

Review of the Active Management of the Norwegian Government Pension Fund Global

20 January, 2014

Andrew Ang

*Ann F. Kaplan Professor of Business
Columbia Business School*

Michael W. Brandt

*Kalman J. Cohen Professor of Business Administration,
Fuqua School of Business, Duke University*

David F. Denison

Former President and CEO of the Canada Pension Plan Investment Board

[This page is intentionally left blank]

Preface

It is a great honor for the authors to conduct a review of the active management of the Norwegian Government Pension Fund Global (“GPFG” or the “Fund” hereafter). We are mindful that this is a service for the Norwegian people and that it addresses questions of importance to their economic future. The Fund is an extraordinary endowment, and the evaluation we undertake on their behalf seeks to provide input to help realize the full potential of this investment.

Contents

Preface	1
Contents	2
Mandate	4
Approach	5
Executive Summary	7
Section I: Theoretical Foundation and Review	13
I.A: Diversification and Rebalancing	13
I.B: Examples of the Rebalancing Premium	17
I.C: Role of Benchmarks	19
I.D: Dynamic Factors	23
I.E: Security Selection	27
I.F: Recommendations of Ang, Goetzmann and Schaefer (2009)	28
I.G: Summary	29
Section II: Review of Active Management of the Fund	31
II.A: NBIM's Approach to Active Management Since 2008	33
II.B: Fund Returns	35
II.C: Active Returns	36
II.D: Choice of Factors	41
II.E: Active Factor Exposures	47
II.F: Operational Reference Portfolio	51
II.G: Benchmark Deviation Risk	56
II.H: Summary	61
Section III: Active Mandate	63
III.A: Comparative Advantages and Disadvantages	63
III.B: Frameworks for Delegating Investment Management	69
III.C: Investment Framework at Canada Pension Plan Investment Board	71
III.D: Challenges and Approaches of Investments in Alternative Assets	77
III.E: Private Equity	82

III.F: Infrastructure	86
III.G: Summary	90
Section IV: Recommendations	91
IV.A: Report Each Stage of Value Added	92
IV.B: Adopt the “Opportunity Cost Model” for Active Management	100
IV.C: Increase Risk Taking of Active Management	107
IV.D: Summary	114
Appendix A: Empirical Robustness	116
Appendix B: Total Portfolio Approach	118
Appendix C: Benchmarking Real Estate	120
References	124
Tables	130
Figures	140

Mandate

The report should include:

- i) Review of Norges Bank's historical performance in the management of the GPFG, including:*
 - *Analyses of the Fund and the equity and fixed income portfolio's return and risk relative to the benchmark, with emphasis on the last five years*
 - *Breakdown of performance by main strategies/activities, within the limits of available data*
 - *Analyses of risk-adjusted performance, including a discussion of choice of methodology for risk adjustment and choice of relevant risk factors*
- ii) Review of how delegation to Norges Bank can improve GPFG's expected return and risk relative to the current strategic benchmark. The review should cover investment opportunities both within and outside of GPFG's current investment universe and include:*
 - *Theoretical and empirical foundation*
 - *Return and risk characteristics and investment capacity (scalability)*
 - *GPFG's comparative advantages or disadvantages based on GPFG's special characteristics*
 - *Comparable investors use of strategies to exploit these investment opportunities*
- iii) Based on the analysis in i) and ii), discuss potential implications for the Ministry's mandate to Norges Bank. This part of the report should include a discussion of benchmarking, relevant risk measures and risk budgets, reporting requirements and how other funds have implemented comparable strategies.*

Approach

In addressing the mandate for this report, we begin with a section on the theoretical foundation covering how value is created at different stages in the investment process. The rationale for this section is to lay the groundwork to interpret the investment decisions of the asset owner represented by the Ministry of Finance and the fund manager: Norges Bank, the Central Bank of Norway, through its asset management unit, Norges Bank Investment Management, NBIM. The notion of “active” vs. “passive” management is not cleanly demarcated in any delegation model, including the one taken by Norway. Rather, there exists a continuum of decisions—which are all active, but to varying degrees—involved in the investment process. The section also reviews the concepts advocated by the previous report on active management for the Norwegian Government Pension Fund Global (hereafter referred to as the “Fund”) written by Andrew Ang, William Goetzmann, and Stephen Schaefer in 2009.

Section II reviews the active performance of NBIM relative to the benchmark given by the Ministry of Finance. We concentrate on the last five years of active performance, although the short sample, the changing benchmark, and the fact that NBIM has changed its investment management structure mean that the results over this short sample have to be interpreted with caution. For this reason, we also conduct a qualitative review in addition to our quantitative analysis. The section summarizes the relevant factor exposures of NBIM’s active management, and also characterizes the risk of the active returns in terms of volatility (often called “tracking error”) and other higher moments.

In Section III, we review the current delegation model between the Ministry and Norges Bank, and we compare the delegation structure with other comparable funds. We discuss the Fund’s comparative advantages and disadvantages, and highlight how funds with similar characteristics have been organized to exploit investment opportunities. The section concentrates on private equity and infrastructure because the Fund is not currently invested in these asset classes. We discuss how peer funds undertake these investments.

In the final Section IV, we make recommendations to use the Fund’s comparative advantages to benefit the investment management process. We discuss how the mandate given by the Ministry of Finance to Norges Bank could be developed further.

Data and Other Inputs

In preparing the report we have relied upon historical monthly return numbers from NBIM on an aggregate fund basis as well as on performance broken down by fixed income and equity strategies. The data covers the period from January 1998 to June 2013. Our performance numbers are stated in the Fund’s currency basket. Given the sensitive nature of this

information it is not all detailed in our report. In addition we have used return information from global capital markets including fixed income and equity indexes, volatility indexes, liquidity indexes, and other external data we deemed relevant to assess risk and return. These have been converted where necessary into the Fund's currency basket for appropriate comparison.

We have had several meetings with NBIM management to discuss NBIM's investment approach, especially with regards to the Operational Reference Portfolio (ORP). In addition, NBIM has been responsive to our regular requests for data and addressing technical issues. These meetings form part of our qualitative assessment of active management and have been incorporated in our recommendation for the mandate. We thank NBIM for their time, openness, and responsiveness during this process.

Acknowledgements

We thank Gerlof de Vrij, Cam Harvey, Antti Ilmanen, Sung Cheng Chih, and Ludovic Phalippou for helpful comments on the analysis.

Executive Summary

In theory, long-horizon investors can potentially add value through some combination of broad diversification, regular rebalancing, being exposed to factor risks, and security selection. The Fund benefits to some extent from all of these. The Fund's benchmark is assigned by the Ministry of Finance and ensures broad diversification of investments; beyond this, NBIM has used insights to counteract some of the limitations of standard index benchmarks in order to achieve additional diversification benefits. The rebalancing guidelines in the mandate given from the Ministry of Finance cause the Fund to act counter-cyclically ("buying low and selling high"), maintain optimal portfolio weights, and capture a rebalancing premium. NBIM has recently implemented a factor-based approach with initial focus on size and value-growth premiums in equities and duration risk in bonds in order to harvest resulting returns over the long term. Finally, we observe that NBIM continues its long standing focus on security selection in public markets, both through internal programs and external managers. It has expanded into real estate, which provides diversification benefits not obtainable in the stock and bond universes and is an additional venue to practice security selection.

Our review of the active performance of the Fund primarily focuses on the past five years and excludes any consideration of the real estate program given its nascent state. We find that NBIM's risk profile has reduced significantly: the maximum ex-ante allowable volatility of the benchmark deviation ("tracking error") was reduced by half a percentage point to 1.0% per year by the Ministry of Finance in 2009, and NBIM currently utilizes only half of that limit. Based on requirements in the mandate given by the Ministry of Finance, Norges Bank has imposed restrictions on leverage which do not permit NBIM to increase the exposure to risky assets. NBIM has reduced its use of derivatives. It has also sharply restricted its roster of external managers, who are chosen to complement its internal capabilities.

The returns generated by active management have been consistently positive in the post-2009 sample contributing on average 0.10% per month, which is highly statistically significant. Over this period, the Fund's monthly Reward-to-Risk ratio (the ratio of average returns to realized volatility) was 0.47. We caution that it is dangerous to extrapolate this excellent performance going forward because there were many investment opportunities due to financial market dislocations during and after the financial crisis which are no longer available. We believe a monthly Reward-to-Risk ratio in the range 0.10 to 0.15 is more realistic for the Fund in the long run.

It is inappropriate to compute direct performance measures that adjust for systematic risk factors because the fund manager was not given a benchmark with factor risk exposures. For the same reason, it is also inappropriate to measure the fund manager's performance with a

levered benchmark. Nevertheless, it is relevant to investigate how the Fund's active returns relate to well-known sources of systematic risk. We find that approximately 60% of the variance of active returns is attributable to dynamic risk factors, although these exposures have varied in form and extent over the period. This proportion of factor risk exposure is not out of line with other types of investors. We believe that this significant factor risk is highly suitable because collecting systematic factor premiums plays to the comparative advantages of the Fund: they can be harvested in scale, some have verification horizons stretching to a decade, and skill is required to implement factor investing strategies. As an example of the latter, in 2011, NBIM initiated an "Operational Reference Portfolio" (ORP) designed to harvest factor premiums over long horizons. While still in its early development and changing over time, NBIM has made impressive progress in implementing a scalable platform. We regard the ORP as a significant advancement in NBIM's investment capabilities.

As specified in our mandate, we review approaches to active management and delegation of authorities of institutional investors comparable to the Fund. There are relatively few investors with similar attributes to the Fund, so rather than a wide canvass we focus on a select group with members who are widely regarded as global leaders in this sector. These leaders share the same "structural" advantages of scale, long time horizon, and ability to hold illiquid assets. In addition, they exhibit a set of common "developed" advantages arising from a fund's investment management capabilities. Developed advantages are important enablers of success in active management, especially for private market investing. NBIM has built some developed advantages already. We see no institutional impediments that would prevent NBIM from creating other capabilities; indeed, although we did not examine the real estate program in detail, we surmise that some of these are either in place or under development in that area.

We posit a series of considerations that are relevant for determining the appropriate degree of delegation across the three parties involved in a typical investment management framework: the asset owner, the oversight body of the asset manager (the "board"), and the asset manager itself. In particular, we note that organizations such as large U.S. endowments and many of the large Canadian pension plans that combine an orientation to active management, a belief in sources of risk to be harvested outside standard public benchmarks, the time-varying nature of those risk premiums, and the presence of an experienced and highly qualified board and investment manager, have high levels of delegation from the asset owner to the board and asset manager.

We explore in detail one such investment and delegation framework, the "Opportunity Cost Model," which is particularly well suited to long-horizon investors. This model is extensively used by the Canada Pension Plan Investment Board (CPPIB) and GIC Private Limited (GIC), formerly known as the Government Investment Corporation of Singapore. In the Opportunity

Cost Model, the asset owner retains responsibility for the most important decisions about the fund: in particular, the asset owner chooses a Reference Portfolio that provides highly scalable and low-cost, passive exposures to equity and bond indexes. The Reference Portfolio represents the necessary amount of systematic risk to achieve the fund's objectives. The asset owner also specifies any investment constraints and determines an active risk limit governing the extent of deviation from the Reference Portfolio. In turn, responsibility to invest beyond the Reference Portfolio is given to the board and the asset manager. The fund manager is incentivized to make investments that provide superior risk-adjusted returns net of costs relative to the Reference Portfolio, where the cost implicitly accounts for the appropriate, from a risk perspective, funding of the deviation from the Reference Portfolio. For example, any dollar that could be invested in private real estate is benchmarked against the opportunity costs of investing that dollar in a mix of public equities or bonds. Thus, any active investment that deviates from the Reference Portfolio is benchmarked net of fees against public market securities in the Reference Portfolio used to fund that investment.

The Opportunity Cost Model represents a compelling alternative to traditional asset class portfolio construction and investing. It provides a consistent and coherent framework for analyzing and benchmarking investment decisions across all private markets as well as within some areas of public markets. The focus of active management becomes the component of returns that cannot be obtained in public market investments as captured in the Reference Portfolio benchmark. This raises the bar and accountability for active management.

Traditional "endowment" models for strategic asset allocation specify fixed, or slowly-moving, allocations to different asset classes, each having their own benchmark. In the Opportunity Cost Model, the manager is free to take any deviations from the Reference Portfolio based on a fair-valuation outlook, rather than being forced to maintain positions when the asset class valuations are very expensive or cheap. There is also no need to separately choose different benchmarks for each asset class—although these benchmarks can be used in alternate ways like assessing the skill of the fund manager's real estate division, for example. Since all active positions are benchmarked against the zero-cost, easy-to-implement passive stock and bond exposures in the Reference Portfolio, the difference between the fund and Reference Portfolio returns directly represents the value-added of active management.

We caution that while the Opportunity Cost Model is conceptually relatively simple, it is one that is challenging to operationalize. Two preconditions are a long horizon of the asset owner (one of the Fund's structural advantages) and that all parties in the delegation process require expertise (a developed advantage). The long horizon is necessary because there can be significant differences, sometimes stretching to a decade, between the returns in private markets and the corresponding funding mix of Reference Portfolio assets. The fund manager

must have considerable experience and skill, along with access to excellent information systems, to find, evaluate, and monitor attractive investments with return components that cannot be obtained in public stock and bond markets. We believe that NBIM already has many of these resources and could develop the remaining capabilities to implement the Opportunity Cost Model successfully.

