NORGES BANK
INVESTMENT MANAGEMENT

01 | 2020

PORTFOLIO DELEGATION AND THE EFFECTS OF BENCHMARKS

DISCUSSION NOTE

We review the literature on portfolio delegation and examine the impact of benchmarked compensation on manager incentives, portfolio choice, and asset prices.

Date 27/08/2020
ISSN 1893-966X

dn@nbim.no
www.nbim.no

SUMMARY

- Delegated portfolios are common in the investment industry and benchmarks are widely used in management mandates and contracts. Benchmarks may be useful for guiding a manager's portfolio choice, constraining their ability to take risks, or incentivising managers to acquire information. We review the literature that focuses on how benchmarked compensation influences the incentives of managers.
- Naturally, benchmarked compensation incentivises managers to optimise their portfolio relative to their benchmark. When holding the benchmark, the manager's relative returns and compensation do not vary, and so the benchmark portfolio is effectively a risk-free asset for the manager. This can lead to portfolios that are overly risky when the investor cares about total portfolio returns.
- We show that it may be possible to improve portfolio outcomes for the investor by constraining benchmarked managers, or altering the composition of their benchmark. This is less feasible, however, if the manager can acquire additional information and construct portfolios with better risk-return properties compared to the investor.
- The incentive issues of benchmarked compensation and the widespread use of benchmarks can have implications for asset prices. A recent academic literature provides evidence that widespread benchmarking can lead to higher equilibrium prices, and lower expected returns, for assets that are widely included in benchmarks of delegated portfolios.

1. Introduction

Delegated portfolio management is pervasive across the investment industry. Individuals delegate management of their wealth to advisors, corporate pension plan sponsors delegate management of their assets and liabilities. Within investment organisations, chief investment officers delegate asset class, region or sector coverage to different portfolio managers. In all these delegation settings, investors provide managers with contracts and mandates that guide and constrain the management of their portfolios. While there are many ways in which these contracts vary, a common feature is the inclusion of a benchmark against which a manager's performance is evaluated.

In this note, we focus on how the incentives of managers are affected by benchmarks, highlighting some key insights from the portfolio delegation literature. We use a simple framework to model the portfolio choice of a benchmarked manager, and evaluate the portfolios they construct from the point of view of the investor. For a benchmarked manager, their benchmark acts as a risk-free asset. When holding the same portfolio weights as their benchmark, the manager constructs a portfolio that has no relative return variation, and is therefore riskless in terms of their compensation. This leads to a misalignment between the incentives of the investor and manager. Since the manager takes no risk when holding their benchmark portfolio, they are not incentivised to account for risk within the benchmark when choosing their portfolio weights. This can lead to manager constructing sub-optimal portfolios that are too risky from the point of view of the investor.

We outline some simple approaches that have been proposed to address these sub-optimal portfolio outcomes. We show that constraining the manager to a total portfolio volatility target can induce the manager to take relative positions that improve the risk-return trade-off of the portfolio. In addition, changing the composition of the manager's benchmark can potentially offset the inefficiencies that arise from relative return incentives, or from differences in risk aversion and horizons between the manager and investor. These solutions do not directly address the misalignment in incentives that arises from benchmarking compensation, however. They are also less feasible in an environment where the manager has additional information and can construct portfolios with better risk-return properties compared to the investor.

There are a range of possible benefits to benchmarking managers that may explain the prevalence of benchmarks in the investment industry. In our model, using total returns instead of relative returns in manager compensation can help align incentives between the investor and the manager. We emphasize, however, that this does not preclude the benefits of using benchmarks in portfolio management mandates and contracts. We outline an active literature that explores this issue, where recent studies have demonstrated how benchmarks can be useful in a range of delegated portfolio management settings. For example, investors may view benchmarks as a tool for guiding the manager's portfolio choice, and perhaps also as a way to constrain the manager's ability to take risks. Benchmarks can also be useful for incentivising managers to acquire information that improves portfolio performance, in particular when benchmarks are used alongside other sources of compensation, or when managers are constrained by certain mandate requirements or restrictions.

We outline a recent literature that examines the implications of widespread benchmarking for market pricing and equilibrium outcomes. The incentive effects that arise from benchmarking, combined with a prevalence of benchmarking in the investment industry, can have significant implications for asset prices. We describe a range of studies that show how widespread benchmarking can lead to higher equilibrium prices and lower expected returns for assets that are widely included in benchmarks of delegated portfolios. Furthermore, this can lead to persistent mispricing and a deterioration in the informational efficiency of asset prices.

The note proceeds as follows. In the next section, we outline a simple delegation framework, and in Section 3 we use the framework to illustrate how benchmarking influences incentives and portfolio choice. In Section 4, we consider how additional constraints and benchmark composition can be used to improve benchmarked portfolio outcomes. In Section 5, we consider the use of total return compensation, and discuss the role of benchmarks in manager compensation. In Section 6 we consider the implications of benchmarking for asset prices. Section 7 concludes.

2. Portfolio Delegation Framework

In this section, we outline the framework we use to examine how benchmarked compensation influences the incentives and portfolio choice of a manager. We first outline a basic model where the investor and the manager observe the same set of investment opportunities. We then extend the model where the manager possesses additional information and can construct portfolios with better risk-return properties compared to the investor.

Investor vs. Manager Portfolio Choice

To assess the effects of delegation and benchmarking on portfolio choice, we compare portfolio outcomes with and without delegation. We compare a case where an investor delegates full responsibility for portfolio construction to a manager, to a case where the investor forms a portfolio themselves. To do this, we need to specify the preferences of the investor and the manager, which we assume are described by the following utility functions:

$$u_I = -exp(-\gamma_I W_I) \tag{1}$$
$$u_M = -exp(-\gamma_M W_M) \tag{2}$$

where W_I is the value of the investor's portfolio without delegation. γ_I and γ_M are the coefficients of absolute risk aversion of the investor and manager, respectively. These exponential utility functions are commonly used in the portfolio delegation literature, though the analysis we present in this note are also in line with studies that use alternative utility functions.[1]

W_M is the manager's compensation when portfolio management is delegated, which is determined as follows:

$$W_M = ar_P + b(r_P - r_B) \tag{3}$$

where r_P is the return on the portfolio, and r_B is the return on a benchmark portfolio. This compensation function contains two variable components, a component that depends on the total returns of the portfolio, and a component that depends on the performance of the portfolio relative to a benchmark. The a and b parameters can be used to vary the manager's exposure to the total and relative performance incentives.[2]

With no delegation, the investor chooses their optimal portfolio weights, x_I, and the returns on the portfolio are given by $x_I'r$, where r is a vector of asset returns. We assume that asset returns are normally distributed, where the vector of expected returns and return covariance matrix are denoted by μ and Σ, respectively. For now, we assume that neither the investor nor the manager has access to a risk-free rate, which does not affect the analysis we present. The investor's optimal portfolio weights, x_I are given by:

$$x_I = \theta + \frac{1}{\gamma_I}\Delta\mu \tag{4}$$

where $\theta = (e'\Sigma^{-1}e)^{-1}\Sigma^{-1}e$ and $\Delta = \Sigma^{-1} - (e'\Sigma^{-1}e)^{-1}\Sigma^{-1}ee'\Sigma^{-1}$. μ is the vector of expected returns and e is a vector of ones. θ is the minimum variance portfolio, and combined with $\frac{1}{\gamma_I}\Delta\mu$ optimally balances expected returns and risk. These weights represent the investor's portfolio choice in the absence of delegation, and can be considered as an optimal portfolio against which to compare the manager's portfolio choice.[3]

When portfolio management is delegated, the manager chooses a vector of portfolio weights, x_M, and is provided with a vector of benchmark weights, x_B. The returns on the portfolio and benchmark are given by $x_M'r$ and $x_B'r$, respectively. The manager optimally chooses their portfolio subject to the compensation function in equation (3). We initially confine our discussion to the case where the manager only has a relative performance incentive, where $a = 0$. We later consider compensation with a total portfolio return incentive.

