90 bps*	*Beta = 1 TE = 3%*	*Risk < BM TE = 90 bps*	*Risk < BM TE = 3%*
Equity β	0.6	0.2	0.6	0.7	0.5	0.3	0.5	0.2
Bond β	2.0	4.6	2.9	5.1	3.0	5.3	2.9	4.6
α β correl		−33%	78%	78%	0%	0%	−10%	−33%
Beta	1	0.78	1.16	1.53	1.00	1.00	0.98	0.78
Total risk	4.4%	4.4%	5.1%	7.0%	4.5%	5.3%	4.4%	4.4%
Exp. ret.	0.68%	0.87%	0.86%	1.28%	0.80%	1.06%	0.78%	0.88%
Sharpe ratio	0.54	0.69	0.58	0.63	0.61	0.69	0.61	0.69
Alpha		0.19%	0.18%	0.60%	0.11%	0.37%	0.10%	0.19%
Tracking error		2.91%	0.90%	3.00%	0.90%	3.00%	0.90%	2.92%
Info. ratio		0.23	0.69	0.69	0.43	0.43	0.37	0.23

Source: Author's calculations. Note: All data are annualised and based on hypothetical returns and risks. See appendix for details of the simulation

10.2 Adding Constraints to Improve Performance

Conventional investment doctrine suggests that relaxing constraints is a way to improve performance. To test that dictum, we introduced two possible constraints (which we discuss below) on the benchmark-relative portfolio construction. Predictably, they reduced the amount of expected alpha for a given amount of tracking error, as shown in Fig. 10.2. Here we

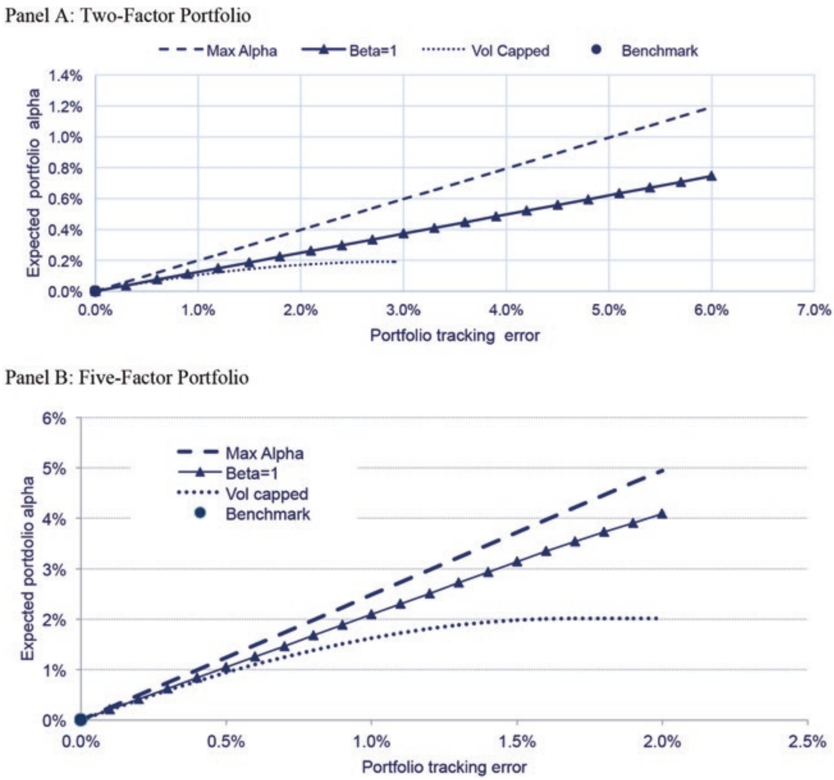

Fig. 10.2 Constraints tend to undermine information ratios—alpha and TEV for two- and five-factor portfolios

plot the unconstrained benchmark-relative optimal frontier from Fig. 10.1 in the space of tracking error vs. expected portfolio alpha, along with the frontiers for the two constrained portfolios, which we have called beta = 1 and vol-capped.

The first constraint we looked at, beta = 1, was originally proposed by Roll (1992) and can be formally defined in Eqs. 10.1 and 10.2 as:

$$\beta = \frac{\sigma_P \sigma_{BM} \rho}{\sigma_{BM}^2} \tag{10.1}$$

where $\sigma_P \sigma_{BM} \rho$ is the covariance of the portfolio with the benchmark and $\sigma^2{}_{BM}$ is the variance of the benchmark. In matrix terms using portfolio sensitivities, this is measured as:

$$\beta = {F_P \Sigma F'_{BM}} \big/ {\sigma_{BM}^2} \tag{10.2}$$

where F_P is the set of portfolio factor sensitivities or betas and Σ is the covariance matrix of factor variances.