All the leading global institutional investors we examine have extensive and well-established private market investment programs including private equity and infrastructure in addition to real estate. We are not aware of any definitive study that shows large institutional investors have obtained sustained positive, risk-adjusted returns in private market investments, although there are some recent reports that provide some confirmatory evidence. The institutions in the Fund's peer group confirm achievement of value-added returns over multi-year periods and intend, in general, to increase their exposures to private markets. We describe the special selection skills these investors have developed to choose managers, the various approaches to direct investing they employ in private markets, and their benchmarking and measurement practices. If the Fund were to move into private equity and infrastructure, a best practice of benchmarking is to measure the investment opportunities foregone in public markets. The Opportunity Cost Model is also an appropriate model of benchmarking the Fund's nascent real estate program.

We believe there is scope to increase the risk taking of active management. The empirical analysis shows that active returns constitute a tiny contributor to overall Fund performance: since 2009, benchmark returns account for over 99% of the variance of total Fund returns. We recommend an increased volatility limit of benchmark deviations given that increased risk taking can be devoted to areas which add long-term value for the Fund: superior diversification outside benchmarks based on market-capitalization weights and allocations to systematic factor risks. Under the Opportunity Cost Model, active risk budgets should also increase because real estate investments would be included in the deviations from benchmark. The Fund has had a positive experience to date with active management, and some of the increased risk taking limit would not be taken because it would be used as a prudent risk buffer. Increased risk taking should be accompanied by greater transparency about the areas where the risk of active management is being deployed. At the moment, the asset owners' risk appetites are mostly expressed through the typical size of deviations from the benchmark ("tracking error"), which reflects only one dimension of risk, and greater total active risk should be accompanied by explicit limits on the risks that really matter—downside risk relative to the benchmark.

We conclude our report with the following recommendations:

1. The Fund should report risk and returns from each discrete phase of its investment process.
2. The Fund should adopt the Opportunity Cost Model and corresponding delegation framework. This can be done in a staged process over several years, and should initially be implemented in the Fund's real estate program.
3. The amount of risk taking by active management should be increased, and the deviations from public market benchmarks should include real estate exposures. This should be combined with risk limits on the whole distribution of active returns, especially limits on downside losses relative to the benchmark.

[This page is intentionally left blank]

Section I: Theoretical Foundation and Review

The voluminous asset pricing literature emphasizes three key concepts relevant for long-term investors: diversification, rebalancing, and taking on exposures to factor risks which yield long-run risk premiums. In addition, investors with superior skill can add further value via security selection. Passive index funds enable investors to hold a diversified portfolio at low cost, and a rebalancing rule allows their expected utility to remain optimal over time. Implementing both requires making active decisions. Financial theory and empirical evidence suggests that investment management can add value beyond diversification and rebalancing by exposing the portfolio to factors which earn long-run risk premiums, but they come at the expense of potential short-term losses relative to a passive index portfolio. Finally, a talented fund manager can select stocks with superior risk-return trade-offs—these stocks have expected returns that are not fully explained by their factor risk exposures. Identifying and harvesting appropriate factor risk premiums were a special focus of the previous report on the active management of the Fund by Ang, Goetzmann and Schaefer (2009).

I.A: Diversification and Rebalancing

Diversification and rebalancing add value over long periods.

Diversified portfolios that hold assets spread over many different industries and sectors, geographies, and styles of investing enable investors to obtain superior risk-return trade-offs compared to concentrated portfolios comprising only a few assets. When asset returns are not perfectly correlated, bad outcomes of some assets can be offset by good outcomes of others, which leads to reductions in risk at the portfolio level. Thus, some of the random fluctuations of individual returns can be diversified away by holding combinations of many assets. This insight led to the development of the Capital Asset Pricing Model (CAPM), where in equilibrium only non-diversifiable risk is priced.¹

The most diversified portfolio in the CAPM paradigm is labeled the “market portfolio,” but in practice this special portfolio is unobserved. In fact, the early literature questioned whether the CAPM could be empirically tested because the true market portfolio includes assets not found in liquid, public markets.² Under the CAPM, the market is the only factor which determines the expected returns of assets: securities with larger exposures to market risk have higher average

¹ See, among others, Markowitz (1952), Sharpe (1964), Lintner (1965), and Mossin (1966).

² This is referred to as the Roll (1977) critique. Stambaugh (1982) and others argue that as long as the set of traded assets is sufficiently highly correlated with assets that are not traded, then inferences about the CAPM using only traded assets are virtually identical to those obtained using the true market which includes non-traded assets.

returns and they tend to lose more, on average, when the market factor declines. To compensate for these losses during bad times, investors require that the assets with large market exposures have high expected returns.

Today, investors can hold well-diversified portfolios at low cost thanks to the rise of index funds pioneered by Wells Fargo and Vanguard during the 1970s. It is interesting to note that index funds were developed 10-20 years after the CAPM was developed in academic finance. Passive index funds are based on market indexes, which are usually constructed by independent index providers who make decisions on the securities to be included in the indexes, their constituent weights, and how these weights change over time. Different index providers have different criteria for selecting securities and build different weighting schemes. While somewhat arbitrary, diversified portfolios based on these indexes are useful benchmarks because passive fund management based on these indexes can be run at effectively zero cost. We emphasize that the selection of an investable market portfolio is a choice that must be made by an investor—there is no unique *investable* market portfolio.

In theory, a passive market benchmark requires no trading if there is no entry or exit of stocks over time. That is, most of the trading associated with a passive market benchmark occurs when new securities are created or redeemed, or companies leave the index. (There are some other minor rebalancing considerations due to changing liquidity and other corporate events.) In this sense, the market portfolio is a “static” factor. While the passive market reflects the equilibrium outcome of the *average* investor, Samuelson (1969) and Merton (1969, 1971) show that it is optimal for an *individual* investor to rebalance to fixed asset weights, or risk exposures, over time under the realistic assumption that asset returns are not predictable. The fixed asset class weights are determined by a solution to a one-period portfolio choice problem, and they are often proxied by diversified index benchmarks in different asset classes.


Rebalancing across asset classes ensures that an investor can maintain optimal expected utility. It is an *active* strategy: it entails increasing risk positions in securities that have recently declined in price and decreasing exposures to asset classes that have recently appreciated in value. Rebalancing is thus a *counter-cyclical* investment strategy. Rebalancing is often hard to implement because behavioral biases frequently cause investors to react in the opposite way. Most investors are *pro-cyclical*: they are fearful to buy securities which have lost value and reluctant to part with securities whose prices have rapidly increased. Rebalancing is a simple, but powerful, tool to ensure investors act counter-cyclically.

In the presence of mean-reverting returns, long-term investors have additional strategies available to them that short-horizon investors do not, which academics often label “long-term

hedging demands” following Merton’s terminology.³ While rebalancing forces investors to “buy low and sell high,” long-term investors can opportunistically buy even more when prices are very depressed and sell when securities are grossly over-valued.

While rebalancing is optimal from the perspective of maximizing an individual’s expected utility, a rebalancing strategy also allows wealth to increase exponentially even when the prices of individual assets are stationary. This “rebalancing premium” is a consequence of diversification at each rebalancing date lowering a portfolio’s variance, while a buy-and-hold portfolio is eventually dominated by one, or a small number, of assets and consequently has higher risk. The lower variance allows wealth to cumulate faster under rebalancing.⁴

We can decompose a portfolio return into several components measuring the choice of the diversified index benchmark and rebalancing, as shown in the following figure:



Denoting a portfolio return as r , the true, but non-investable market portfolio as r_{mkt}^{true} , the choice of an investable diversified index return as $r_{bmk}^{passive}$, and the rebalanced index return

³ See Brandt (2009) and Wachter (2010) for recent summaries of the dynamic portfolio choice literature. A buy-and-hold passive strategy has payoffs that are linear in the stock price. Rebalancing is a short volatility strategy and produces payoffs that are concave over the stock price. See chapter 4 of Ang (2014). A canonical model with “positive hedging demands” where long-term investors increase their portfolio holdings of equities relative to short-term investors when expected returns are high and prices are low is Campbell and Viceira (1999).

⁴ The mathematical finance literature also terms this “volatility pumping” or the “variance drain” (see, for example, Evstigneev and Schenk-Hoppé, 2002). Erb and Harvey (2006) is a reference in the finance literature.

as r_{bmk}^{rebal} , we can represent the figure mathematically as:

$$r = \underbrace{r - r_{bmk}^{rebal}}_{\text{"Active" return}} + \underbrace{r_{bmk}^{rebal} - r_{bmk}^{passive}}_{\text{Rebalancing gain}} + \underbrace{r_{bmk}^{passive} - r_{mkt}^{true}}_{\text{Choice of passive benchmark relative to true market}} + r_{mkt}^{true} . \quad (1)$$

In equation (1), the first two stages of a diversified benchmark and a rebalancing rule are often considered to be “passive” management. These actually involve active choices. A choice must be taken to implement the true, diversified market portfolio because the full passive market portfolio is unobservable. Once the diversified portfolio is chosen, the optimal investment strategy is to periodically rebalance to this passive benchmark. How often rebalancing is done and how it is implemented are again active decisions. Since the true market is a theoretical construct, only the rebalancing gain is directly observable. However, different passive benchmarks are not the same, and certain investors might prefer one over another depending on the rules for security inclusion, entry and exit, and liquidity screens. We define a “passive” investment strategy as one that can be implemented by pure systematic trading, involving low turnover, and approximately zero cost.

The final decision in equation (1) is the traditional “active” return. This is a return in excess of the rebalanced, diversified benchmark. The fund manager deliberately deviates from the benchmark to take advantage of further investment opportunities. This could well involve superior diversification not available in the benchmark and rebalancing gains that are hard to capture systematically. It could also involve taking on factor risk and acting on superior information, which we discuss further below. The deviation from benchmark is a more “active” decision than the choice of the diversified benchmark and the rebalancing rule because it typically involves more frequent discretionary decisions. Another difference from “passive” index strategies is that active management usually involves higher turnover and cannot be done at zero cost.

In the Fund’s context, the rebalanced benchmark is chosen by the Ministry of Finance. Section II reviews NBIM’s active return relative to this benchmark. However, the breakdown in equation (1) makes clear that the rebalancing decision and the choice of the passive benchmark are also active investment decisions; the whole process of investment management is a series of active decisions. In Section IV, we recommend more transparency associated with each of these stages.

I.B: Examples of the Rebalancing Premium

Rebalancing has added value in practice.

We show how the rebalancing decision has added value in historical samples. This exercise is useful because the benefits of rebalancing tend to be less well known than the benefits of diversification. In later sections, we discuss how the rebalancing regime has become a comparative advantage for the Fund. As part of our recommendations in Section IV, we advocate additional transparency of the Fund's investment process to highlight the choices made to capture the rebalancing premium, which can be done along the lines of the analysis in this section.

I.B.1: U.S. Stocks and Bonds

Figure 1 illustrates the gains to rebalancing. We rebalance over stocks and bonds in the United States. We have total returns and market capitalizations for each asset class. We take all stocks listed on the NYSE, Amex, and NASDAQ exchanges from CRSP. For bonds, we use all U.S. Treasury securities (bonds, notes, and bills) reported by Ibbotson Associates. We rebalance every quarter to fixed 60% equity and 40% bond weights. Panel A plots the cumulated wealth of the rebalancing strategy starting with \$1 at the beginning of June 1961 and ending in June 2013. For comparison, we plot the return of a passive strategy which holds market capitalization weights in stocks and bonds.⁵ The returns on the passive strategy are computed with market capitalization weights at the beginning of each month. Technically, the passive strategy rebalances only to the extent that the indexes change market weights in response to "free float" adjustments and entry and exit of issuers, otherwise it does not rebalance. The passive strategy represents the returns accruing to the average "market" investor, as represented by the CRSP and U.S. Treasury positions. To compare the rebalancing and passive strategy on the same risk basis, we rescale the passive strategy so that it has the same ex-post standard deviation as the rebalanced strategy.

Panel A shows that rebalancing has added value over the sample period: \$1 in the rebalanced strategy grows to \$79 at the end of the period whereas \$1 in the pure passive strategy cumulates to \$39. The annualized Reward-to-Risk ratio, defined as the ratio of the annualized average return to annualized standard deviation of returns, of the 60/40 rebalanced strategy is

⁵ Sharpe (2010) refers to such a passive portfolio as an "adaptive" asset allocation policy. Another way to compute the value of rebalancing is to compare a rebalancing strategy with a strict buy-and-hold strategy, as done in chapter 4 of Ang (2014), which is dependent on an assumed starting point. Another advantage of the comparison done here is that it explicitly compares the investor who rebalances with the equilibrium market investor who does not rebalance.

higher, at 0.92, compared to the Reward-to-Risk ratio of 0.78 in the passive, non-rebalanced strategy.⁶

Rebalancing powerfully induces counter-cyclical behavior. Panel B of Figure 1 plots the portfolio weight in stocks in the solid black line. By construction, we bring back the portfolio weight to 60% in the rebalanced strategy. The passive, un-rebalanced strategy, in contrast, slowly wanders around as the market capitalizations of equities and bonds change. During the late 1990s, equity weights increase as equities rapidly rose in price during the internet boom. The rebalanced strategy cuts back on equity exposure during this time. In the early 2000s, the passive equity weight decreases as the prices of internet stocks fall, while the rebalanced strategy buys equities to maintain its 60% holding. During the financial crisis in 2007-2008, the passive weights in equities decrease markedly as Treasuries outperform. The rebalanced strategy moves in the opposite direction and increases equity holdings. This leads to larger gains than the passive strategy as equity prices rebound after the world financial system stabilizes.