[1] These exponential utility functions, together with the assumption that returns are normally distributed, produce the same optimal portfolio weights as mean-variance preferences. It is well-known that these functions exhibit constant absolute risk aversion and imply that an investor holds the same amount in risky assets as wealth increases, and that the proportion of wealth in risky assets falls as wealth increases. This is inconsequential for the analysis we present.

[2] There are many possible extensions to this compensation function, such as fixed pay or a bonus component that cannot be negative, and implicit rewards through manager reputation and career concerns. The mechanisms and incentives that we highlight in this note are also present when using more complex compensation functions. We follow the majority of the portfolio delegation literature, however, in considering linear and symmetric compensation, but note that there are potentially important incentives effects that arise from the limited liability of managers, which is the subject of a large literature.

[3] We focus on the comparison between the portfolio the manager chooses and the portfolio the investor would choose in the absence of delegation. We exogenously define contract parameters and compare the portfolio outcomes of these parameters. We do not consider the full contracting problem where the investor optimally chooses the contract parameters.

The optimal portfolio weights for the benchmarked manager are:

$$x_M = x_B + \frac{1}{b\gamma_M}\Delta\mu \tag{5}$$

The optimal manager weights are conceptually similar to the optimal investor weights, though a key difference is that the minimum variance portfolio is replaced with the manager's benchmark weights. We can see from this expression that, if the manager was extremely risk averse, they would hold the same weights as the benchmark ($\lim_{\gamma_M \to \infty} x_M = x_B$). For benchmarked managers, holding the same weights as the benchmark ensures that they take no relative risk, and the benchmark portfolio acts as a risk-free asset for the manager, an idea we return to in the analysis that follows.[4]

The manager also chooses different portfolio weights to the investor to the extent that they choose a different amount of risk based on the coefficient γ_M, compared to the investor's risk aversion coefficient, γ_I. This risk scaling is further impacted by the manager's exposure to the relative performance of their portfolio, captured by the b parameter.

Portfolio Choice with an Informed Manager

Within the framework so far, we have assumed that the investor and manager observe the same expected returns and covariance matrix. We also consider an extension where the manager possesses additional information. In many studies, an investor chooses to delegate portfolio management in order to access this additional information. It is common to assume that the manager is able to observe a signal that is correlated with the distribution of asset returns, that enables them to construct portfolios with better risk-return properties compared to what the investor can achieve on their own.

Following studies such as Bhattacharya and Pfleiderer (1985), Stoughton (1993) and Admati and Pfleiderer (1997), we assume the manager observes a vector of signals, s, that are correlated with the vector of returns, r:

$$s = r + \epsilon \tag{6}$$

where ϵ is a noise term, that is uncorrelated with r. The value of the signals depends on the relative variability of r and ϵ. If the variability of returns is large relative to the noise terms, the signal is more precise. For simplicity, we assume that the signals are uncorrelated with one another.

Importantly, the manager observes the set of signals *prior* to the formation of their portfolio. This implies the manager can use the signals to inform their estimates of expected returns and the return covariance matrix. The manager observes a set of *conditional* expected returns and covariance matrix, which

[4] The importance of the benchmark for the manager's portfolio choice is not sensitive to our choice of utility function. While the optimal benchmarked portfolio can vary based on different utility functions, the benchmark invariably serves as a risk-free asset. This also the case with the inclusion of a risk-free rate. While a risk-free asset delivers constant total returns, these returns are a source of volatility relative to a benchmark.

we denote as μ_c and Σ_c, respectively. For asset i, these conditional moments are given by:

$$\mu_{ic} = \frac{\sigma_{ir}^2}{(\sigma_{ir}^2 + \sigma_{i\epsilon}^2)} s_i + \frac{\sigma_{i\epsilon}^2}{(\sigma_{ir}^2 + \sigma_{i\epsilon}^2)} \mu_i \tag{7}$$

$$\sigma_{ic}^2 = \frac{\sigma_{ir}^2 \sigma_{i\epsilon}^2}{\sigma_{ir}^2 + \sigma_{i\epsilon}^2} \tag{8}$$

where σ_{ir}^2 is the return variance of asset i and $\sigma_{i\epsilon}^2$ is the variance of the noise term which determines the amount of noise in the signal. A more precise signal implies that $\sigma_{i\epsilon}^2$ is small relative to σ_{ir}^2. The conditional expected return is a weighted combination of the signal and *unconditional* expected returns, and a more precise signal leads to a larger weight placed on the signal. We define ρ as the precision of the signals, which is equal to the ratio of the variance terms (for asset i, $\rho = \sigma_{ir}^2 / \sigma_{i\epsilon}^2$).

The manager constructs portfolios based on conditional expected returns, μ_c, and the conditional covariance matrix, Σ_c. Assuming the same utility function as earlier, and benchmarked compensation $W_M = b(r_P - r_B)$, the manager's optimal conditional portfolio choice, $x_{M,c}$, is given by:

$$x_{M,c} = x_B + \frac{1}{b\gamma_M} \Delta_c \mu_c \tag{9}$$

where Δ_c and μ_c are defined as earlier, based on conditional moments. The signal and conditional moments are not observed by the investor, and the investor must employ the manager to make use of this information. Here, even with $b = 1$ and $x_B = \theta$, the manager chooses a different portfolio to the investor. It is common to assume that the manager also needs to bear some cost, for example in terms of their time and effort, in order to acquire or improve the quality of this signal. We initially assume that the manager can costlessly observe the set of signals and cannot influence their precision, though we later relax this assumption.

Based on the delegation framework presented in this section, we proceed in the next section to illustrate the incentives and portfolio choice of the manager, and to evaluate the portfolio outcomes from the investor's point of view. In the sections that follow, we use both the simple framework with symmetric information and the extended model with additional manager information to explore the portfolio delegation problem.