This forces the beta of the portfolio and that of the benchmark to be the same, which makes intuitive sense on many levels. Most importantly, it forestalls any attempt to substitute beta returns for alpha by making the portfolio a leveraged version of the benchmark. Any alpha will therefore be the result of genuine skill in stock selection or market timing and will be uncorrelated with beta. In fact, many active managers proclaim their objective to provide "uncorrelated alpha" so the constraint is within the spirit of active management.

The resulting portfolios at different levels of tracking error deliver lower alpha (and hence lower information ratios), as shown by the line in Fig. 10.2, but the overall Sharpe ratio of the portfolio is improved, and the set of possible portfolios is more efficient in terms of risk and return, as shown in Fig. 10.3. The reason for the improvement is that the Sharpe ratio combines three elements: the Sharpe ratio of the benchmark, the information ratio of the portfolio, and an element that equates to the correlation between the two, beta and alpha. If this correlation falls, as it is forced to in the beta = 1 portfolio, then the risk also falls and the total risk-adjusted return (Sharpe ratio) goes up.

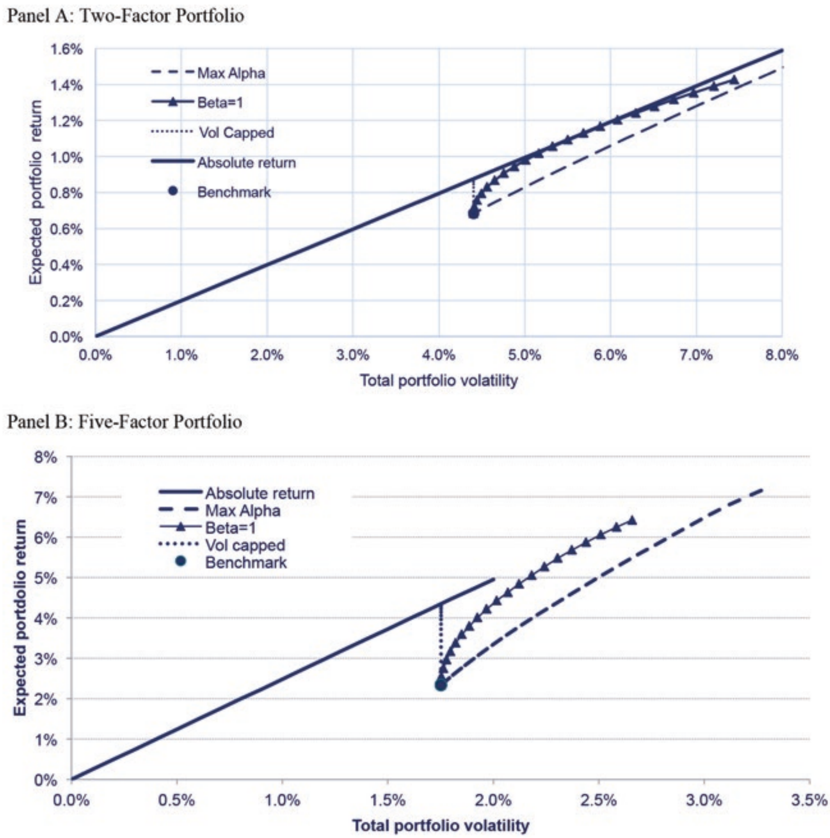

Fig. 10.3 Risk and return for two- and five-factor portfolios

Our second constraint is to restrict the total portfolio risk to a level no higher than that of the benchmark as originally proposed in Jorion (2003). We define the portfolio risk, benchmark risk, and the constraint as follows:

$$\sigma_p = \sqrt{F_P \Sigma F'_P} \quad (10.3)$$

$$\sigma_{BM} = \sqrt{F_{BM} \Sigma F'_{BM}} \quad (10.4)$$

$$F_P \Sigma F'_P - F_{BM} \Sigma F'_{BM} \le 0 \quad (10.5)$$

This is also intuitive since it allows active positions, so long as the overall portfolio risk is not increased. Alpha-beta correlation under this scenario is typically zero or negative, which is also an attractive quality. As in the previous case, the total alpha delivered is lower for the same amount of tracking error as compared to an unconstrained portfolio, that is, the information ratio falls. But, again, the overall risk-return characteristics are improved and the set of possible portfolios is more efficient than the unconstrained approach, that is, their Sharpe ratio goes up.