While Figure 1 highlights that rebalancing adds value relative to the passive strategy because of mean reversion in returns, mean reversion is not necessary for rebalancing to add value. As long as the two assets are not perfectly correlated, a fixed-weight portfolio of stocks and bonds can reduce variance. It is the reduction of return variance that contributes to higher expected wealth over the long term. In practice, there is some mean reversion, albeit small, in asset returns, and in this case, rebalancing contributes even more value than in environments where returns are not predictable.⁷ Counter-cyclical strategies like rebalancing can further help if one asset class becomes over- or underpriced relative to another. Baker and Wurgler (2000), among others, report that companies opportunistically issue more equity when stock prices are high. These are times when future stock returns tend to be low. Changing relative valuations also result from time-varying risk premiums; when the average investor becomes more risk averse, equity prices fall relative to bonds, and vice versa. Long-term investors with more stable risk appetites can benefit from these fluctuating prices.⁸

⁶ We use the Reward-to-Risk ratio, as opposed to the Sharpe ratio where the risk-free rate is subtracted from the average return, as a way to relate the total return generated to the amount of risk taken. One needs to be aware, however, that the Reward-to-Risk ratio contains a cash return which has substantially varied over our sample period. The fact that the Sharpe ratio requires specifying a risk-free return is another disadvantage: most commonly used risk-free rates like U.S. T-bills are not the risk-free return for the Fund.

⁷ See Rapach and Zhou (2013) for a recent summary.

⁸ A seminal model along these lines is Campbell and Cochrane (1999). In equilibrium, not everyone can rebalance and those that do not miss out on collecting a rebalancing premium, see Chien, Cole and Lustig (2012).

I.B.2: World Stocks and Bonds

A global example is shown in Figure 2. We take global bonds in the Barclays Global Aggregate Bond Index and global stocks in the FTSE World Index. Like Figure 1, we consider a strategy with fixed 60% equity and 40% weights with quarterly rebalancing. Figure 2 displays the same patterns for world stocks and bonds as Figure 1 does for U.S. stocks and bonds for rebalanced and passive strategies. Again we can interpret the passive strategy as the returns to the “average” market investor, as proxied by the Barclays and FTSE indexes.

Panel A of Figure 2 plots cumulated wealth starting with \$1 at the beginning of February 1994 for the rebalanced strategy in the solid, black line, and the passive strategy with market capitalization weights in the red, dashed line. We scale the passive return to have the same ex-post volatility as the rebalancing strategy. Wealth in the rebalancing strategy is larger at the end of the sample, September 2013, for the rebalanced strategy. The end-of-period wealth is \$4.06 for the rebalancing strategy compared to \$3.84 for the passive, un-rebalanced strategy. The annualized Reward-to-Risk ratios for the rebalancing and passive strategies are 0.82 and 0.70, respectively.

Panel B of Figure 2 plots the fraction of the portfolio held in stocks for both the rebalancing and passive strategies. We observe the same counter-cyclical pattern for the rebalancing strategy as in Figure 1, but the dip in the passive equity weight is even more pronounced as many non-U.S. stock markets fell dramatically during the financial crisis. The large jump in the equity portfolio weights at the end of 2012 is due to the relative increased coverage in the FTSE benchmark universe at this time. This highlights that indexes used to proxy the diversified market portfolio are also active decisions.

In summary, although the rebalancing benchmark is traditionally considered “passive,” it consists of active decisions that govern rebalancing and the choice of the passive index. Rebalancing rules can add value, on average, compared to non-rebalanced, passive holdings.

I.C: Role of Benchmarks

An ideal benchmark is well diversified, able to be rebalanced, and can be cheaply implemented.

We have discussed the first two stages in the investment process—choosing a diversified portfolio and a rebalancing regime. Both stages involve setting benchmarks and maintaining optimal exposures to them.

There are several properties of an ideal passive benchmark, listed in order from the most to the least important:

1. Well diversified

A good starting point is a (float-adjusted) market-capitalization benchmark. Abstracting from the specific rules adopted by different index providers, a market index represents the investment opportunity of the average, or representative, investor. In terms of implementation, market weights also are an excellent approximation for investment capacity. An important shortcoming of a market index is that it is restricted to securities in liquid, public markets. Thus, there are diversification benefits available by moving to non-traditional asset classes. The Fund has started to move into these areas with the first real estate investment in 2011.

There is a second source of diversification benefits not available in traditional market indexes. Index providers have to make somewhat arbitrary decisions on countries or types of securities to include in an index, and their weighting schemes. An investor may consider a country investable even though it is not included in an index. It may be better for an investor to continue holding securities that exit an index rather than immediately selling them. Long-term investors not needing immediate liquidity may benefit from taking larger positions in securities than the “free float” weights in market indexes. Although this is a second-order effect compared to adding private market assets, it is not negligible. NBIM attempts to take advantage of these further diversification benefits, which we discuss in Section II. In particular, Section II.G demonstrates that there are significant differences between the risk properties of indexes with different geographic construction rules.

2. Reflect optimal exposure to risk premiums available in liquid markets

Market capitalization weights may not represent the optimal diversified exposure for a given investor because of the presence (or lack) of liabilities, the properties of an investor’s income stream, or because there are factor risks that command a premium other than the market risk premium (see Section I.D) which the investor has a comparative advantage in harvesting. (We discuss some of the Fund’s comparative advantages in Section III.A.) These considerations, combined with the investor’s risk bearing capacity, lead to investor-specific benchmarks which are different from pure market capitalization weights. The most important deviation from the passive market is the proportion of equities and bonds in the strategic benchmark.

In setting the Fund’s benchmarks, the Ministry of Finance starts with off-the-shelf market-capitalization indexes for stocks and bonds from FTSE and Barclays, respectively.

The actual benchmark over-weights Europe and down-weights the U.S. and Canada in the equities portfolio and uses GDP weights in the sovereign bond portfolio. In equities, the weighting scheme allows the Fund to obtain greater geographical diversification than allowed by the market; non-U.S. capital markets generally have lower ratios of publicly traded capital to GDP. In bonds, the tilts away from market capitalization emphasize economic importance; market values of sovereign debt also reflect other considerations such as seigniorage and reserve requirements. Together with the rebalancing framework, these conscious choices in diversification are the most important investment choices for the Fund and are appropriately anchored at the highest level—Parliament.

3. Replicable

The benchmark should be transparent, in that there are clear rules that allow the index to be replicated. Obtaining an index from a third-party index provider ensures it cannot be manipulated by a fund manager.

Ideal indexes represent holding period returns, and thus reflect the actual amount of wealth that can be accumulated by an investor. Only benchmarks in traditional equity and bond markets meet this requirement. Benchmarks in real estate, private equity, infrastructure, and other private markets do not represent investable returns. These benchmarks are often based on fund or security values that are estimated or appraised, and thus do not immediately reflect true market values. It is possible to buy all the securities in a public market index, but impossible to buy a small slice of all constituent properties in a direct real estate index. Consequently investors in private markets face more idiosyncratic risk than the indexes for these markets. In alternative asset markets, the roles of traditional active and passive management cannot be separated—everything is active. While illiquid market indexes are useful in the investment process, they are less useful in benchmarking value-added performance for a long-term investor. We discuss an alternative framework in Section III.C.

4. Can be rebalanced

A related property to replication is that the index should be straightforward to dynamically rebalance. For very large investors, this is not a trivial consideration even for market-capitalization indexes because of transactions costs and adverse selection. The latter refers to traders who, knowing that an investor is forced to trade securities at a certain time, create temporary scarcity forcing the investor to buy at higher prices or sell at lower prices than would otherwise be the case without the need to rebalance. A fund manager should therefore have some leeway to optimally implement rebalancing.

As rebalancing forces investors to go against their pro-cyclical natural tendencies, the rebalancing regime must also be robust.

Because they are not investable, private market indexes cannot be rebalanced. Illiquid assets can, however, be incorporated into a rebalancing scheme. A simple, but highly inferior, method is to ignore illiquid asset positions and rebalance only over publicly traded assets. We discuss alternative mechanisms undertaken by the Fund's peers in Section III.

5. Implementable at close to zero cost

The costs of investability for a rebalanced benchmark should be so negligible that a good working approximation is that the benchmark can be achieved at zero cost. In this way, the benchmark represents a viable passive alternative to active management. The market-weighted indexes in liquid, public markets are designed with this goal in mind. As investors move away from market capitalization weights to non-market weighted exposures, it becomes more difficult to create a tradable passive strategy at near-zero cost.

No index proxy for illiquid, private markets meets this criterion; private markets are characterized by incomplete information, high transaction costs, and the lack of centralized markets where all participants can trade at the same prices.

Perhaps the most important advantage of a benchmark implementable at close to zero cost is that when a fund's returns are compared to the benchmark, the net deviation from benchmark after costs directly represents the value-added from active management.

6. Has a long history

A long time series of benchmark returns is desirable for risk management and to help set the strategic benchmark. With long return histories, investors can better estimate risks—especially on the downside—allowing them to anticipate potential losses. This helps create more robust strategic allocations. Indexes can be “backfilled,” so one concern is the creation of an index that, even if it fulfils all the previous criteria, is that it is based on an investment strategy that has recent good performance which may not persist. One way to mitigate such “data mining” is a compelling economic story explaining why the benchmark's losses in bad times should be rewarded with a long-run risk premium. In Section II, we examine the factor exposures in the Fund's active management strategy taking into account these considerations.

I.D: Dynamic Factors

By taking on risk exposures that other investors seek to avoid, a long-term investor can harvest dynamic factor risk premiums.

I.D.1: What are Factors?

A dynamic factor strategy is a style of investing which systematically takes positions away from the market portfolio and maintains these deviations over time. In doing so, investors take on non-diversifiable risk and earn long-run average returns. Investors with positive exposure to factor risks reap premiums for being exposed to risk that other investors seek to avoid. It is similar to collecting an insurance premium during good times and suffering losses during bad times.

In this report, the term “dynamic factor,” or simply “factor,” refers to an investment strategy that consistently goes long securities with certain characteristics which are offset by short positions in securities with the opposite characteristics. The securities with similar characteristics tend to move with each other, both in under- and out-performing the market portfolio. Academic theory, empirical studies, and investing experience have identified classes of securities that have consistently higher average payoffs than the market portfolio. On the flipside, there are stocks that consistently underperform the market when the first group of securities outperforms. Stocks with low prices relative to fundamentals (value stocks), for example, beat stocks with high prices relative to fundamentals (growth stocks), on average, giving rise to a *value-growth* premium. Over the long run, stocks with past high returns (winners) outperform stocks with low or negative past returns (losers), leading to *momentum* strategies. Securities that are more illiquid have high average excess returns compared to more liquid securities reflecting an *illiquidity* premium.⁹

Dynamic factors combine long positions in these specific securities which outperform with short positions in other securities that underperform. These are often referred to as “Fama-French” factors for the researchers that initially formed long-short portfolios designed to capture the value-growth premium and the size premium:¹⁰

Value-Growth Premium = Value stocks minus growth stocks

Size Premium = Small stocks minus large stocks

⁹ See Ang, Goetzmann and Schaefer (2011) for a summary of these risk premiums and their economic stories. See also Ilmanen (2011) and Ang (2014).

¹⁰ See Fama and French (1993).

We can collect other premiums in a similar manner:

Momentum Premium = Winning stocks minus losing stocks

Illiquidity Premium = Illiquid securities minus liquid securities

Credit Risk Premium = Securities with high default risk minus securities with low default risk

Low Volatility Risk Premium = Stocks with low volatility minus stocks with high volatility

Implied Volatility Risk Premium = Selling put options offset by stocks or calls to produce market-neutral positions

These are examples of many long-short factors identified in the finance literature.¹¹ In Section II, we use several of these factors to describe the risk exposures of active management taken by NBIM.


These risk premiums are not a free lunch, however, because while they offer attractive rewards over the long term, they can underperform in the short run. Factor risk premiums are compensation for enduring more severe losses during bad times than the average investor. Factor risk premiums can also arise through the actions of non-rational investors which are not arbitrated away. Losses from factor investing strategies can persist for several years: value stocks persistently underperformed growth stocks, for example, for the last few years of the 1990s, and high yield bonds underperformed investment-grade bonds for the first few years of the early 2000s (see Section II.D). The underperformance of factors can be particularly severe, as was the case for credit and volatility risk factors during the 2007-2008 financial crisis.

Factor investing is “dynamic” as opposed to the “static” market-capitalization equity and bond benchmarks. The latter require no trading when there is no entry or exit of companies. To harvest factor risk premiums, an investor moves away from a passive market-weighted portfolio within an asset class. The factors require dynamic trading because their holdings of securities vary: as companies rise in price relative to their fundamental value, for example, they cease to be value stocks, and are dropped from a value factor. Growth stocks initially have high prices relative to fundamental value, but have low average returns. When their prices sufficiently decrease, they no longer become overpriced and exit from a growth factor. Thus, the value-growth factor requires dynamic trading to maintain long exposures in value stocks and short exposures in growth stocks.

¹¹ Harvey, Liu and Zhu (2013) survey hundreds of such long-short factors.

I.D.2: Dynamics Factors as Part of the Investment Process

Factors based on pervasive systematic risks are attractive for many investors because they manifest across many different asset classes and geographical regions. Augmenting the previous diagram to include risk factors, we have:



If we denote the factor tilts away from the rebalanced benchmark portfolio as r_{bmk}^{factor} , then we can extend the return decomposition in equation (1) to

$$r = \underbrace{r - r_{bmk}^{factor}}_{\text{"Active" return}} + \underbrace{r_{bmk}^{factor} - r_{bmk}^{rebal}}_{\text{Factor return}} + \underbrace{r_{bmk}^{rebal} - r_{bmk}^{passive}}_{\text{Rebalancing gain}} + \underbrace{r_{bmk}^{passive} - r_{mkt}^{true}}_{\text{Choice of passive benchmark relative to the true market}} + r_{mkt}^{true}, \quad (2)$$

where now the factor exposure term, $r_{bmk}^{factor} - r_{bmk}^{rebal}$, makes explicit that some of the portfolio's return comes from taking on dynamic factor risks. The "active" return is now interpreted as the returns of actively managed portfolio generated in excess of the fund's factor benchmark. We will refer to this as "security selection" in excess of a fund's factor risk exposures, and comment further on this component below.