3. Benchmarked Manager Incentives and Portfolio Choice

In this section, we use the basic investor-manager framework where the manager has no additional information compared to the investor. We also use a simple numerical example to further illustrate the intuition underlying the manager's incentives. The risk-free function of the benchmark described in Section 2 leads to a misalignment in incentives between the investor and the manager. An early study describing this effect is Roll (1992), which considers the portfolio choice of a manager that focuses on maximising returns relative to a benchmark, while minimising tracking error volatility. The paper shows that when the manager chooses portfolio weights that optimise *relative* returns and risk, this does not improve the *total* portfolio. The environment in Roll (1992) is equivalent to our framework when $b = 1$, which implies that the optimal portfolio choice of the manager is:

$$x_M = x_B + \frac{1}{\gamma_M}\Delta\mu \tag{10}$$

To understand how this leads to sub-optimal portfolios for the investor, we can compare the manager's optimal relative positions, $x_M - x_B$, with the optimal relative positions the investor would take, $x_I - x_B$. For now, we assume that the risk aversion coefficients, γ_I and γ_M are the same and equal to γ. Subtracting x_B from equations (4) and (10) gives the optimal investor and manager *relative* positions:

$$x_I - x_B = \theta + \frac{1}{\gamma}\Delta\mu - x_B \tag{11}$$

$$x_M - x_B = \frac{1}{\gamma}\Delta\mu \tag{12}$$

The optimal relative positions of the manager do not align with the optimal investor positions. The difference between the two sets of relative positions is $\theta - x_B$: the difference between the minimum variance portfolio and the benchmark.[5] In the remainder of this section, we use a simple numerical example based on these equations to further illustrate the portfolio choice implications of benchmarking a manager.

A Numerical Example

We assume an investment universe that consists of three assets, to provide a stylised representation of a broad asset allocation problem. The manager constructs a portfolio, and we define a benchmark, based on this the set of available assets. For simplicity, we assume all assets are uncorrelated.

For added intuition, we label the three assets as Equity (EQ), Fixed Income (FI) and Real Estate (RE). We do not attempt to describe realistic dynamics of these asset classes, however, and simply use labels to differentiate between more or less volatile assets in our asset space. For Real Estate in particular, we abstract from the well-known issues involved in the measurement of returns for this asset class.

We choose a small set of aggregated assets and disregard correlations in order to easily communicate the key incentives and portfolio implications

[5] As shown earlier, the optimal portfolio weights differ between the manager and investor in that the lowest risk for the manager is holding the benchmark weights, as opposed to the minimum variance portfolio for the investor. The minimum variance portfolio maximises diversification and tends to include all assets to achieve this. When benchmarked, the minimum (zero) variance portfolio for the manager is the benchmark itself. When the benchmark deviates from the minimum variance portfolio, the difference in optimal relative positions for the manager and investor reflects the deviation between the minimum variance portfolio and the benchmark. This implies that the manager will construct an optimal portfolio if provided the minimum variance portfolio as a benchmark.

within a delegated management setting. The manager incentives and portfolio choice that we describe in this environment are robust to alternative parameterisations, however, and also generalise to more complicated asset spaces. The asset expected returns and volatilities are provided in Table 1.

Table 1: Parameterisation

Asset	Expected Return	Volatility
Equity (EQ)	5%	15%
Fixed Income (FI)	1.5%	8%
Real Estate (RE)	4%	14%

Within our example, we assume that the manager is given a mean-variance inefficient benchmark. This means that it possible to construct portfolios with a higher expected return for the same amount of risk as the benchmark (or lower risk for the same expected return). We specify an inefficient benchmark by defining a benchmark that contains only the equity and fixed income assets, therefore capturing only a subset of the full asset universe. We assume the investor provides the manager with a benchmark with 60% equities and 40% fixed income.[6] We assume that the manager's investment universe includes all three assets which can be used to form a portfolio.

Figure 1 (a) shows the optimal investor and manager portfolio weights in our numerical example, where we set $\gamma = 5$. Figure 1 (b) shows the same investor and manager optimal positions expressed relative to the benchmark. The investor would like to include a position in the RE asset because it improves the diversification of the total portfolio. This RE position implies a lower total weight in both the equity and fixed income relative to the benchmark. In terms of positions relative to the benchmark, this implies that a long RE position is financed using short positions in both the equity and fixed income assets. The manager, however, implements a different set of positions to the investor's optimal set. The manager chooses to increase the weight in equities relative to the benchmark, with only a small position in RE, and to hold a much lower weight in fixed income compared to the benchmark.

The manager chooses a total portfolio that is sub-optimal from the point of view of the investor. Figure 2 (a) shows the Sharpe ratios of the optimal investor and manager portfolios. The manager's positions lead to a deterioration in the risk-return profile of the portfolio, relative to the optimal investor portfolio.[7] The figure also shows the information ratios of the portfolios defined as the expected excess return of the portfolio relative to the benchmark, scaled by its tracking error. The information ratio is considerably higher for the manager's positions compared to what the

[6] We choose these weights for illustrative purposes. As long as the benchmark is mean-variance inefficient, the choice of weights in the simplified benchmark has little impact on our illustration. We further explore the role of benchmark composition in Section 4.

[7] In our example, the risk-return profile of the manager's portfolio is still an improvement relative to the benchmark, but this may not always be the case.

Figure 1: Investor vs. manager portfolio choice

(a) Benchmark, investor and manager total positions *(b) Investor and manager positions relative to benchmark*

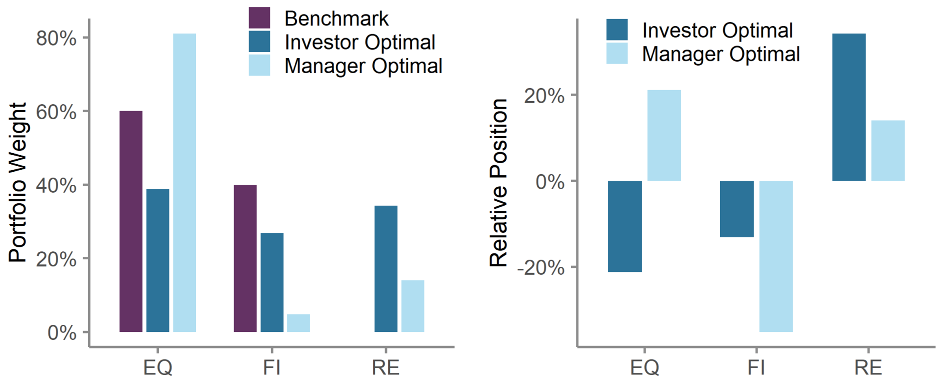

Figure 2: Investor vs. manager risk-return

(a) Sharpe and information ratios of investor and manager positions *(b) Efficient frontier and the risk-return properties of alternative portfolios*