To summarise, the unconstrained approach delivers the highest alpha, but at the expense of overall portfolio efficiency, while the two constrained approaches deliver less alpha, but also much less risk, so that the overall risk-return profile is better. For a given amount of total risk for the end investor, the constrained and absolute-return approaches all deliver higher returns. It is also noteworthy that at a certain level of tracking error, the constrained portfolios are as efficient as the set of possible absolute-return portfolios. We will discuss this in more detail in the next section.

10.3 Convergence of Benchmark-Relative and Absolute-Return Portfolios

We have shown that an investor is better served in the mean-variance framework by introducing a constraint into their mandate, either requiring that beta be equal to one or alternatively that total portfolio risk is never more than benchmark risk. In this section we will show some examples of what representative examples of these portfolios might look like under varying tracking error assumptions. One point to note, however, is that all these hypothetical portfolios assume the investor will receive positive returns from their constituent risk factors. We will deal with bear-market scenarios in the next section.

Revisiting Table 10.2, it is useful to compare the previously described basic portfolios with the constrained benchmark-relative ones. The beta = 1 portfolios (columns 5 and 6) have lower information ratios, but higher Sharpe ratios, and—as discussed earlier—the alpha-beta correlation is zero.

The last two columns display two sample portfolios where the overall risk is limited to the benchmark level (σ) or below: one for a tracking error (TE) of 1% and a second for a tracking error of 3.0%. Like the beta = 1 portfolios, these have higher Sharpe ratios and lower information ratios. The portfolio sensitivities for the TE = 3.0% portfolio are highly significant: *they are identical to the sensitivities of the absolute-return portfolio*. Put another way, a benchmark-return manager, operating within a tracking error and total risk constraint, while maximising alpha, has created an identical portfolio to that of an absolute-return manager (the two shaded columns). One final note. It is possible to show the same convergence for a beta = 1 portfolio, although at a much higher level of tracking error.

Thus far, we have demonstrated that it is possible to constrain a benchmark-relative manager in such a way that it induces them to improve the overall Sharpe ratio of their portfolio. In doing so, the portfolio ends up with identical characteristics to that of an absolute-return manager. However, there is one important proviso: returns for the risk factors must be expected to be positive. In an upcoming section, we will look at how a bear-market scenario affects these conclusions. Before turning to this point, however, the question arises as to how an investor can identify the amount of tracking error necessary to allow the portfolio exposure to be the same as the absolute-return portfolio. We will address this in the next section.

10.4 Identifying Optimal Tracking Error Levels

Figure 10.3 and Table 10.2 show that at some level of tracking error a constrained-alpha maximisation strategy will produce portfolios identical to absolute-return portfolios. The question arises as to what is the determinant of the required level of tracking error. We can borrow from Scott (2011) for the answer for this. Using the simulations from Fig. 10.3, the benchmark-relative portfolios that have identical characteristics to the absolute-return portfolios satisfy the criteria derived in Scott (2011), namely:

$$\lambda * = \frac{IR - \rho SR}{SR - \rho IR} \tag{10.6}$$

where $\lambda*$ is the optimal risk budget, determined as a function of the information ratio (IR), the Sharpe ratio (SR), and the correlation between alpha and beta. The risk budget is the ratio of the tracking error to the benchmark risk. A tracking error of 2% and 4% benchmark volatility would have a risk budget of 2%/4% or 0.5. Unfortunately, it is not possible to derive in advance what the impact of the constraint will be on the information ratio of the portfolio manager. This means that it is not likely practical to compute the optimal risk budget. Nevertheless, it is perhaps useful to indicate the general magnitude of tracking error necessary to produce the most efficient benchmark-relative portfolios. As we shall see in the next section, perhaps the more important decision on tracking error is driven by the desire to protect in a bear-market environment. We will turn to address this important issue in the next section.