For a large, long-term investor, it is reasonable to break out the effect of factors in a decomposition of the return. Like the passive, rebalanced benchmark, the factor risk premiums may be harvested mechanically—dynamic factor strategies follow well-defined, systematic rules. In terms of the criteria for ideal benchmarks discussed in Section I.C., factors offer diversification benefits and factor benchmarks can be optimally constructed to match investors'

preferences and risk tolerances.¹² A market-weighted passive benchmark is dominated by one source of risk—equity market direction (which is the only risk factor in the simple CAPM), and relaxing the restriction of market weights allows other dynamic factors to come to the fore and thus improve diversification. Most factor indexes are replicable and can be rebalanced. The verification horizons of some strategies span decades, and there exist large literatures in finance investigating the equilibrium origins of the factor risks premiums. These characteristics make factors more like traditional “passive” equity and bond market capitalization-based indexes.

However, only some factors have large capacity and can be invested in at scale. Only some factors have relatively long histories. Dynamic factor investing requires more skill than tracking benchmarks based on market capitalization weights. Factor strategies cannot be implemented at zero cost, although they are much cheaper than the fees commanded by most active mutual funds and hedge funds. Factors cannot be implemented as cheaply as market index funds; these costs must be taken into consideration when deciding whether to harvest a particular factor risk premium. For a very large investor, the costs may be particularly onerous because of small capacity and large adverse selection. These characteristics make factor strategies more like traditional “active” management.

Because of the tension between these two, dynamic factors occupy a “middle ground,” so we have drawn the diagram above with factors sitting in between the traditional rebalanced passive benchmarks and “active” management. To date, the Fund has not adopted value-growth, size, momentum, illiquidity, short volatility positions, or other similar factors in the strategic benchmark. The Ministry of Finance has investigated the feasibility of including factor exposures in the Fund’s benchmarks, like de Jong and Driessen (2013) on liquidity premiums and the MSCI report, “Harvesting Risk Premia for Large Scale Portfolios,” examining value, size, low volatility, and momentum. At the very least, a component of active returns—deviations from the passive benchmark—can be attributed to factors. In the 2012 Report to the Storting, the Ministry’s assessment was that the benchmark index in the mandate from the Ministry to Norges Bank should not be adjusted for systematic risk factors and any strategies for exploiting systematic risk factors should form part of the operational management of NBIM. In Section II, we investigate how much of the Fund’s return in excess of benchmark is related to factor risks.

¹² For a review of the academic literature, see Ang, Goetzmann and Schaefer (2011). For practitioner reviews of the recent performance of factor strategies and their diversification benefits, see Bender et al. (2010) and Ilmanen and Kizer (2012).

I.E: Security Selection

Successful security selection requires skill.

The last part of the investment process is security selection. This involves taking positions in securities in excess of a fund's benchmark reflecting static diversification opportunities, dynamic rebalancing, and systematic, time-varying tilts away from market capitalization weights to gain exposure to various factor risk premiums.

If markets were purely efficient, then it is impossible to add value by security selection. According to the "efficient markets" view, security prices are not forecastable, which is sometimes termed the "random walk" model.¹³ This view is obsolete. The modern view, due to Grossman and Stiglitz (1976), is that it is impossible for financial markets to be purely efficient; there are opportunities for smart investors skilled in collecting, processing, and analyzing information to select securities exhibiting expected returns in excess of their factor risk exposures. These investment opportunities are often fleeting and limited in size. The actions of informed investors drive the market back to a dynamic equilibrium of near-efficiency.

The literature identifies several characteristics of successful investors in undertaking security selection. Investors with stable funding sources and few constraints on liabilities can pursue these opportunities without the risk that they have to disinvest at the wrong time. While many mispricing opportunities disappear after several days, only a patient, long-term investor can pursue those that might take months—or years—to yield a reliable risk-adjusted profit. Hiring high-quality people with appropriate investment skills, building good systems, and trading to minimize transaction costs are prerequisites to playing the security selection game well. At the same time, the literature offers repeated warnings on how difficult it is to beat risk-adjusted benchmarks. Given the difficulties in generating profits on a risk-adjusted basis from security selection, active management based on cheaply harvesting factor risk premiums may be more sustainable in the long run for many investors.

We will show that the Fund's comparative advantages, discussed in Section III.A, give it some scope to successfully engage in security selection. Opportunities for security selection abound in private markets—where the advantages in superior information, analyzing that information, and finding and monitoring skilled managers are magnified. In fact, the original motivation for holding illiquid, alternative assets as originally espoused by Swensen (2000) (often called the "endowment model"), is that private markets offer more scope for security selection because they are less efficient than public markets. The lack of appropriate benchmarks means that

¹³ The Efficient Markets Hypothesis traces back to Cowles (1933) and was formalized by Fama (1965) and Samuelson (1965). This is reviewed at length by Ang, Goetzmann and Schaefer (2011).

measuring security selection value-added in private markets is challenging (see Section I.C). We also discuss how other comparable funds tackle these issues in Section III, which inform our recommendations for the Fund’s active management mandate in Section IV.

I.F: Recommendations of Ang, Goetzmann and Schaefer (2009)

The previous review of active management emphasized the importance of factor risks and appropriate verification horizons for reporting.

A study of the active management of the Fund was conducted in 2009 by Professors Andrew Ang, William Goetzmann, and Stephen Schaefer. They examined the Fund’s performance from January 1998 to September 2009, and found that active management had a small impact on the Fund’s overall returns because the permitted deviations from the benchmark were small. Active management did not detract from the value of the Fund and prior to 2008, contributed a significantly positive portion to the Fund’s returns.

The small amount of active management, however, had large exposure to a parsimonious number of well-recognized, systematic factors. This was detected statistically ex post, and Ang, Goetzmann and Schaefer (2009) also showed that it was possible to have measured the factor exposure on an ex-ante basis. Factor exposures were extremely important in explaining the active losses during the financial crisis over 2008-2009, and largely responsible for the rebound in active returns in the latter part of 2009. The most important factors were liquidity, volatility, credit, and a size factor. Importantly, the study found that the active losses during the financial crisis were broadly in line with what should have been anticipated given the Fund’s exposure to these factors. If the Fund’s asset owners had some knowledge of the potential risks associated with these factor exposures, then the Fund’s losses during the financial crisis might have been within expected loss limits. Active losses in active management during the financial crisis were surprising partly because the Fund’s factor exposures, implicitly generated by NBIM’s investment activities, were not made public.

Ang, Goetzmann and Schaefer’s main recommendation was that the Fund should focus more attention on factor risk. Since the Fund’s active management generated significant factor risk loadings, the Fund should adopt a more intentional approach to managing its strategic and dynamic factor exposures. Put another way, the Fund should move beyond the Fund’s passive asset class benchmarks to also adopt dynamic factor benchmarks. In the context of equation (2), the report’s recommendation was to manage the contribution of the factor risks,

$r_{bmk}^{factor} - r_{bmk}^{rebal}$, to the overall Fund’s return.

Ang, Goetzmann and Schaefer claimed that factor investing has several advantages. Making clear the factor exposures allows better communication of the risks being taken by active management. In the optimal case, the amount of factor risk would be informed by the preferences of the asset owner, the Norwegian people. NBIM's skill in building indexes and tracking them, its professionalism, and scale, makes it well suited to creating factor benchmarks. The factor benchmark also allows a clearer separation of the value added by active management, since it would separate the active return into a portion based on delivering a factor risk benchmark cheaply and a portion based on active decisions to hold securities with weights different from the factor benchmark. In the latter, value could also be added in timing the exposure of the factors. The report, however, did not advise on an explicit mandate for the Fund incorporating factors or specifically state which factors would be appropriate to adopt in the Fund's benchmarks. As we discuss the recent performance of the Fund in Section II, we highlight the role of factors and discuss how NBIM has managed factor risk exposures—a choice that has so far been delegated to it and not placed in the Fund's benchmark.

Ang, Goetzmann and Schaefer identified several drawbacks in changing the Fund's management to including a factor-based approach. There is no consensus on a complete set of factors and how to optimally construct them. Not surprisingly, there are no widely accepted factor benchmarks and a dearth of factor risk investments available in the very large scale required for the Fund. There are additional complications in communicating the factor risks to the public.

Since different factors have different verification horizons and risks, Ang, Goetzmann and Schaefer also recommended that different components of the overall return be reported in different categories reflecting their horizon-specific risk profiles. This allows certain strategies which only tend to pay off over many years, possibly decades, to be benchmarked appropriately. Even the simple equity premium requires decades to exhibit a significantly positive return in excess of bonds. Having different categories for the different investment strategies facilitates better risk management.

I.G: Summary

Long-horizon investors can create value by:

1. Holding a well-diversified portfolio, but since the most well-diversified market portfolio is not observable, the investor must choose a benchmark to proxy the market,
2. Rebalancing, which is a counter-cyclical strategy that tends to “buy low” and “sell high,”

3. Taking on (potentially time-varying) exposures to factor risks that offer attractive rewards over long horizons, but are subject to short-term drawdowns, and
4. Security selection, which requires skill.

All of these are active decisions, to a greater or lesser degree.

Section II: Review of Active Management of the Fund

This section describes NBIM's active management of the Fund both empirically using historical return data and descriptively based on conversations with NBIM personnel. We measure active management as the Fund's returns in excess of the benchmark. We refer to this as the "active return," and it is the same as the "benchmark deviation" defined in Section I. NBIM's mandate is to "seek to achieve the highest possible return, net of cost, measured in the currency basket of the investment portfolio" subject to several constraints. The most important of these is the volatility of the deviations allowed from the benchmark ("tracking error").¹⁴ The other constraints include limits on the investment universe, a restriction that leverage cannot be used to increase the Fund's exposure to risky assets, and a requirement to take into account a country's fiscal strength when investing in sovereign bonds, among others. We were provided with monthly benchmark and actual returns for the equity portfolio, the bond portfolio, and the combined Fund consisting of the equity and bond portfolio for the period January 1998 to June 2013.

In the context of Section I, the benchmark provided by the Ministry of Finance embeds the diversification and rebalancing choices of the Fund, but does not treat separately the factor risk exposures or security selection. Part of NBIM's active management is to take on factor exposure. Its investment mandate states: "The equity and bond portfolios should be composed in such a way that the expected relative return is exposed to several different risk factors." As Section I discusses, harvesting dynamic factor risk premiums is a complement, but not a substitute, for plain-vanilla static equity and bond risk premiums. In NBIM's current mandate, the Fund's strategic benchmark does not take into account such systematic risk premiums. Thus, part of our analysis entails selecting appropriate factors to which the Fund's active returns can be compared, and measuring the magnitude of the factor risk exposures.

With the exception of our discussion of the Fund's active management approach in Section II.A, our review is necessarily backward looking and covers a period where NBIM's investment strategies and management style have developed considerably. There are at least three distinct regimes in the sample: the period prior to the financial crisis in 2007-2008 where active returns were significantly positive, the losses during the financial crisis which consolidated toward the end of 2008, and the period post 2008. The first two of these regimes were studied at length in Ang, Goetzmann and Schaefer (2009).

¹⁴ In this report, we use the term "benchmark deviation" or "active return" to refer to the Fund return minus the benchmark return. We do not like the term "tracking error" even though it is pervasive in industry because it has the connotation that the benchmark deviation is an "error." In practice, deviations from benchmark are deliberately taken and are not "errors." In this report we only use the term "tracking error" in the context of the Fund's current risk constraints.

As per our mandate, we focus on the last five years, which we define to be the period from January 2009 through June 2013 when our data end. This period does not constitute exactly 60 months, but we choose January 2009 as a start date of our main analysis in order to have a clean post-financial crisis sample. Likewise, we choose the pre-crisis period from January 1997 through December 2006 as a comparison reference sample. Even with these carefully chosen subsamples, the fact that the management of the Fund has developed new capabilities partly in response to its experience during 2007-2008, and that these developments are on-going, means that our conclusions based on past performance need to be interpreted in the context of an evolving strategy. Stated another way, our findings may be more relevant looking backward than going forward.

The short sample period, particularly for our analysis on the last five years, inevitably leads to low statistical power. Moreover, the analysis we conduct is “ex post” meaning that the Fund’s risk factor exposures are measured over the entire sample or subsamples. The financial crisis’s extreme factor returns and resulting large realization of tracking error make it difficult to perform a reliable ex-ante performance evaluation for the post-crisis period. This is because the factor loading estimates to evaluate performance in 2009 and 2010 would be dominated by what happened during 2007 and 2008. In contrast, we are more interested in what the factor exposures were post crisis and going forward. For this reason, we refrain from conducting inference on “rolling samples,” which would be dominated by this episode and are not a good indication of true conditional behavior at different points during this period.

Another important qualification is that the choice of our factors is made at least partially on an ex-post basis, in that we have selected our factors taking into account return histories and the current state of the finance literature. We may have chosen different factors if we were to conduct our study at an earlier point in time. Since most of our factors were used in the Ang, Goetzmann and Schaefer (2009) study, this potential factor snooping bias is less of a concern for the post-crisis than for the pre-crisis sample. Nevertheless, although many of our factors are the same as the earlier study, the fact that we choose a slightly different set of factors to attribute the Fund’s active returns demonstrates this point clearly. We discuss the choice of factors in Section II.D.