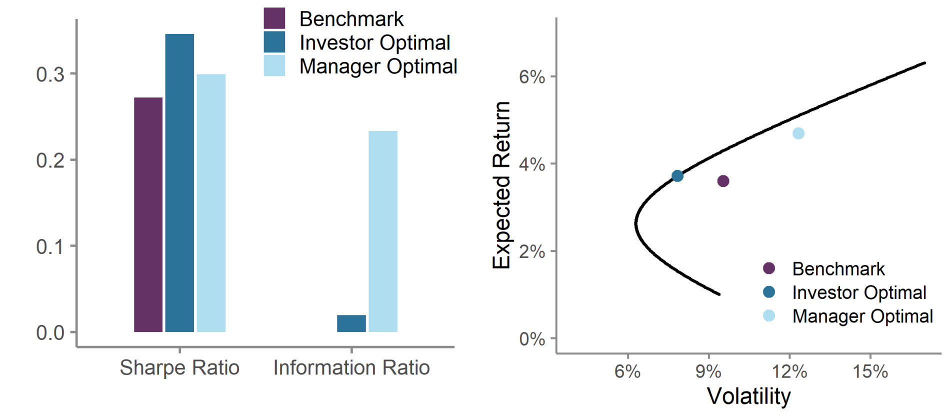

investor would have chosen. When the manager is benchmarked, they maximise relative returns while minimising the variability in relative returns, in order to optimise their compensation. In other words, the manager maximises their information ratio, while the investor would like to maximise the Sharpe ratio of the portfolio. As described above, the benchmark is a risk-free asset from the point of view of the manager. As a result, the manager does not take into account risks that are present within the benchmark, and chooses to take an equity position from this risk-free starting point. This means that the manager adds risk on top of those already within the benchmark, which are not accounted for by the manager. This leads to a risk-shifting issue from the manager to the investor, and this leads to overly risky portfolios from the point of view of the investor. Figure 2 (b) shows the expected return and volatility of the manager's portfolio, the benchmark, and the investor's optimal portfolio, alongside the efficient frontier in this asset space. The inefficient benchmark lies within the frontier, and efficiency improvements would result in a movement upwards and/or to the left. The manager's optimal portfolio, however, leads to an additional risk compared to the benchmark and lies below the efficient frontier.

4. Improving Benchmarked Portfolio Choice

Motivated by the inefficient portfolios that result from benchmarking in Section 3, many studies explore how these outcomes can be improved. In this section, we describe an example where imposing a risk constraint on a manager can improve the portfolio from the investor's perspective. We then explore the possibility of changing the composition of the benchmark to improve the total portfolio. We also consider the effectiveness of these measures when extending our framework. In addition to the possible risk-return improvements from investing in a wider set of assets than the benchmark, we consider a case where the manager has additional information that can further improve the portfolio's risk-return properties.

Constraining Managers

Many studies consider whether additional risk constraints placed on benchmarked managers can improve their portfolio choice. Roll (1992) suggests that one way to mitigate the excessive total portfolio risk that managers take is to constrain the beta of the portfolio. Alternatively, Alexander and Baptista (2010) show that providing an alpha target to the manager can lead to efficiency improvements for the total portfolio, where alpha is the intercept from a regression of portfolio returns on the returns of the benchmark. Additional proposed alternatives include constraints on tail risk measures such as Value-at-Risk (e.g. Alexander and Baptista (2008) and Palomba and Riccetti (2012)), expected shortfall (e.g. Stucchi (2015)) and maximum drawdown (e.g. Alexander and Baptista (2006)).

Using the numerical example from Section 3, we illustrate how a constraint on the total risk of the portfolio can improve portfolio choice from the point of view of the investor. This constraint was proposed by Jorion (2003), who recommends requiring the manager to construct a portfolio which has total volatility equal to the volatility of the benchmark. When constrained in this way, the manager chooses relative positions that improve the total portfolio risk-return characteristics.

Within our numerical example, we constrain the manager's portfolio choice problem, such that the variance of the manager's portfolio $\sigma_M^2 = x'_M \Sigma x_M$ equals the variance of the benchmark $\sigma_B^2 = x'_B \Sigma x_B$. Figure 3 shows the manager's optimal total and relative positions that satisfy this constraint, alongside the optimal total and relative positions shown earlier in Figure 1. These positions are more closely aligned with the optimal positions from the point of view of the investor, where the additional total risk constraint helps to align the manager's weights with their optimal portfolio choice. Figure 3 (b) shows the effect of the constraint on the manager's relative positions. The constraint leads the manager to increase the RE position, and they no longer take a further equity position relative to the benchmark. We can see the intuition behind this by noting the following relationship between the variance

Figure 3: Investor vs. manager portfolio choice: total risk constraint

(a) Benchmark, investor and manager total positions *(b) Investor and manager positions relative to benchmark*

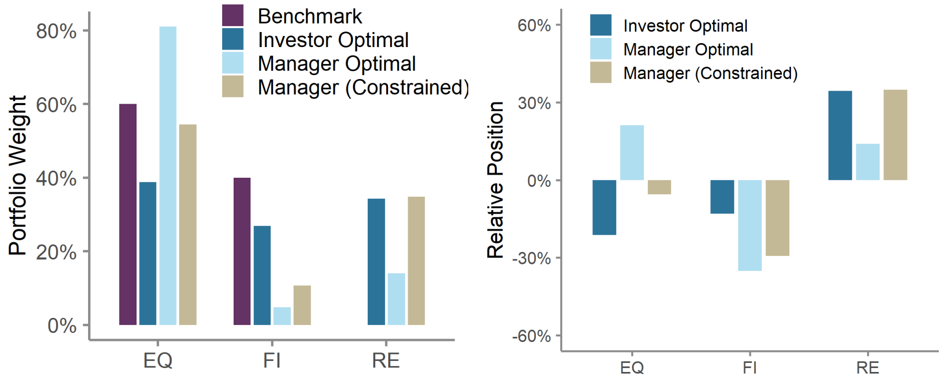

Figure 4: Investor vs. manager risk-return: total risk constraint

(a) Sharpe and information ratios of investor and manager positions: total risk constraint *(b) Efficient frontier and the risk-return properties of alternative portfolios: total risk constraint*

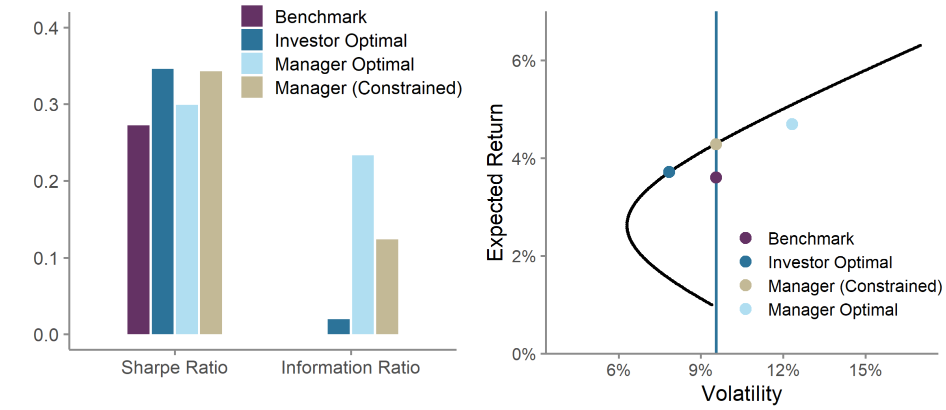

of the manager's portfolio, σ_M^2, and the variance of the benchmark, σ_M^2:

$$\sigma_M^2 = \sigma_B^2 + (x_M - x_B)'\Sigma(x_M - x_B) + 2x_B'\Sigma(x_M - x_B) \tag{13}$$

The variance of the manager's portfolio returns is equal to the benchmark portfolio variance plus two additional terms. The first, $(x_M - x_B)'\Sigma(x_M - x_B)$, is variance that results from the manager's positions relative to the benchmark, which is always positive. The second term, $2x_B'\Sigma(x_M - x_B)$, measures the covariance of the manager's relative positions with the benchmark.