10.5 How to Avoid Tracking Bears

As mentioned at the outset, one of the primary motivations for switching to an absolute-return strategy is to benefit from downside protection during a bear market. In principle, benchmark-constrained investments should be dragged into negative territory when the market falls. Even if the active manager has added alpha, (s)he may still have made losses in absolute terms. By contrast, an absolute-return manager with market-timing skill aims to anticipate bear markets and shift the portfolio into cash to avoid negative returns. The question then arises, what would a benchmark-relative manager do if they had the same skill and anticipated the same bear market? Depending on the tracking error, the optimal portfolio construction would be one as close to cash as the tracking error would allow. How do our constrained portfolios measure up to this ideal? To find out, we re-examined the outcomes in Fig. 10.3 under a bear-market scenario.

In all cases, we assumed that the absolute-return and the benchmark-relative managers had both correctly anticipated a bear market and had shifted to a portfolio structure consistent with their investment objectives. The former, since they are focused on capital preservation, would shift the portfolio into cash in an extreme case. Without the same room for manoeuvre, the latter would have to do different things, depending on the constraints they were working under.

In the simple case, where the (unconstrained) benchmark-relative manager is maximising alpha subject to a limit on tracking error, they would shift as close to cash as the tracking error would allow. This would be

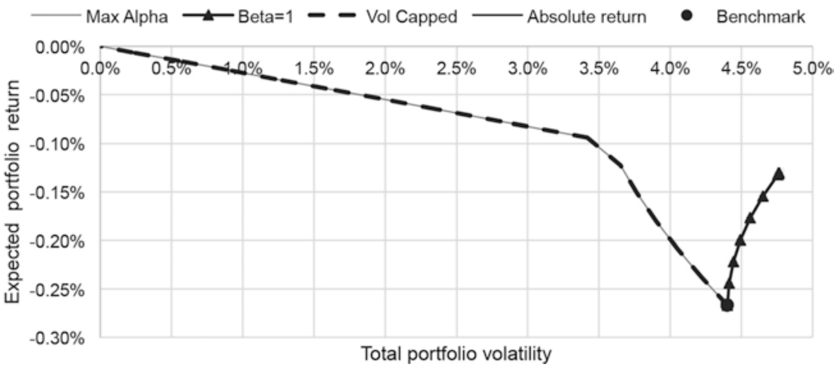

Fig. 10.4 The bear-market test

represented by the solid line in Fig. 10.4. The greater the tracking error, the further back up the solid line they would go and the lower the losses they would suffer. The constrained portfolio, where volatility must be no more than the benchmark volatility, would deliver the same portfolio as the unconstrained benchmark-relative strategy. Again, a higher tracking error would allow them greater leeway to move into cash. The reason they are identical is that both portfolios would be aimed at reducing risk in a bear market. However, the manager who has to hold the beta equal to one is labouring under an obvious disadvantage. Their performance must, perforce, be in line with the benchmark and therefore likely to be negative, depending of course on how much alpha they can derive from their asset mix and their security selection. The absolute-return portfolio is not visible on the graph, since, barring the ability to go short, the manager would be sitting completely in cash assuming all markets are producing negative returns.

The addition of one of the two constraints in a bull market environment clearly improves the efficiency and end-investor risk-adjusted return over an unconstrained benchmark-relative approach. In a bear market, however, the beta = 1 constraint is at a clear disadvantage to the total risk constraint. The total tracking error required to allow for an all cash position, however, is equal to the volatility of the benchmark, something that is higher than the conventional mandates might allow.

10.6 Implications for Investors and Conclusions

Investors who are interested in pursuing an absolute-return strategy either to improve portfolio efficiency or to avoid losses in bear markets are well served by making the switch, so long as the manager has the necessary market-timing skills. For those who would like the same benefits, but might wish—or be forced—to remain in a benchmark-relative framework, there are other options. This might be the case where the institution performs a strategic asset allocation and has budgeted risk and return to different investment teams for benchmark risks/returns and excess active risks/returns. The simplest prescription is to consider increasing tracking error of the mandate, allowing more defensive positions in a bear market. They could even consider non-traditional approaches like having asymmetric tracking error limits where the limit is large so long as the portfolio beta or total risk is being decreased. If the single most important element of absolute-return is loss-avoidance, then allowing enough tracking error to position in or close to a 100% cash holding would accomplish this.

Alternatively, the investor could add one of the restrictions mentioned in this chapter, while also allowing for enough tracking error to permit the benchmark-relative portfolio manager to move to the highest Sharpe ratio portfolio. The second constraint of limiting the total portfolio risk to no more than the benchmark risk has the added benefit of allowing the manager to move closer to cash ahead of an anticipated bear market.