Given these caveats regarding the statistical analysis, we augment our quantitative discussion with commentary on the development of the Fund’s active management since 2008, economic background of the sources of factor risks behind the active returns, and general descriptions of the market environment over different periods in the sample. In particular, our assessment in Section II.F of NBIM’s newly developed framework for investing in factors is largely qualitative. Our qualitative analysis is based on academic work, publically available studies produced and distributed by NBIM, and our discussions with senior personnel at NBIM.

We conduct our analysis on currency basket denominated returns for the Fund. Our factors are converted to the currency basket by combining NOK denominated returns on the Fund and currency basket returns provided by NBIM. For the return attribution to factors, we also denominate factors that are tradable (which is all but one, more details below) in the currency basket. That said, since most computations are done using excess returns, i.e., portfolios that are long one dollar of notional currency exposure and short one dollar of notional currency exposure, the returns are naturally relatively currency hedged and therefore the currency denomination is of second order importance (except where explicitly noted below). Since we were not provided with individual security positions and trades of the Fund, or the ex-ante risk estimates for different active strategies, we do not make an assessment of the Fund's risk budgeting process. We have not reviewed NBIM's management structure, or the management capabilities of its senior officers.

Finally, we do not review the Fund's performance in real estate. The Fund was authorized to invest up to 5% of its assets in real estate in 2010, with the allocation to fixed income (which has a long-run target of 40%) being reduced accordingly. As of June 30, 2013 the Fund held 0.91% of its total portfolio in real estate. Given the short history of real estate investments and its relatively small scale, we were not provided with real estate returns. As reported in our conversations with NBIM, the investments in real estate are very different from investments in liquid, public markets and NBIM has spent much effort in building internal capacity and competency. It is proceeding to grow its real estate investments with caution and waits opportunistically for appropriate joint venture partners. The tracking error of the Fund does not include its real estate investments. It is possible the Fund may possibly better use its good track record of professional management in liquid markets in evaluating and managing its investment opportunities in illiquid asset markets. We discuss such a framework employed by other large funds in Section III.

II.A: NBIM's Approach to Active Management Since 2008

Since the financial crisis, the Fund has de-risked, only judiciously employs external managers, and has created the Operational Reference Portfolio.

We begin with a general discussion and assessment of NBIM's strategic plans for active management and risk budgeting with a specific focus on how these have changed since 2008. This discussion is based on conversation with NBIM management as well as published reports by NBIM and the Ministry of Finance.

While there are parts of NBIM’s investment policy that have changed—some as part of learning to cope with the extraordinary demands of the financial crisis and others that were already in place and were implemented during and after this period—NBIM’s core philosophy for active management has not: NBIM continues to aim to select securities based on fundamental analysis to generate value in excess of its benchmark. NBIM has changed its investment practices in several areas, however: it reduced its use of derivative strategies, simplified and reduced its use of external managers, lowered its active risk both ex ante and ex post, and has started to explore systematic factor strategies.

In the first area, the Fund concentrates on simpler instruments with lower leverage. The Ministry of Finance allows the Fund to use derivatives related to equity and fixed income. Based on a requirement in the mandate from the Ministry, the investment mandate issued by Norges Bank’s Executive Board to NBIM states that leverage may not be used “with a view to increase the investment portfolio’s exposure to risky assets in the equity and fixed income portfolio,” and that leverage associated with fund structure “may not be with the aim to increase the investment portfolio’s financial exposure to risky assets.” Net leverage as a percent of market value was upwards of 10% prior to 2009, and is now close to zero.¹⁵

Second, there are many fewer external managers employed since the financial crisis. As of June 2013, external managers represent approximately 4% of the Fund. In comparison, Ang, Goetzmann and Schaefer (2009) found that prior to 2009, externally managed funds reached a peak of approximately 20% of assets in fixed income and 40% in equities. NBIM seeks to use external managers only where it does not have internal competency, which includes emerging and frontier markets where first-hand knowledge of domestic conditions is essential. Because of the small proportion of external management employed by the Fund, we do not separately examine external active equity returns.¹⁶

Third, the Ministry of Finance lowered the Fund’s maximum volatility of benchmark deviations (“tracking error”). The ex-ante tracking error limit was reduced from 1.5% per annum to 1.0% in 2009. The prior 1.5% was a strict maximum limit. Under the new mandate, although the 1.0% tracking error is a limit, the risk budget is understood to apply during normal times. During extremely volatile times, the 1.0% tracking error may be exceeded. Even within these boundaries, the amount of ex-post active risk taken by NBIM has generally fallen substantially since 2008, as we show below in Section II.G.

¹⁵ See <http://www.nbim.no/en/About-us/governance-model/executive-board-documents/ceo-investment-mandate-for-the-government-pension-fund-global/> and <http://www.nbim.no/en/press-and-publications/Reports/2012/nbim-performance-results-gips-report-/>

¹⁶ Ang, Goetzmann and Schaeffer (2009) extensively reviewed external managers in the period up to 2009 and found significant inferior performance of external fixed income managers.

Finally, the development of factor strategies is a significant new development by NBIM. A new “Operational Reference Portfolio” (ORP) includes factor tilts away from the benchmark index, but it also includes other elements including smart implementation of rebalancing, expansion of the investment universe to take advantage of diversification opportunities beyond the benchmark constituents, and pursuing some required elements of the Ministry of Finance’s mandate to take into account the fiscal strength of sovereign issuers. Our view is that the construction of the ORP is an important strategic development of the Fund, and one that is consistent with the Fund’s comparative advantages. We assess the ORP in Section II.F.

II.B: Fund Returns

The Fund’s Reward-to-Risk ratio was 0.20 per month since January 1998, and the Reward-to-Risk ratios were significantly higher in the pre- and post-crisis periods.

To put the quantitative analysis of the active returns of the Fund in perspective, Panel A of Table 1 summarizes the performance of the Fund over the full sample period (January 1998 to June 2013), the pre-crisis period, and the post-crisis period which is the main focus of our analysis. We present these statistics not to comment on the overall performance of the Fund, which is dominated by the performance of the benchmark given by the Ministry of Finance, but to provide a reference point for the active returns analysis. Since the beginning of 2009, the Fund returned 11% per annum with an annual standard deviation of 9.3%.¹⁷ These statistics are each nearly twice what the Fund experienced during the pre-crisis period, partially reflecting the growing allocation to equities over time. The Reward-to-Risk ratio of the Fund is nearly identical during the pre- and post-crisis periods, suggesting that the Fund is being proportionally or better compensated for the additional risk it has taken on over a relatively long period of time.¹⁸

Panel A of Table 2 breaks the Fund’s performance into its fixed income and equities components (which are currently weighted approximately 40%/60% in the benchmark portfolio, with the real estate portion coming from the Fund’s bond allocation). The fixed income allocation has very similar total return statistics during the pre- and post-crisis samples, with average annualized returns of 5.3% and 6.5% and annualized return standard deviations of 3.1% and 3.5%, respectively. In contrast, the equity portfolio of the Fund experienced considerably better post-crisis performance, returning 14.2% per annum with an annual

¹⁷ We annualize the monthly frequency numbers in Table 1 by multiplying both the mean and variance by 12.

¹⁸ Since the return on holding cash was substantially higher pre-financial crisis than post, the Sharpe ratio is actually significantly higher for the post crisis sample. In that sense, the additional return earned for taking additional risk has increases more than proportionally.

standard deviation of 15.6%. The monthly Reward-to-Risk ratio since the beginning of 2009 is 0.26 as compared to 0.16 in the pre-2007 sample.

It may seem inconsistent that, at first glance, the Fund's overall Reward-to-Risk ratio was not larger during the post-crisis period since a greater fraction of the Fund was allocated to equities—an asset class that performed nearly twice as well as during the pre-crisis period. The explanation lies in the correlation between the fixed income and equity benchmarks. In the pre-2007 period, stocks and bonds were negatively correlated at -0.38, meaning they were good diversifiers in a balanced portfolio. From 2009 onwards, the correlation between stocks and bonds is instead close to zero. (During the financial crisis the stock-bond benchmarks have a slightly positive correlation of 0.07.) Therefore, post 2009, the relative outperformance of the increased equity allocation was almost exactly offset by the diminished diversification benefits between the fixed income and equities benchmarks. While it is not in our mandate to comment explicitly on the appropriateness of the benchmark portfolio mix, it is worth noting that the correlation structure on which traditional benchmark constructs are based are not set in stone and should be monitored over time.

II.C: Active Returns

Active management has significantly added value relative to the Fund's benchmark over the full sample, and pre and post the financial crisis.

We define the active return as the difference between the returns of the Fund (or its fixed income and equity components) and the benchmark. This definition is consistent with the active mandate, as NBIM is tasked with beating the benchmark given by the Ministry of Finance. As part of its active management, NBIM may seek to increase or reduce its exposure to the benchmark or take on other systematic factor exposures. We consider both of these actions as part of the active return, and we therefore do not adjust the active returns for the resulting benchmark or other factor exposures. We elaborate on this point in Section II.E.

II.C.1: Variance Attribution of Total Returns

We first put the (absolute) magnitude of the active returns in perspective relative to the overall Fund returns. Ang, Goetzmann and Schaefer (2009) suggest a variance decomposition which measures the fraction of variance of the overall Fund returns, or its fixed income and equities components, attributed to the benchmark portfolio returns, with the remainder being attributed to active returns. Note that for this calculation, the active returns should ideally be uncorrelated with the benchmark returns, which is not the case when the active decision

involves leverage or deleverage relative to the benchmark by over-or under-allocating to higher beta securities.¹⁹ In that case, the active returns, computed as the difference between the Fund returns and benchmark returns, would be perfectly correlated with the benchmark returns, and our variance decomposition would attribute 100% of the Fund return variance to the benchmark suggesting no active management.

Panel B of Table 1 reports the variance decomposition for the Fund. Over the entire sample, active returns accounted for 0.7% of the variance of fund returns. The variance contribution of active management is even smaller in the post crisis sample at 0.4%. It is clear from these results that the Fund returns are overwhelmingly dominated by benchmark returns, an observation also made by Ang, Goetzmann and Schaefer (2009). The fact that the Fund is as a first-order approximation a well-diversified index fund does not, however, belittle the role of active management. Given the size of the Fund, even a little active management can add substantial economic value to its beneficiaries. The tiny proportion of active risk in the Fund's total risk does suggest that there is scope to increase the amount of active risk, a point we return to in Section IV.

Panel B of Table 2 repeats the variance decomposition for the fixed income and equities components of the Fund. These results show more striking differences between the portfolio components. First, the contribution of active returns to the variance of fixed income returns is considerably larger. Over the whole sample, active management contributed 10.7% of the variance of fixed income returns and only 0.3% of the variance of equity returns. Second, there are dramatic changes through time. Pre-crisis active returns contributed 0.6% and 0.4% to the variance of fixed income and equities, respectively. Post crisis, these contributions to variance changed to 16.2% and 0.0%, respectively. The latter is 0.05% with two significant digits. The last number is particularly striking—effectively none of the variance of equity returns is due to active management.

The percentage attribution of the variance of the Fund's returns to active management for fixed income is higher than found by Ang, Goetzmann and Schaefer (2009). In Table 2, active returns account for 10.7% of the variance of fixed income returns compared to approximately 3% found by Ang, Goetzmann and Schaefer. This difference is due to the currency return: the previous report on active management used returns expressed in NOK, and we use returns expressed in the Fund's currency basket. The R^2 of a regression of fixed income returns onto the benchmark over the full sample is 98.0% when returns are converted to NOK, which gives a variance

¹⁹ The Fund is explicitly prohibited from using notional leverage, where actual or synthetic borrowing is used to increase the notional exposure to equities. We use the term leverage in a risk sense where the exposure to risky assets can be increased by overweighting more risk sensitive securities and underweighting less risk sensitive securities.

attribution similar to Ang, Goetzmann and Schaefer. Currency denomination has a relatively larger effect on fixed income returns because of the comparably low volatility of the asset class.²⁰ (Equity returns are substantially more volatile than fixed income, and currency denomination plays a much smaller role.)

The dramatic increase in the relative variance of active fixed income returns during the crisis is familiar from Ang, Goetzmann and Schaefer (2009), who documented large active losses in the asset class through the financial crisis. That elevated variance contribution continued in 2009 except, as we will illustrate shortly, manifested mostly as active gains. Equally striking is the considerably lower contribution of active management to the variance of equity returns (less than 0.1%). As our subsequent analysis will show, this result is partly due to less active risk taken and also partly due to the increased correlation of active risk with the benchmark portfolio. We return to both these effects below.

II.C.2: Active Performance

Panel C of Table 1 presents summary statistics of the active returns for the Fund and Figure 3 plots the cumulative active returns. In this and other figures, the vertical dotted line marks the beginning of the post-crisis subsample. On average, the active return on the fund was three basis points but the post-crisis period enjoyed a larger active return of 10 basis points per month. The monthly standard deviation of active returns was 22 basis points for the whole sample, 11 basis points before the financial crisis, and 21 basis points afterward. The resulting monthly Reward-to-Risk ratio was 0.12 for the whole sample period, 0.35 pre 2007, and an impressive 0.47 post crisis. Note that these numbers are as large as for the overall Fund returns in discussed in Section II.B.