In order to offset the variance that arises from positions relative to the benchmark, the manager must choose relative positions that co-vary negatively with the benchmark. This ensures that the variance of the portfolio is equal to the variance of the benchmark. In our simple example, this requires the manager to take relative positions that are opposite to the benchmark, that is, underweighting the equity and fixed income assets.

It is worth noting that introducing additional risk constraints on managers does not address the misalignment in incentives between the investor and the manager. The manager still optimises positions relative to the benchmark, but the additional risk constraint forces the manager to take relative positions they would not otherwise take. Figure 4 (a) compares the Sharpe and information ratios of the risk constrained portfolio with the unconstrained positions from earlier. Naturally, the constrained relative positions have a lower information ratio compared to the case where the manager is unconstrained.

Figure 4 (b) shows the risk and return of the manager's portfolio choice when subject to the total risk constraint, where the risk constrained portfolio lies directly above the benchmark. The constraint prevents the manager from increasing portfolio risk, as we saw in Figure 2 (b), and instead leads to an efficiency improvement through a higher expected return for the same volatility as the benchmark.

Risk Constraints and Additional Manager Information

Additional concerns may arise with the use of a total risk constraint when the manager has additional information relative to the investor. To explore this further, we use the extended framework where the manager can construct portfolios with better risk-return properties due to this additional information. Based on the numerical example presented so far, the investor and manager are able to form portfolios based on *unconditional* moments. We now assume that the manager can observe a signal with precision, ρ, equal to 0.05, that translates into expected outperformance or alpha of approximately 2% per year.[8]

Table 2 shows Sharpe ratios for portfolios that are constructed using unconditional and conditional moments.[9] Earlier, using unconditional moments, we showed that a benchmarked manager chooses a portfolio that is overly risky from the investor's point of view. Compared to the optimal Sharpe ratio of 0.35 that the investor could achieve without delegation, the manager constructs a portfolio with a Sharpe ratio of 0.30 when provided a simple benchmark. We also showed that the manager's portfolio choice could be improved by imposing an additional constraint on the total risk of the manager's portfolio. This results in a Sharpe ratio of 0.35.

The conditional Sharpe ratios in Table 2 show the extent to which portfolios can be improved using the additional information provided by the signals. With no delegation and assuming the investor possesses the additional information, they can achieve a conditional Sharpe ratio of 0.45. This optimal investor case fully utilises the value of the additional information. If instead we assume that only the manager observes the set of signals and is provided with a simple benchmark as before, this leads to a conditional Sharpe ratio of 0.37.

[8] This is the alpha of a regression of the investor's returns when using conditional moments in their portfolio choice, on returns using unconditional moments.

[9] We simulate 10,000 sets of signals and returns to evaluate the performance for the portfolio alternatives.

Table 2: Unconditional vs. conditional Sharpe ratios

	Investor Optimal	Manager Optimal	Manager (Constrained)
Unconditional	0.35	0.30	0.35
Conditional	0.45	0.37	0.37

When applying a total risk constraint to a manager with additional information, we don't see the same improvement in Sharpe ratio that we saw in the unconditional analysis. The conditional Sharpe ratio under the total risk constraint is approximately the same as in the case of a benchmarked manager with no constraint, and below the optimal investor case that fully utilises the value of the additional information. With the total risk constraint, the manager is often forced to ignore the value of the information in order to meet the constraint. This is partly the result of a small number of assets in our numerical example restricting the scope for alternative portfolio weights that satisfy the total risk constraint, though we achieve similar results when using simulations with a larger number of assets.

The use of risk constraints to improve the manager's portfolio choice potentially comes at a cost when the manager has additional information. When the investor delegates management to access the information of a manager, they can improve the unconditional portfolio risk-return using a risk constraint, but this is not necessarily effective in improving the conditional risk-return.[10]

Changing the Composition of the Benchmark

In addition to manager constraints, it may be possible to change the composition of the benchmark to improve portfolio outcomes for the investor. An important result in our framework is that the manager's choice of relative positions is independent of the benchmark composition. We can see this by again examining the manager's optimal relative positions by rearranging equation (10) (again assuming that $b = 1$):

$$x_M - x_B = \frac{1}{\gamma_M}\Delta\mu \tag{14}$$

The manager's optimal relative positions do not depend on the benchmark weights themselves. This implies that the manager would choose the same positions for any alternative benchmark, assuming that they have no capacity constraints.[11] On the other hand, the composition of the benchmark still determines the total portfolio weights that result from the manager's positions. It is therefore possible to offset the sub-optimal portfolio outcomes in our framework through changing the benchmark composition.

We relax the assumption that the investor and manager have the same risk

[10] Also, in an environment with time-varying investment opportunities, constraining the manager based on short-term risk metrics may be problematic. For example, short-term total portfolio volatility is an inappropriate measure of risk for a long-term investor.

[11] For example, if the manager is unable to take outright short positions, the maximum short relative position is equal to the total weight of the asset within the benchmark.

Figure 5: Optimal benchmark and total portfolio weights

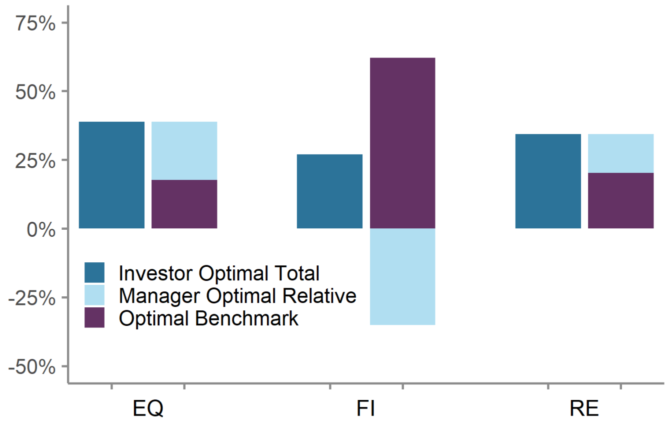

aversion coefficients. We can solve for the set of optimal benchmark weights, x_B^*, that ensure the manager's weights equal the investor optimal weights:

$$x_B^* = \theta + \left(\frac{1}{\gamma_I} - \frac{1}{\gamma_M}\right)\Delta\mu \tag{15}$$

The composition of the optimal benchmark can understood by comparing this expression with the optimal investor portfolio weights in equation (4):

$$x_I = \theta + \frac{1}{\gamma_I}\Delta\mu$$

The optimal benchmark is simply the investor's optimal portfolio minus $\frac{1}{\gamma_M}\Delta\mu$, which is the set of optimal relative positions given in equation (14). By altering the benchmark to reflect the anticipated portfolio choice of the manager, the manager's relative positions lead to a total portfolio that aligns with the optimal portfolio for the investor. This effect is demonstrated, using the numerical example, in Figure 5. As described above, the manager's optimal relative positions are the same regardless of the benchmark. The manager's relative positions combined with the optimal benchmark weights align the total managed portfolio and the optimal weights for the investor. This approach to benchmark design can also take into account possible differences in risk aversion between the investor and the manager. Similar to the risk constraint solution presented earlier, the benchmark composition approach also does not address the misalignment in incentives between the investor and the manager.