Options for converting benchmark-relative mandates into absolute-return-like mandates:

1. Constrain total portfolio risk to being less than or equal to benchmark risk. Allow tracking error to be as large as the benchmark volatility. The large tracking error could result in aggressive positions, but only in the direction of defending the portfolio against losses. The downside is that the risk constraint tends to force a negative correlation between alpha and beta.
2. Constrain beta to be equal or less than one. Allow for a large tracking error. Constraining beta to one is fine in a bull market, but we saw that this was detrimental in a bear market. Changing the restriction to an inequality allows the manager to decrease overall risk in anticipation of a bear market.
3. Increase tracking error. In the absence of other constraints, the single easiest method for protecting downside in a bear market is to

allow the manager enough latitude to position the portfolio in cash without hitting any guideline constraints. Following this route alone does allow for the possibility of more severe losses in a bear market if the manger fails to correctly anticipate the decline.

None of these restrictions is commonplace. And they are likely to be met with resistance by some portfolio managers since they will force them to deliver a lower information ratio and perhaps lower alpha, which is often the basis for fees. Nevertheless, at a minimum, these arguments open a crack in the hitherto solid consensus that a benchmark-relative manager who maximises alpha is perfectly aligned with the interests of the end investor. There is perhaps room for improvement. One final note: This analysis is based on the assumption that there is market-timing skill. The decision to move from benchmark-relative to absolute-return will not in itself protect from losses. This is entirely dependent on a skilled portfolio manager correctly anticipating a bear market. These structures discussed above simply provide a framework to allow the skilled decisions to best be reflected in the construction of the portfolio.

Appendix: Simulation Details

Imagine a simple 60/40 stocks bonds portfolio where the stock component of the benchmark has a beta of one, meaning the benchmark has a beta of 0.6 (60% × 1) and the bond component is a simple 0–10-year universe of government bonds with a duration of 5, giving a benchmark duration of 2.0 (5.0 × 40%). We could simply describe this as a two-factor portfolio, and the decision for the portfolio manager is what the appropriate beta and duration are for the investment. There is a risk for each asset class (assumed to be 21% for the equity component and 3% for the bond component), and an expected return component. For equities, we have assumed an expected excess return over the risk-free rate of 7% and for bonds, 3%. Furthermore, we assume a correlation of 25% between stock and bond returns. It is important to note that the comparative results of this simulation are not sensitive to the actual expected returns, risks, or correlations (so long as they are not extremes, such as perfect positive or negative correlation, etc.). In an active process, the expected returns would change as the portfolio manager's views change, as well as possibly the expected correlation and volatilities. This information represents the minimum necessary to construct the best possible portfolio given a set of market views.

Under the absolute-return scenario, the possible portfolios are created using the highest expected return subject to a target or maximum portfolio volatility. The frontier of available portfolios then is the set of best possible portfolios assuming different levels of target risk. The simple benchmark-relative positions are the sensitivities that give the highest possible expected excess return over the benchmark (alpha), subject to a tracking error limit. It is important to point out here that these portfolios are based on the same market views. It is not feasible to have equities deliver 7% over cash for an absolute-return manager, and some other amount for a benchmark-relative manager. The market only has one outcome, although it can be measured against differing reference points. The constrained benchmark-relative simulations are based on the same framework and set of views as the unconstrained simulation but with the addition of beta = 1 in the first case and portfolio volatility ≤ benchmark volatility in the second case.

The simulation was repeated for five risk factors to ensure that the results were not unique to a two-asset portfolio, which produced similar results and identical conclusions.

References

Haugen, R. A., & Baker, N. L. (1991). The efficient market inefficiency of capitalization-weighted stock portfolios. *Journal of Portfolio Management, 17*(3), 35–40.

Haugen, R. A., & Baker, N. L. (2010). Case closed (Chapter 23). In J. Guerard Jr (Ed.), *The handbook of portfolio construction: Contemporary applications of Markowitz techniques* (pp. 601–619). New York: Springer.

Jorion, P. (2003). Portfolio optimisation with tracking-error constraints. *Financial Analysts Journal, 59*(5), 70–82.

Roll, R. (1992). A mean/variance analysis of tracking error. *Journal of Portfolio Management, 18*(4), 13–22.

Scott, R. (2011). Simple and optimal alpha strategy selection and risk budgeting. *Journal of Asset Management, 12*(3), 214–223.