Figure 3 illustrates these findings graphically. Pre 2007, the active return cumulated steadily. During 2007 and 2008, there was a large drawdown of the active return causing the cumulated active return to be negative at the end of 2008. Post crisis, there was a complete recovery of the losses and gains continued to the end of the sample. Figure 3 shows that the recovery from the financial crisis accounts for the majority of the active returns over the last five years. This suggests we need to be cautious about extrapolating the Fund's recent performance going forward. Many of the financial market dislocations that facilitated the post-crisis recovery are no longer available, particularly in active fixed income. The financial crisis was also a period where the low asset prices coincided with very high expected returns and high volatility, and expected returns have decreased since that time. We believe a more realistic Reward-to-Risk

²⁰ A large literature on optimal currency hedging investigates this issue. Black (1990) is an early reference. A more recent reference is Campbell, Serfaty-de Meiros and Viceira (2010).

ratio over the long run is closer to the Fund's record over the whole sample period, which is around 0.10 to 0.15.

Figure 4 presents the active returns of the Fund from a different perspective. It shows in red dots the Fund's monthly active returns (as opposed to the cumulative returns in Figure 3). We highlight the Fund's post-financial crisis active returns in Figure 5. In both Figures 4 and 5, we clearly notice the large negative returns in 2008 and then a mirroring sequence of large positive returns throughout the recovery in 2009. The figures also illustrate the extent to which the standard deviation of active returns is dominated by the 2008-2009 period. Post 2010, the distribution of active returns resembles that observed in the pre-2007 sample. The large losses in 2008 and subsequent large gains in 2009 also account for the negative skewness of active returns over the whole sample and the positive skewness of active returns in the post-crisis subsample as reported in Table 1.

Ang, Goetzmann and Schaefer (2009) noticed and commented on the high autocorrelation of active returns through the financial crisis. Figures 4 and 5 show clear sequences of positive and negative returns, and the autocorrelations reported in Panel C of Table 1 confirm this observation. The autocorrelation of active returns was 0.52 over the whole sample and 0.62 post crisis. These statistics are considerably larger than the pre-crisis autocorrelation of 0.22. We return to this observation below.

In Figures 6, 7, and 8 we contrast the cumulative and monthly returns of fixed income and equities. Figure 6 plots the cumulative returns for fixed income in Panel A and for equities in Panel B. It is clear from these plots that the Fund-level drawdown in the financial crisis as well as the subsequent recovery in 2009 was largely driven by fixed income investments. We also note that the fixed income recovery has leveled off over the last few years, presumably as dislocations in global fixed income markets realigned and active opportunities have become less abundant, whereas the returns on active management in equities are relatively stronger. Figure 7 again plots as red dots the monthly returns on active fixed income and equities in Panels A and B, respectively. These plots show the active returns of the two components of the Fund on the same scale, showing that the larger variance attribution for fixed income is not only caused by the fact that fixed income investments are less volatile, but also that active returns in that asset class are more volatile.

Figure 8, which zooms in on the post-crisis period, also illustrates another fact we previously alluded to. The crisis subsample period is short, and the distribution of active returns over this period is not stable. The year 2009 stands out as an unusually volatile year for active fixed income. Starting in 2010, however, there has not been a single monthly active return in either fixed income or equities that exceeded 0.5%. The same is observed at the overall Fund level in Figure 5. Not coincidentally, 2010 is the year the Ministry reduced the tracking error limit and

methodology that controls the amount of active risk NBIM can take. We return to the issue of risk in Section II.G.

Despite the relatively short samples, the statistical support for a positive contribution from active management is strong for the pre- and post-crisis periods, with p-values less than 0.01 as reported in Panel C of Table 1. The support for a positive contribution over the whole sample is borderline statistically significant with a p-value of 0.04, obviously due to the large underperformance and increased volatility of active returns during the financial crisis. Panel C of Table 2, which considers fixed income and equities separately, shows that the full-sample versus subsample dichotomy is driven by fixed income. For equities, the average contribution from active management is statistically significant for all three sample periods, whereas for fixed income it is only significant if we exclude the crisis period. While the overall sample mean is positive, the variance of active fixed income returns was so large during the crisis period that the p-value on the mean is a relatively large 0.21 for the whole sample period.

Panel C of Table 2 also sheds further light on the autocorrelation of active returns. The Fund-level autocorrelations discussed above are predominantly driven by the high persistence of fixed income returns. This persistence, in turn, is probably due to the illiquidity of certain fixed income securities during the financial crisis and subsequent gradual recovery of liquidity in the post crisis period. (We do not, however, have access to security-level holdings to further test this hypothesis.) As we will illustrate below in Section II.D, the elevated autocorrelations of active returns also coincide with relatively higher autocorrelations in the returns on the systematic risk factors that the Fund is exposed to. In fact, the autocorrelation of active equity returns in Panel C of Table 1 is low compared to the recent persistence in systematic equity risk factors.

Abstracting from the financial crisis, we conclude that active management in both fixed income and equities has positively contributed to the returns of the Fund, and when the financial crisis is excluded, the outperformance is statistically significant. Although the outperformance is a few basis points, which appears small in a relative sense (recall the average return on the Fund was 92 basis points per month over the last five years), it is important to consider the size of the Fund. In absolute terms (meaning in NOK or currency basket valuation), the gains are large. The outperformance in terms of returns is partly small due to the tracking error risk limits imposed on the Fund; there may be scope to increase the contribution of active management by raising these limits, which we take up in Section IV. While the recent performance of fixed income has clearly been outstanding, we caution extrapolating the recovery from the financial crisis. We consider a monthly Reward-to-Risk ratio of 0.2 for fixed income as being more realistic going forward.

Having observed a statistically significant and economically meaningful contribution by active management to the total returns of the Fund, we now examine to what extent this contribution can be attributed to exposures to established factors explaining the cross section and time series of fixed income and equity returns. In contrast to the academic literature on performance evaluation, our goal is not to reduce the “alpha” of the manager by labeling part of the active return as “alternative beta” or “smart beta.” Quite the contrary, we view active returns originating from exposures to long-standing and well-established systematic factors as a sustainable source of value-added for a long-run investor—perhaps more sustainable than unidentifiable and potentially less structural sources of outperformance like security selection. As Section I discusses, factor risk premiums represent a valuable source of returns that active management can provide relative to a passive, rebalanced benchmark.

II.D: Choice of Factors

We select factors reflecting systematic risk premiums in fixed income, equity, and derivative markets.

While a large literature in financial economics has documented the existence of various systematic factors other than the market portfolio, there is no consensus on the number and types of systematic factors which are appropriate for investing, benchmarking, or risk management. We choose our factors with the following considerations:

1. Consistency with Ang, Goetzmann and Schaefer (2009)

Some of these factors are slightly different due to data revisions, or the discontinuation and development of a new index in one case.

2. The factors are recognized as having large spreads in average returns by the academic literature and industry practice and represent long-short positions

Along this line, we select factors with benchmarks that can be traded in large size. Several indexes reflecting some of our factors have been recently developed—including some investigated by a study done by MSCI commissioned by the Ministry of Finance—but these are not constructed to be zero-cost investment portfolios. We also prefer factor indexes in existence at January 1998 (the start of the Fund returns).

3. The factors span the activities undertaken by NBIM to generate value

Given our short sample, we choose a parsimonious set of factors which aims to span the range of active strategies taken by NBIM. On the one hand, including factors that reflect mispricing opportunities or factor risk premiums not deliberately chosen by the fund manager can lead to “spurious” attributions. On the other hand, the factors should capture broad patterns in active management activities, even though the fund manager is not directly trading a particular factor. For example, fixed income structured products or asset-backed securities are highly exposed to credit risk, as are sovereign bond issues of certain emerging markets. Many active strategies are exposed either directly or indirectly to volatility risk. Asset overlays and even rebalancing decisions different from benchmark also have exposure to volatility risk.

We use the following factors:

Term: Difference between long- and short-maturity U.S. Treasury bond returns

CreditAa: Difference between Aa and Treasury bond returns

CreditBaa: Difference between Baa and Aa bond returns

CreditHY: Difference between high yield and Baa bond returns

Liquidity: The difference in the spread between illiquid and liquid U.S. Treasury bonds, reflecting periods of high and low liquidity

ValGrth: Difference in returns between “value” stocks and “growth” stocks

SmLg: Difference in returns between small and large stocks

Mom: Difference in returns between stocks with past high returns (“winners”) and stocks with low past returns (“losers”)

LowVol: Difference in returns between stocks with stable returns (“low vol”) and stocks with volatile returns (“high vol”)

SellVol: Captures excess returns to a strategy of selling implied volatility in derivative markets

Table 3 lists the definitions of these factors and their data sources. We categorize the factors into fixed income factors (Term, CreditAa, CreditBaa, CreditHY), the Liquidity factor, and equity factors (ValGrth, SmLg, Mom, and LowVol), and an implied volatility factor (SellVol). Because of data availability, especially at the start of the sample period, all the factors are based on U.S.

data, except the ValGrth and SmLG factors. Figure 9 plots cumulated returns on the factors expressed in the Fund's global currency basket, except for the Liquidity factor which is not a return series and is therefore not redenominated.

II.D.1: Fixed Income and Liquidity Factors

The fixed income factors Term, CreditAa, CreditBaa, and CreditHY are identical to those used by Ang, Goetzmann and Schaefer (2009). They reflect returns to investment strategies capturing duration, or maturity risk, in bonds (Term), and different credit risks (CreditAa, CreditBaa, and CreditHY).²¹ We plot cumulated returns on the fixed income factors in Panel A of Figure 9. The duration premium reflected in Term, as shown in the solid black line, has been the most consistent performer among the fixed income factors. The reward to holding high yield bonds relative to Baa bonds is negative over the sample, as shown in the blue long-dashed line. Panel A shows the large drawdowns in the credit factors during the 2008-2009 financial crisis and a subsequent strong recovery after the financial system stabilized. In contrast, there were strong returns to the Term factor over 2008-2009 as investors sought safety in long-maturity U.S. Treasuries.

The Liquidity factor level is an updated version of the on-the-run/off-the-run Treasury spread used in Ang, Goetzmann and Schaefer (2009).²² Newly issued, or on-the-run Treasuries, are more liquid than seasoned, or off-the-run Treasuries. The Liquidity factor level reflects the spread between the relatively illiquid off-the-run bonds and the relatively liquid on-the-run bonds. We graph it in Panel B of Figure 9, which shows a large spike during the financial crisis. This corresponds to the pronounced withdrawal of liquidity in markets around the world at this time. Since 2009, liquidity has recovered and U.S. Treasury markets are even more liquid, as measured by the on-the-run/off-the-run spread, than during the pre-crisis period.

Since the Liquidity factor level is not a return series, unlike the other factors, it requires special attention in the context of our return attribution. The series appears stationary and could serve as an explanatory variable as is. However, in the spirit of Merton's (1973) ICAPM we first-difference it, thereby focusing on changes in liquidity, as opposed to levels. We also flip the sign with the intuition that portfolios which experience negative returns when liquidity deteriorates

²¹ Note that the credit risk factors may also contain a term premium to the extent that the maturity or duration of the long bond basket is not the same as those of the short bond basket.

²² This is provided by the Federal Reserve Board and is originally constructed by Gürkaynak, Sack and Wright (2007). There are other illiquidity factors in the literature, including those constructed from equities, like Pástor and Stambaugh (2003), and from bonds, like Bao, Pan and Wang (2011). The former does not perfectly isolate credit and other risks and the latter requires transaction data for bonds that is not available at the start of the sample period. There is, however, large commonality across different liquidity measures as documented by Chordia, Roll and Subrahmanyam (2000) and others. Ang, Goetzmann and Schaefer (2009) compare the on-the-run/off-the-run spread with other measures of illiquidity.

should earn a positive premium over the long run. This treatment is in line with Ang, Goetzmann and Schaefer (2009).

II.D.2. Equity Factors

The equity factors ValGrth, SmLg, and Mom plotted in Panel C of Figure 9 were used by Ang, Goetzmann and Schaefer (2009) and reflect the most common cross-sectional anomalies studied in finance since the Fama and French (1993) model. Fama and French constructed their model to capture value-growth and small-large effects, to which is commonly added a momentum factor (see Jegadeesh and Titman, 1993; Carhart, 1997). All three factors exhibited negative returns over 2008-2009, with the losses associated with momentum being particularly severe. Momentum shorts stocks with past low returns; many financial companies experienced negative returns over 2008 but performed strongly in 2009 as policymakers undertook efforts to inject liquidity and rescue certain financial firms. Momentum strategies experienced losses in these short positions when financial firms rebounded.

The SmLg factor deserves further comment. Banz (1981) was the first to document that small stocks delivered higher risk-adjusted returns than large stocks. Since the mid-1980s, however, the magnitude of the size effect has diminished; small stocks do have higher average returns than larger stocks, but small stocks do not seem to have higher risk-adjusted returns than large stocks after controlling for their market risk exposure (betas). Many small stocks have betas greater than one. Several commentators, including Dimson, Marsh and Staunton (2011) and Fama and French (2012) note that there is now no significant difference between the risk-adjusted returns of small and large companies among international stocks. Panel C of Figure 9 shows that the SmLg factor has earned positive average returns over the sample, but it would diminish significantly after controlling for market risk. We include a SmLg factor because leveraging market beta risk entails an active decision. The Executive Board of Norges Bank has issued guidelines restricting leverage, but not barring, the use of leverage in the Fund's management. NBIM has been given a mandate to outperform a benchmark and takes active bets relative to that benchmark. Taking on a beta greater than one with respect to that benchmark, which requires leverage, is an active decision.