This idea of optimal benchmark design underlies some of the main results presented in Van Binsbergen, Brandt, and Koijen (2008). This study emphasises the role of benchmark composition in tackling the misalignments and inefficiencies that arise when delegating portfolio management.[12] In addition to adjusting the benchmark composition for the manager's relative

[12] Van Binsbergen et al. (2008) consider an environment where a CIO delegates portfolio management to multiple portfolio managers. They show that benchmark design can be used to offset diversification losses that arise from portfolio managers optimising within asset classes, or from differences in risk aversion between the investor and manager.

return incentives and differences in risk aversion, Van Binsbergen et al. (2008) further show that benchmarks can be used to align differences in investment horizon between the investor and manager.[13]

Benchmark Composition and Manager Information

Amending the composition of the benchmark in anticipation of the manager's positions becomes infeasible when the investor does not have full information about the manager. In order to correctly change the composition, the investor needs to know the manager's preferences and risk aversion, and to observe the same expected returns and return covariances. In our extension with additional manager information, the benchmark composition solution is not possible. This issue underlies the results presented in Admati and Pfleiderer (1997), an influential study on delegated portfolio management that shows that there are few circumstances under which benchmarks can be used to achieve optimal portfolios. In their framework, there are only two ways to align the manager's portfolio choice: either the benchmark must be set equal to the (conditional) minimum variance portfolio, or the manager must be compensated based on total portfolio returns, rather than returns relative to a benchmark portfolio. This raises the question of whether a manager's compensation should include a benchmark component at all, which we discuss in the next section.

5. Should Managers be Benchmarked?

So far, we have discussed how a manager's incentives change when they are benchmarked, and examined how benchmark composition and risk constraints may improve portfolio outcomes. We have noted that these proposed solutions do not address the misalignment in incentives that arises from benchmarking. In this section, we show that one way in which to better align incentives between the manager and the investor is to base the manager's compensation on total portfolio returns. In other words, we could remove or replace the benchmark component in their compensation. We also describe an active literature that documents a range of possible benefits to benchmarking and attempts to explain their widespread use. Despite the incentive issues caused by benchmarks, benchmarks may be useful for guiding a manager's portfolio choice, constraining their ability to take risks, or incentivising managers to acquire information.

In order to understand compensation based on total returns, we revisit the compensation function in equation (3):

$$W_M = ar_P + b(r_P - r_B)$$

The a and b coefficients determine the balance between the manager's

[13] In our framework, there is no difference in horizon between the investor and manager. Van Binsbergen et al. (2008) extend a similar framework to allow expected returns to vary over time, where the optimal portfolio for a long-term investor includes hedging demands in addition to the standard mean-variance weights. They suggest that benchmark composition can be used to account for the manager's shorter horizon through changing their starting point.

exposures to the total and relative performance of the portfolio. Under this compensation scheme, the manager's optimal portfolio choice is given by:

$$x_M = \frac{1}{a+b}(bx_B + a\theta + \frac{1}{\gamma_M}\Delta\mu) \tag{16}$$

The introduction of a total return component to the manager's compensation leads to the addition of the minimum variance portfolio, θ, in their optimal portfolio choice. As shown earlier, the investor's optimal weights are given by:

$$x_I = \theta + \frac{1}{\gamma_I}\Delta\mu$$

Here, we can see that one way to better align the portfolio choice of the investor and the manager is to set $b = 0$. In this case, the manager's optimal weights become:

$$x_M = \theta + \frac{1}{a\gamma_M}\Delta\mu \tag{17}$$

In this case, if the manager and investor have the same risk aversion, their portfolio choice is equal with $a = 1$. Or more generally, if the manager's risk aversion is known the investor can align portfolio choice by setting $a = \frac{\gamma_I}{\gamma_M}$.[14] Naturally, the more aligned the preferences of the manager and investor, the more aligned their incentives and portfolio choice.

Benchmark "Irrelevance"

In light of the alignment in incentives that can be achieved based on total returns, it is worth noting a stark result in the portfolio delegation literature that argues that benchmarks are *irrelevant*. In an important contribution by Admati and Pfleiderer (1997), they show that there is little scope for benchmarks to induce managers to acquire information relevant to improving the portfolio, or to help infer manager skill. In fact, using a model similar to our extended framework with additional information, they show that the manager's contract and the benchmark in fact play no role in aligning a manager's portfolio choice with the optimal portfolio of investor.

So far, we have assumed that the manager has access to a set of signals with a given level of precision. Following Admati and Pfleiderer (1997), we now consider a framework where the manager is able to improve the level of precision by exerting higher 'effort'. A stark result in the portfolio delegation literature is that the manager's contract is irrelevant when it comes to incentivising a manager to exert higher effort and improve portfolio performance. Admati and Pfleiderer (1997) show that it is not possible to affect the manager's effort through changing their exposure to relative returns in their compensation function.

We now refer to ρ as the level of the manager's effort, which as before directly

[14] This assumes that the investor can observe the manager's risk aversion. In addition, there may still be a misalignment in portfolio choice if the horizons of the investor and manager differ.

translates into the level of precision of the signals:

$$\rho = \frac{\sigma_{ir}^2}{\sigma_{i\epsilon}^2} \tag{18}$$

In contrast to earlier, ρ is a choice variable for the manager, and effort is costly for the manager to exert. We capture the costs of effort through the generic function $V(\gamma_M, \rho)$ so that the utility of the manager becomes:

$$u_M = -exp(-\gamma_M W_M + V(\gamma_M, \rho)) \tag{19}$$

Studies such as Stoughton (1993) and Li and Tiwari (2009) show that the 'irrelevance' result holds for a wide set of cost functions.[15] The manager's optimal behaviour balances the higher signal precision that results from higher effort with the costs of exerting that effort.

It seems intuitive that given a higher value of b in the manager's contract, the manager would be more sensitive to their performance relative to the benchmark, and would therefore be incentivised to exert effort and improve the precision of their signals. In contrast, the b contract parameter is actually irrelevant for the manager, who will exert the same level of effort regardless of its value.

Within our framework, for any value of b, the manager is able to form a portfolio that has the same compensation profile for any set of realised returns. We can see this by combining the manager's compensation function, $W_M = b(r_P - r_B)$, with their optimal portfolio weights in equation (9). Given that $r_P = x'_{M,c}r$ and $r_B = x'_B r$, the manager's compensation function can be written as follows:

$$W_M = b\left(\frac{1}{b\gamma_M}\Delta_c\mu_c\right)' r \tag{20}$$

Here, the b terms offset such that the manager's compensation does not depend on the exposure to relative performance. The first b term captures the effect of increasing the manager's exposure to relative returns when b is increased in their compensation function. However, the manager's optimal portfolio choice is also scaled by b. In other words, the manager optimally reduces their portfolio risk to offset any increase in the b parameter in their compensation. The manager is able to undo any incentive that is provided to them, and their effort choice balances the costs and benefits of increased precision independently of the value of b. Similarly, the composition of the benchmark does not affect the manager's compensation.[16]

Revisiting Benchmark Relevance

Following the benchmark irrelevance result, there has been an ongoing literature examining the case for benchmarking managers. Several studies

[15] The function needs to satisfy the conditions described in Stoughton (1993), that is the cost function is an increasing convex function of effort starting from zero.