LowVol reflects the phenomenon that low volatility stocks have higher returns, on average, than high volatility stocks. Since Ang et al. (2006, 2009), the literature has labeled this the "risk anomaly." This is a factor that was not considered in the Ang, Goetzmann and Schaefer (2009) report. Since then a sizeable academic literature has investigated this anomaly and found it to be pervasive: it is observed in foreign exchange markets, international equities, fixed income, option markets, and commodity markets (see Ang, 2014, for a summary). The asset management industry has introduced an array of risk anomaly products, including ETFs and institutional funds. This is a factor that has also been previously considered by the Ministry of

Finance (see MSCI's "Harvesting Risk Premia for Large Scale Portfolios" report). We construct LowVol by sorting U.S. stocks into quintiles based on their idiosyncratic volatility relative to the Fama and French (1993) model using daily returns over the past three months. LowVol represents the difference in returns between the first and fifth quintiles, and is rebalanced at the quarterly frequency. We scale the LowVol factor to have a volatility of 15% before it is converted to the Fund's global currency basket.²³

II.D.3: Implied Volatility Factor

Our final factor, SellVol, captures the returns from selling equity volatility protection in derivatives markets. Specifically it is the excess return on a variance swap which is the difference between an up-front rate tied closely to the difference between at-the-money implied volatility and ex-post realized volatility on the equity market. This factor was also included by Ang, Goetzmann and Schaefer (2009) but their series, which was sourced from Bloomberg, was discontinued soon after it experienced extreme losses in 2008. Since then several "advanced" volatility selling indices have been introduced by broker-dealers deploying surely ex-post discovered filters that avoid the 2008 drawdown. To conduct a "real-time" analysis as much as possible, we use the series from Ang, Goetzmann and Schaefer (2009) through 2008 and splice on a representative "advanced" return series. As can be seen clearly from the dashed-dotted green line in Panel C of Figure 9, short volatility strategies, which are captured by the SellVol factor, earn steady premiums during stable times. When volatility spikes, investors who are short volatility lose money. These losses represent gains to the investors on the other side of the trades who purchase volatility protection.

II.D.4: Correlations

Table 4 presents correlation matrices of the bond benchmark returns, stock benchmark returns, and the systematic risk factors for the full sample (Panel A), the pre-crisis subsample (Panel B), and the post-crisis subsample (Panel C). First we notice the changing correlation between the bond and stock benchmark returns, ranging from -0.38 pre finance crisis to -0.05 afterward.

Second, there are some noteworthy correlations between factors. Among the fixed income factors, we observe the credit factors are highly correlated with each other and negatively

²³ Another "low risk" anomaly is the low returns to stocks with past high betas. This manifests in a Security Market Line that is flatter than what is predicted by the CAPM, and was been documented even in the early literature testing the CAPM like Black, Jensen and Scholes (1972). Frazzini and Pedersen (2011) formulate a "betting against beta" factor that goes long stocks with low past betas and shorts enough stocks with high past betas to obtain a net beta neutral portfolio. Constructing this factor requires borrowing at the risk-free rate, which is usually taken to be the yield on U.S. Treasury bills. We choose not to use this factor because U.S. T-bills are not the risk-free asset for the Fund. Moreover, the Fund is prohibited from employing such explicit form of leverage. If instead we consider a long-short portfolio that is dollar neutral, and hence does not require explicit leverage, the spread in returns between stocks with past low and high betas is very small compared to portfolios sorted on past volatility.

correlated with the Term, consistent with pattern in the cumulative returns in Panel A of Figure 9. Among the equity factors, ValGrth and Mom are negatively correlated as value firms tend to be firms that have underperformed their peers for some time. More surprisingly, over the recent five years, Mom is almost as negatively correlated with SmLg, whereas in the pre-crisis sample this correlation was positive. Finally, the LowVol factor exhibits strong correlations with other equity factors, particularly in the post-crisis sample: -0.47 with ValGrth, -0.62 with SmLg, 0.71 with Mom, and -0.33 with SellVol.

Third, the off-diagonal blocks show the correlations between the benchmark returns and other risk factors are important. The bond benchmark is highly positively correlated with Term, fairly strongly negatively correlated with the credit risk factors, and relatively uncorrelated with all other factors. The average correlation of the bond benchmark returns with the factors is 0.01 for the whole sample and -0.02 for the most recent five years. Similarly, the stock benchmark is strongly positively correlated with the credit risk factors and SellVol, strongly negatively correlated with Term, Mom, and LowVol, and relatively uncorrelated with the remaining factors. The average correlation of the stock benchmark returns with the factors is 0.11 over the whole sample and 0.16 since the financial crisis. These overall low correlations between the stock and bond risk factors and the other systematic risk factors illustrate that systematic factor exposures should not be thought of as an alternative to traditional bond and stock investments but rather as a diversifying complement. Said differently, factors provide significant incremental diversification benefits to traditional passive market capitalization-based indexes (see Section I).

In Table 5 we report monthly autocorrelations of the bond benchmark returns, stock benchmark returns, and other risk factors. The purpose of this table is to shed more light on the high autocorrelations of the active returns mentioned in Section II.C. There are a number of risk factors with high autocorrelations during and post the financial crisis, including Term, CreditHY, Liquidity, Mom, LowVol, and SellVol (which had extremely high persistence during the crisis). To the extent that the active returns of the Fund are partially attributable to systematic exposures to these factors, the persistence in the active returns documented in Panel C of Tables 1 and 2 can be explained by the persistence in the systematic factor returns.

II.E: Active Factor Exposures

The Fund's active returns have economically large and statistically significant exposures to factor risks.

Having identified a set of systematic risk factors that the academic literature broadly agrees are robust sources of diversifying returns, we now evaluate to what extent the Fund's active returns are exposed to these factors. Our goal is to measure the Fund's exposures with the belief that the more the Fund's active returns are associated with established factors, the more likely it is that active management will continue to add value over the medium to long term through the systematic risk premium channel if the factor exposures are maintained over the long run (see Section I.D). This objective is different from the performance evaluation literature which generally tries to measure active returns net of certain factor exposures. Factor risk exposure can occur as a result of deliberate decisions to engineer direct exposure, such as taken through the ORP, which we discuss below, or indirectly as a result of an investment process which gives rise to those exposures. In either case, it is relevant to gauge the ex-post factor risk generated by active management.

It is common in the academic literature to define a risk-adjusted or residual return, which is computed by regressing a fund's return onto its benchmark and other factors. We believe this approach is inappropriate to assess the Fund's active performance. First, and most importantly, the Fund manager is asked to beat the raw benchmark and thus generates returns relative to that benchmark. It should be assessed on ex-post active returns. The Fund manager was not given a benchmark incorporating other systematic factors, nor was it given a benchmark that levered up or down the exposure to the equity and bond indexes—an exposure that is only measured ex post and not stated ex ante as part of the mandate. NBIM certainly makes investment management decisions which gains it exposure to systematic factors not in its basic benchmark. We therefore investigate whether the active returns are correlated with various systematic factors, and if the movements in the active returns can be to any extent be explained by systematic factors. We do not use the factors, or a de-levered or levered benchmark, to measure the value-added from active management.²⁴

Following Ang, Goetzmann and Schaefer (2009) we analyze the active returns of the Fund in two ways. First, we present *partial* correlations of the active returns with the systematic factors in Panel D in each of Tables 1 and 2. Partial correlations control for the effect of other factors

²⁴ For completeness, we examine residual returns which do take into account a leveraged and potentially time-varying exposure to the benchmark in Appendix A.

when computing the marginal relationship between the active returns and one factor.²⁵ We compute partial correlations of active returns with dynamic factors. Both are long-short positions: the active return measures the deviation from the benchmark taken by NBIM, and the factors represent differences in returns of one group of stocks with high or low levels of one characteristic, like credit ratings, the value orientation of a stock, or past return or volatility characteristics, etc., compared to another group of stocks with the opposite characteristics (the exception is the Liquidity factor, which is not a return). The factors implicitly measure the returns resulting from taking long-short positions relative to the typical security. Thus we compare long-short active returns with long-short dynamic factors. This analysis would not be appropriate if the active return had large exposures to the benchmark itself, but Appendix A shows that the loading of active returns on the benchmark is very small.²⁶

Second, we regress the active returns jointly on the factors. In Figures 4, 5, 7, and 8, we plot the fitted values of the regression along with the realized active returns of the Fund (Figures 4 and 5) as well as its fixed income and equities components (Figures 7 and 8). We report the regression R^2 s in the figures' legends. For the overall Fund returns, we include all 10 factors. For the fixed income returns we only include the factors Term, CreditAa, CreditBaa, CreditHY, and Liquidity, whereas for the equity returns we only include the factors ValGrth, SmLg, Mom, LowVol, and SellVol.

We prefer looking at partial correlations of active returns rather than reporting coefficients of a multivariate regression. The factors are long-short returns (with the exception of Liquidity) and so can be arbitrarily scaled to any volatility. This free scaling affects estimates of factor loadings, but the correlations are scale-free. Moreover, we focus on *partial* correlations, as opposed to simple bivariate correlations, because they allow us to control for the effect of other variables. That said, we also acknowledge that partial correlations place a greater demand on the data, of which we do not have much to work with, especially with high (absolute) correlations between factors and time-varying relationships (Table 4 shows that several of the factors have correlations of +/-0.5 or stronger). This means that the point estimates need to be interpreted cautiously and that the R^2 s potentially represent estimates of upper bounds for the amount of factor exposure. As a robustness check, we also report simple correlations, which are less problematic computationally, but are in turn more difficult to interpret, in Appendix A.

²⁵ Formally the partial correlations are defined as follows. If the set of factors is partitioned into $F = \begin{bmatrix} F_1 & F_2 \end{bmatrix}$

with corresponding covariance matrix $\Sigma_F = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}$, then the partial covariance of F_1 controlling for F_2 is

given by $\Sigma_{F_1|F_2} = \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}$.

²⁶ This is not surprising because according to the Fund's mandate, leverage is used only "with the aim to implement the investment mandate in an effective manner."

II.E.1: Fund-Level Active Factor Exposures

Starting at the Fund level, we observe in Panel D of Table 1 that over the full sample the active returns of the Fund load significantly on CreditAa, Liquidity, SmLg, and SellVol. Active returns had a more modest negative loading on ValGrth. However, comparing the full-sample exposures to those measured over the pre- and post-crisis subsamples, it is clear that the full-sample partial correlations are heavily influenced by the crisis period. Looking at the post-crisis results, the active returns of the Fund loaded positively on Term, the credit factors, Liquidity, and Mom. The active returns loaded negatively on LowVol and SellVol, but the latter is statistically insignificant. Combined, these factors explain 60.5% of the active returns on the fund and even fit the changing magnitude of the active returns from 2009 to the more recent years remarkably well. This is shown in Figures 4 and 5.

It is useful to interpret the Fund's 60% factor R^2 in the context of other investment vehicles. Fama and French (2010) compute equal- and value-weighted mutual fund returns and regress them onto market, size, value, and momentum factors. They find regression R^2 s of 98-99%, and find that value-weighted mutual fund returns have significant tilts toward small, growth, and momentum stocks. These mutual fund R^2 s are extremely high because of the large exposure of their returns to aggregate market movements. Perhaps a more relevant comparison is the universe of hedge funds, many of which take long-short positions to minimize exposure to the market. A large literature has also documented that factors play an important role in explaining hedge fund returns. Factors explain around half, and often more, of the variance of hedge fund returns, which is similar to the Fund's factor R^2 . Fung et al. (2008) construct a series of equally-weighted hedge fund returns and regress it onto various systematic factors. They estimate an R^2 of 74%. In Jagannathan, Malakov and Novikov's (2010) performance regressions with hedge fund returns onto systematic factors, the R^2 s range from 40-50%. In Fung and Hsieh (2004), factors explain 50-80% of the variance of hedge fund returns.

That approximately two-thirds of the Fund's variance of active returns can be attributed to factor risks indicates that a large amount of the active management is related, either directly or indirectly, to risk that is *systematic*. As Section I discusses, systematic risk premiums arise because some investors seek to avoid holding such risk and that risk cannot be diversified away. Investors, like the Fund, who have larger exposures to these risks endure losses that are more severe than the typical market participant. To compensate for these risks, these investors receive factor risk premiums over the long run. Since these factor risks are relatively well understood and have exhibited returns over long histories, they can be more stable sources of active management than security selection. The factor risk premiums also operate over low frequencies and have long verification horizons—they require a long horizon to ride through

short-term losses—exploiting them is optimally done using the comparative advantages of the Fund.

It is interesting to contrast the loadings of the Fund's active returns pre and post financial crisis. The Fund did not exhibit any meaningful fixed income factor exposure pre crisis. This is not to say that the factor exposures were not being taken, but identifying factor exposures in periods which exhibit low volatility and no large factor losses or gains is econometrically challenging. The flipside is that it is easier to estimate significant factor exposures, if they are present, during the financial crisis—as is the case for the Fund. Pre crisis, the Fund's active returns also had negative value exposure, substantial small firm exposure, and a positive volatility selling exposure which have all changed significantly since then, in part through conscious choices by NBIM and the external asset managers it employs.

II.E.2: Fixed Income and Equity Active Factor Exposures

In Panel D of Table 2, we report partial correlations for active returns in fixed income and equities. The results are relatively clear cut for fixed income, with large and significant loadings on Term, CreditAa, and CreditHY post crisis. Combined, the fixed income factors explain 41.7% of the active fixed income returns (see Panel A in each of Figures 7 and 8). The results are more noisy for the active returns on equities. Matching the Fund-level findings in Table 1, Panel D, the active equity returns load positively on the size factor and negatively on the low volatility factor. The exposure of the active equity returns to ValGrth is negative, and statistically significant, over all three sample periods, while it is positive, but insignificant, at the Fund level. Further analysis reveals that, at least for the last few years, this apparent inconsistency is caused by extreme positive returns on the value factor being associated with positive and large active returns in both fixed income and equities. This causes the Fund-level returns to load positively while equities load negatively on ValGrth. Specifically in the crisis subsample, value stocks experienced extremely large returns in April and August of 2009 (returns of 4.8% and 3.8%, respectively). In those months, the active returns in fixed income were 79 and 86 basis points, and equities returned an excess of 26 and 12 basis points. With these two observations removed from the post crisis sample, the partial correlation with ValGrth at the Fund level is also negative.