[16] In Section 4, we showed that the manager's relative positions do not depend on the composition of the benchmark for a fixed value of b. For the irrelevance of the benchmark for the manager's compensation, the relative positions of the manager vary for different values of b.

re-consider the role of benchmarks when extending beyond the simple delegation setting we have used, in part to try and explain the widespread use of benchmarking in practice. There are a range of studies that have shown that benchmarks may indeed play a useful role in incentivising managers to exert effort to acquire information relevant to future asset returns.

Some studies have shown that it is optimal to include a benchmark in portfolio management contracts when the manager's compensation also includes option-like components or when compensation is non-linear. For example, Ou-Yang (2003) and Li and Tiwari (2009) show that benchmarks feature in the optimal contract alongside a fixed fee and fractions of assets under management. In the analysis in Li and Tiwari (2009), it is optimal to compensate the manager relative to a benchmark with option-like compensation that has a floor at zero. For an appropriately designed benchmark, this provides a strong incentive for the manager to exert effort in acquiring information about future asset payoffs.

Other studies have emphasized the importance of additional contracting features or manager activities that affect the relevance of benchmarking. Studies such as Gómez and Sharma (2006), Dybvig, Farnsworth, and Carpenter (2009) and Agarwal, Gómez, and Priestley (2012) highlight the importance of mandate requirements and trading restrictions given to the manager in restoring the case for benchmarking manager compensation. Intuitively, benchmarks can be relevant if an additional constraint on a manager prevents them from fully adjusting their portfolio to undo the incentive provided by the benchmark that we showed in equation (20). In addition, Kashyap, Kovrijnykh, Li, and Pavlova (2019b) argue that the risk-free nature of benchmarks can enable managers to generate higher returns for the investor. In their model, when the manager is able to enhance portfolio returns through activities such as securities lending or transaction cost reduction, a benchmark can incentivise the manager to enhance returns on a larger scale.

Recent studies such as Sockin and Xiaolan (2019) and Breugem and Buss (2019b) have also taken an equilibrium perspective in assessing manager contracts and the incentives they provide. Sockin and Xiaolan (2019) model a delegation setting where investors allocate between passive and active funds, and need to incentivise active managers to exert effort to acquire information about asset payoffs. They show that in equilibrium, benchmarking arises endogenously as part of the active managers' compensation as a means to align the manager's portfolio choice with the investor.

Overall, there appear to be benefits to benchmarking which may justify their widespread use. Interestingly, the incentive issues of benchmarks can still persist alongside these benefits, and this may have broader implications when many managers are benchmarked in the investment industry. This has led to an active research effort exploring the implications of widespread benchmarking for asset prices, and we briefly outline this research in the next section.

6. Equilibrium Effects of Benchmarks

While there is an ongoing debate on the optimality of benchmarking managers, there is little disagreement that benchmarks are pervasive within the investment industry.[17] As emphasized throughout this note, benchmarks can have a meaningful effect on manager incentives and portfolio choice. These effects, combined with widespread benchmarking, can have significant implications for equilibrium asset prices.

A range of studies develop equilibrium models with delegated benchmarked portfolios, focusing on the equity market. We can use our framework to outline the basic reason why benchmarking has implications for asset prices. For simplicity, we assume that the manager has no additional information relative to the investor. We also make some further assumptions following the simple model presented in Kashyap, Kovrijnykh, Li, and Pavlova (2019a), which is also similar to the models presented in Basak and Pavlova (2013) and Buffa, Vayanos, and Woolley (2019). First, we assume that both the investor and manager are able to invest in a risk-free asset. In addition, we assume that assets are claims on cash-flows next period, D, with prices, P. The distribution of cash flows is described by $D \sim N(\mu_D, \Sigma_D)$. With aligned risk aversion ($\gamma_I = \gamma_M = \gamma$), the investor and manager optimal portfolios become:

$$x_I = \frac{1}{\gamma}\Sigma_D^{-1}(\mu_D - P) \tag{21}$$

$$x_M = x_B + \frac{1}{b\gamma}\Sigma_D^{-1}(\mu_D - P) \tag{22}$$

We assume that the two types of agents, investors and managers, make up proportions λ_I and λ_M, respectively, of the number of agents within the market.[18] Kashyap et al. (2019a) show that in this environment, equilibrium prices are given by:[19]

$$P = \mu_D - \frac{\gamma}{\lambda_I + \lambda_M}\Sigma_D(1 - \lambda_M x_B) \tag{23}$$

The implications of benchmarking can be understood by considering the case with no delegated portfolios, where $\lambda_M = 0$. In this case, asset prices would simply represent expected returns adjusted for risk, a common result in standard asset pricing models. When introducing benchmarked portfolio managers, with $\lambda_M > 0$, assets that are widely included within managers' benchmarks have higher prices. This results in an effective decrease in the importance placed on the riskiness of these assets. This again results from the idea that the benchmark is the risk-free asset for managers, and that it is inherently risky to depart from their benchmark.[20]

[17] BIS (2003) provide survey evidence indicating the prevalence of benchmarking in the asset management industry. Ma, Tang, and Gómez (2019) also find that the large majority of US mutual funds link manager compensation to performance relative to a benchmark.

[18] Implicity, there are two types of investor, where one type choses to delegate their portfolios to managers. The fractions of managers and the two investor types sum to one.

[19] To solve for asset prices, we impose the market clearing condition that $\lambda_I x_I + \lambda_M x_M = 1$, where the supply of each asset is normalised to 1.

[20] For simplicity, the framework describes the intuition of benchmark effects in terms of stocks

This basic implication of benchmarking motivates a range of studies exploring the equilibrium effects of widespread benchmarking, and make further asset pricing predictions within an environment with a large role for delegated portfolios. Studies such as Basak and Pavlova (2013) and Cuoco and Kaniel (2011) show that upward pressure on the prices of stocks within the benchmarks lowers expected returns and Sharpe ratios in equilibrium.[21] Furthermore, Basak and Pavlova (2013) show that trading by benchmarked investors generates excess correlations between stocks in the benchmark.[22]

Recent studies also consider the impact of widespread benchmarking on mispricing and the informational efficiency of asset prices. Breugem and Buss (2019a) show that benchmarking can reduce informativeness of the market. Once again, this is because relative return investors hold benchmark stocks as a means of reducing relative risk. These stock demands are insensitive to information, and this reduces the value of private information. Other studies such as Buffa et al. (2019) and Sotes-Paladino and Zapatero (2018) show that incentives that arise within delegated mandates can prevent managers from trading against mispricing or investing in overvalued assets. Relatedly, Lines (2016) provides evidence that increases in return volatility lead managers to move closer to their benchmarks, which involves buying underweight stocks and selling overweight stocks, which leads to price distortions.

In addition to the literature documenting the impact of benchmarking on asset prices, there is evidence that benchmarking can influence corporate decision-making. Kashyap et al. (2019a) show that the upward price pressure on stocks within benchmarks acts as an effective subsidy to these firms in equilibrium. These price pressures lead to a lower cost of capital for benchmarked firms. As a result, benchmarking influences the evaluation of opportunities such as mergers and acquisitions and initial public offerings, and may affect the funding of smaller innovative firms to the extent they are not widely included in benchmarks.

7. Summary

In this note, we have highlighted incentives issues that arise from benchmarked compensation. Benchmarks incentivise managers to optimise their portfolio relative to their benchmark, and this can lead to portfolios that are overly risky when the investor cares about total portfolio returns. It may be possible to improve portfolio outcomes for the investor by constraining benchmarked managers, or altering the composition of their benchmark, but this is less feasible if portfolio management is delegated on the basis that the

that are either inside or outside of a benchmark. In practice, a typical stock index used as a benchmark will include a large proportion of the total universe of companies and total market capitalisation. Here, it is more intuitive to think in terms of the extent to which a stock is included in benchmarked capital, which will lead to varying magnitude of effects across stocks.

[21] This is also consistent with the literature documenting the importance of index membership for stock prices. Stock prices react significantly to the announcement of the inclusion or removal from the S&P 500 (Harris and Gurel (1986), Shleifer (1986) and Chen, Noronha, and Singal (2004)), and co-movement with index constituents increases following stock inclusions (see Barberis, Shleifer, and Wurgler (2005) and Boyer (2011)).

[22] In a related study, Leippold and Rohner (2011) provide a model and empirical evidence showing that stocks with high institutional ownership earn lower returns on average, and that non-benchmark stocks or stocks with low correlation to the benchmark are rewarded with a premium.

manager has additional information.

While, following the literature, we emphasize the incentive issues that can arise from benchmarking compensation, our analysis does not preclude a range of possible benefits to benchmarking. One way to better align incentives is to base the manager's compensation on total instead of relative returns, but benchmarks may be useful for guiding a manager's portfolio choice, constraining their ability to take risks, or incentivising managers to acquire information.

References

Admati, A. R. and P. Pfleiderer (1997). Does it all add up? benchmarks and the compensation of active portfolio managers. *Journal of Business 70*(3).

Agarwal, V., J.-P. Gómez, and R. Priestley (2012). Management compensation and market timing under portfolio constraints. *Journal of Economic Dynamics and Control 36*(10), 1600–1625.

Alexander, G. J. and A. M. Baptista (2006). Portfolio selection with a drawdown constraint. *Journal of Banking & Finance 30*(11), 3171–3189.

Alexander, G. J. and A. M. Baptista (2008). Active portfolio management with benchmarking: Adding a value-at-risk constraint. *Journal of Economic Dynamics and Control 32*(3), 779–820.

Alexander, G. J. and A. M. Baptista (2010). Active portfolio management with benchmarking: A frontier based on alpha. *Journal of Banking & Finance 34*(9), 2185–2197.

Barberis, N., A. Shleifer, and J. Wurgler (2005). Comovement. *Journal of Financial Economics 75*(2), 283–317.

Basak, S. and A. Pavlova (2013). Asset prices and institutional investors. *American Economic Review 103*(5), 1728–58.

Bhattacharya, S. and P. Pfleiderer (1985). Delegated Portfolio Management. *Journal of Economic Theory 36*.

BIS (2003). Incentive structures in institutional asset management and their implications for financial markets. *Bank for International Settlements, Committee on the Global Financial System*.

Boyer, B. H. (2011). Style-related comovement: Fundamentals or labels? *The Journal of Finance 66*(1), 307–332.

Breugem, M. and A. Buss (2019a). Institutional investors and information acquisition: Implications for asset prices and informational efficiency. *The Review of Financial Studies 32*(6), 2260–2301.

Breugem, M. and A. Buss (2019b). Optimal contracts in delegated portfolio management - an equilibrium analysis.

Buffa, A. M., D. Vayanos, and P. Woolley (2019). Asset management contracts and equilibrium prices. Technical report.

Chen, H., G. Noronha, and V. Singal (2004). The price response to S&P 500 index additions and deletions: Evidence of asymmetry and a new explanation. *The Journal of Finance 59*(4), 1901–1930.

Cuoco, D. and R. Kaniel (2011). Equilibrium prices in the presence of delegated portfolio management. *Journal of Financial Economics 101*(2), 264–296.

Dybvig, P. H., H. K. Farnsworth, and J. N. Carpenter (2009). Portfolio performance and agency. *The Review of Financial Studies 23*(1), 1–23.

Gómez, J.-P. and T. Sharma (2006). Portfolio delegation under short-selling constraints. *Economic Theory 28*(1), 173–196.

Harris, L. and E. Gurel (1986). Price and volume effects associated with changes in the S&P 500 list: New evidence for the existence of price pressures. *The Journal of Finance 41*(4), 815–829.

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

Kashyap, A. K., N. Kovrijnykh, J. Li, and A. Pavlova (2019a). The benchmark inclusion subsidy. Technical report, National Bureau of Economic Research.

Kashyap, A. K., N. Kovrijnykh, J. Li, and A. Pavlova (2019b). Is there too much benchmarking of asset managers?

Leippold, M. and P. Rohner (2011). Equilibrium implications of delegated asset management under benchmarking. *Review of Finance 16*(4), 935–984.

Li, C. W. and A. Tiwari (2009). Incentive contracts in delegated portfolio management. *The Review of Financial Studies 22*(11), 4681–4714.

Lines, A. (2016, November). Do institutional incentives distort asset prices? London Business School.

Ma, L., Y. Tang, and J.-p. Gómez (2019). Portfolio manager compensation in the US mutual fund industry. *The Journal of Finance 74*(2), 587–638.

Ou-Yang, H. (2003). Optimal Contracts in a Continuous-Time Delegated Portfolio Management Problem. *The Review of Financial Studies 16*(1).

Palomba, G. and L. Riccetti (2012). Portfolio frontiers with restrictions to tracking error volatility and value at risk. *Journal of Banking & Finance 36*(9), 2604–2615.

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

Shleifer, A. (1986). Do demand curves for stocks slope down? *Journal of Finance 41*(3), 579–590.

Sockin, M. and M. Z. Xiaolan (2019). Delegated learning in asset management.

Sotes-Paladino, J. and F. Zapatero (2018). Riding the bubble with convex incentives. *The Review of Financial Studies 32*(4), 1416–1456.

Stoughton, N. M. (1993). Moral Hazard the the Portfolio Management Problem. *The Journal of Finance 48*(5).

Stucchi, P. (2015). A unified approach to portfolio selection in a tracking error framework with additional constraints on risk. *The Quarterly Review of Economics and Finance*.

Van Binsbergen, J. H., M. W. Brandt, and R. S. Koijen (2008). Optimal decentralized investment management. *The Journal of Finance 63*(4), 1849–1895.