Comparing the asset class level results across sample periods, the changes in the active exposures we commented on at the Fund level are confirmed at the asset class level. It is interesting to note that neither the size nor the value exposures change meaningfully despite NBIM's active effort to tilt the portfolio into the direction of these two factors through the ORP (which we provide more detail below). We suspect that the sample is simply too short to measure this tilt without observing actual security holdings. Alternatively, the actual loadings

that NBIM has put on these factors in the ORP are small and undone by security selection. Greater transparency would help differentiate between these hypotheses (see Section IV).

II.F: Operational Reference Portfolio

The Operational Reference Portfolio offers the Fund improved diversification and allows it to efficiently harvest factor risk premiums.

Our assessment of NBIM's ORP is drawn primarily from conversations with personnel at NBIM. While NBIM provided some summaries and reports some limited historical returns, which have informed our analysis, our discussion is primarily qualitative. The ORP is a work in progress and has undergone, continues to undergo, dramatic expansion in its scope. In addition, there are several elements in the ORP which are hard to separate quantitatively. It is our opinion that our conversations with NBIM are more informative than analyzing an extremely short, and noisy, non-stationary history of ORP returns.

The ORP was introduced in 2011 and is an internal benchmark constructed by NBIM. It is a tailored benchmark with non-market capitalization weights, and takes as its starting point the benchmarks set by the Ministry of Finance. The ORP takes advantage of the characteristics of the Fund: a large investor with no fixed liabilities, no immediate liquidity requirements, and unconstrained by domestic currency considerations (see also Section III.A).

NBIM uses the ORP for three purposes:

1. Diversify more widely than standard benchmarks
2. Take on systematic factor risk exposure
3. Implement smart rebalancing

We discuss each of these in turn, and explain how each component requires a different verification horizon for risk management. The ORP is internally owned by NBIM and takes on only a limited set of risk factors that have been considered by the Ministry of Finance. A part of NBIM's mandate concerning fiscal strength in the sovereign bond portfolio, however, is not quantitatively expressed in an off-the-shelf benchmark. This mandate is incorporated by NBIM in the ORP. In its current incarnation, the ORP does not separately report the returns from each component—a point we return to in Section IV.

II.F.1: Improved Diversification

As Section I discusses, the true market portfolio—which is the most diversified portfolio under the CAPM—is unobserved. Benchmark providers have constructed proxies to the market in

different asset classes, but these require somewhat arbitrary choices of which securities to include and how they are weighted.

Index providers, for example, do not use the full market capitalization (number of shares times share price) in computing their index weights for all securities because some shares are closely held and not liquid. These shares may be held by founders, family firms, or are subject to lockup or other illiquidity constraints. Index providers compute “free float adjustments” to take account of these illiquidity effects, which are necessary because most investors value liquidity in trading these indexes. For some investors like ETF providers, these liquidity considerations are paramount. The Fund can move closer to full market capitalization weights than public market indexes. By overweighting more illiquid stocks, it collects a liquidity premium. Another consideration is that only certain countries can be included in an index, usually those above a certain size. These markets may represent attractive investment opportunities, but are not included in some off-the-shelf indexes.

In fixed income indexes, duration (or interest rate risk exposure) is a consequence of the bonds included in the index, rather than a deliberate choice by the investor or index provider. Barclays adjusts its fixed income indexes for bond purchases by the Federal Reserve (quantitative easing), but no such adjustments can be done by central banks in other countries as similar data are not published. Fixed income indexes often arbitrarily have minimum maturity cutoffs, so that investors with low tracking error constraints are forced to sell these bonds before maturity.²⁷ These adjustments in fixed income indexes are analogous to the free float adjustments in equity indexes.

NBIM uses the ORP to alleviate some of these disadvantages of public indexes and thus creates a superior portfolio. In the context of Section I, it goes some way to restoring the original intention of diversification by trying to find a better market portfolio; the ORP enables the Fund to obtain more diversification than can be obtained by using a public index. While public indexes fill an important requirement for transparency and simplicity, they do not deliver the best long-term diversification achievable by the Fund.

II.F.2: Factor Risks

The adoption of factor risks in the ORP follows the recommendation given in the previous Ang, Goetzmann and Schaefer (2009) report that the Fund consider investment strategies focused on factor risk premiums (see Section I.F). In equities, NBIM states that it aims to harvest size and value factors. Our estimated factor exposures in Section II.E are consistent with the stated intention for size, reflected in the large and significant size factor loading, but not for value, which carries a negative factor loading. In fixed income, NBIM adjusts the benchmarks to a

²⁷ See DeCosta, Leng and Noronha (2013).

more uniform duration and takes into account the fiscal strength mandate. The former makes NBIM's interest rate risk more explicit and allows it to tailor its active fixed income risk. The latter requires moving sovereign bond holdings away from GDP weights and weights countries based on fundamental and credit risks. Consistent with this aim, Section II.E shows that the fixed income portion of the Fund has a highly statistically significant loading to the Term, or duration, factor.

While these factor risk exposures are well known, there is no unique way to construct optimal portfolios to harvest these risk premiums. Several index providers offer indexes with factor tilts, and the Ministry of Finance has commissioned some studies investigating tailored factor benchmarks. In our opinion, the optimal implementation of factor risk portfolios does not lend itself to transparent, published benchmarks for the Fund, although they are useful for other investors. An optimal factor portfolio for a very large investor must have minimal turnover and small deviations from index weights. Proprietary methods of minimizing transaction costs and rebalancing portfolios are prerequisites for successfully harvesting factor risk premiums in scale.

We believe that having dynamic factor weights like size and value set in a benchmark index given by the Ministry of Finance offsets some of the benefits of these factor positions because of the large liquidity costs associated with following these indexes. It is important to note that the market-weighted indexes used by the Ministry of Finance are good starting points because they can be replicated at almost zero cost even though they do not take dynamic factor risks. Factor risk strategies, in contrast to passive market-weighted indexes, cannot be implemented at effectively zero cost. (Section I.C provides further remarks on characteristics of optimal benchmarks.)

The current mandate of requiring fiscal strength considerations in the bond portfolio has a disadvantage in that it is a qualitative mandate without a quantitative benchmark. That is, the fixed income index is based on GDP weights absent fiscal considerations, but taking these into account entails deviating from benchmark. In the worst case, fiscal strength considerations can use some of the Fund's risk budget without a commensurate increase in expected return. Indeed, since there are generally higher returns for taking on credit risk, closely following the fiscal strength criterion must reduce the Fund's return in fixed income investments. This is an important example of a case where the ORP allows NBIM to implement a mandate, and manage its risk, which is not made quantitatively explicit by the Ministry of Finance.

II.F.3: Efficient Index Tracking

A detailed public rebalancing rule is detrimental for the Fund; its size would encourage front-running of its trades and incur substantial transactions costs. The ORP has procedures which

minimize transaction costs taking into account the Fund's size, anticipated inflows, and spreads the sales and purchases of securities over several periods. This part of the ORP operationally implements the benchmark exposures of the Fund in a cost-efficient manner.

Rebalancing allows a long-horizon investor to harvest the rebalancing premium, as Section I shows, and implementing rebalancing in a smart way prevents all the benefits of rebalancing from being frittered away by unnecessary transactions costs. Smart trading is essential for all active strategies away from the rebalanced, passive benchmark, including the broader diversification and factor risk exposure components of the ORP mentioned earlier. While NBIM uses the ORP to internally mark its replication of the benchmark index, all active strategies—including the first two components of the ORP—also entail transaction costs.

II.F.4 Advantage of Appropriate Verification Horizons

Each of the three stated components of the ORP has a different verification horizon. The effects of improved diversification involve measuring means, and since financial returns are volatile they can take up to a decade or more to manifest. Many factors—including size and value—can take from five to ten years to be reliably measured. During the late 1990s, for example, the internet boom caused value strategies to underperform for several years. Both equity factors suffered drawdowns during 2008 and 2009. In contrast, smart rebalancing decisions can be measured over days to months.

Of the three components of the ORP, only the first—improved diversification—is always theoretically rewarded by superior risk-adjusted long-run returns. The second component of factor risk premiums should be rewarded in the long run for factors based on compelling economic theory. There should be positive risk premiums for taking on exposures to size and value factors (and the other factors covered in Section II.D), but it may be the opposite case for the fiscal strength mandate as countries with lower debt-to-GDP ratios and higher growth should have lower expected returns.²⁸ For the Fund, factors should be chosen on the basis of persuasive academic justification and strong empirical evidence. The third component, smart rebalancing, allows NBIM to track an index. Absent mispricing opportunities and not taking into account stock lending and other balance sheet activities, there is no possibility to implement an index return better than its theoretical construction, which does not take into account transaction costs.

²⁸ Alternatively, tilting the fixed income portfolio away from market weights to take account fiscal strength considerations should not be considered a compensated risk factor.

II.F.5 Governance

The ORP sits entirely within NBIM and is thus an internal decision of the fund manager. It is overseen by the Investment Risk Committee, which takes strategic investment decisions over horizons of five to ten years. The Investment Risk Committee's recommendations are advisory to the CEO. There is no direct communication to the public on the nuts-and-bolts implementation of the ORP, which in our opinion is appropriate because it involves proprietary materials and trades, which if publicly disclosed, can lead to adverse market impact.

The first and third components of the ORP are solely the responsibility of NBIM: selecting securities not in the index universe to improve diversification and implementing the rebalancing regime. In our opinion, certain style directives, like fiscal strength considerations and responsible investment practice, are not solely a decision by NBIM, and thus deviations from benchmark may not always represent decisions to seek the best possible return for a given amount of risk (see also Section III.A).

The choices of the factor risks of size and value in equities, and duration in fixed income, are taken solely by NBIM. The Ministry of Finance has considered other factors which are currently not included in the ORP, including liquidity, low volatility, and strategic timing of factor risks. If the fund manager is not considering factor risks premiums which the ultimate owners of the Fund, represented by the Ministry of Finance and Parliament, desire to be included in the Fund's investment strategy, then an agency problem will arise. NBIM has come to its decisions on taking size and value factors based on its own research and investing experience. Factor decisions made by the Ministry of Finance would entail using publicly available indexes which may not represent the best risk-return trade-offs obtainable by the Fund.

The fiscal strength exposure in fixed income is a mandate given by the Ministry of Finance that is not explicitly measured by the benchmark, and is partly implemented in the ORP. To the extent that these decisions are not made seeking the best possible risk-adjusted returns, this represents a potential opportunity cost of using deviations from benchmark that could be better employed by other profitable investment opportunities.

II.F.6 Summary of the ORP

The ORP is an excellent development of NBIM that capitalizes on the Fund's comparative advantages. It represents a way to diversify beyond public index benchmarks and efficiently harvest factor risk premiums. It comingles costs of tracking the Fund's benchmark; transaction costs are also associated with all active strategies of the Fund. We believe that risk-taking through the ORP should be separately considered relative to security selection (see Section I), as the sources of return NBIM seeks to harvest in the ORP are distinct. Our recommendations involving the ORP in Section IV build on this work achieved to date.

II.G: Benchmark Deviation Risk

Since the financial crisis, the Fund's realized risk of active management has been roughly half of its maximum limit.

The Ministry of Finance sets the Fund's active risk budget in terms of "tracking error," which is the annual deviation of the Fund's return in excess of its benchmark. The Fund has changed its risk limits over the sample. Prior to 2009, the Ministry set a benchmark deviation volatility of 1.5% per annum. The volatility of active returns was temporarily higher than this limit during 2008, and, as we saw above, active returns were also negative over this period. In the spring of 2010, the tracking error constraint was changed from a hard maximum of 1.5% to a tracking error limit of 1.0%. In the new regime, NBIM is understood to take on active strategies which should not exceed a benchmark deviation volatility of 1.0% during normal times. The method of computation of risk was also changed from using daily return data where newer data, where given a greater weight than older data, to using equally-weighted weekly return data over the previous three years. The new lower risk limit is not considered inviolable, so realized benchmark deviations may occasionally exceed this expected target during extremely volatile periods.

II.G.1: Realized Benchmark Deviation Volatility

It is difficult for us to quantitatively assess the Fund's tracking error given (a) we only have monthly data and the formal calculations of tracking error are now done at a higher frequency, (b) the extreme active return experience in 2008 and 2009, particularly in fixed income, and (c) there have been changes in the tracking error level and its computation. Our analysis is also ex post while the Fund's mandate risk restriction is ex ante. We cannot evaluate ex-ante risk management or budgeting as we were not provided with any ex-ante risk estimates. We were also not given high frequency data required for computing such estimates. With these caveats in mind, we plot in Figure 10 the rolling annualized standard deviation of active returns for the Fund (solid blue line) as well as its fixed income and equities components (dashed green and dashed-dotted purple lines, respectively).

We estimate the standard deviation with three years of monthly-frequency returns and shift it forward by 18 months to better align the estimates with time. That is, for January 2009, for example, our estimate of the tracking error uses returns from July 2007 through June 2010. In this way, the estimate in January 2009 more accurately reflects the realized benchmark deviation volatility surrounding that date as opposed to the previous three years. Even with this adjustment, it is important to note that the tracking errors in Figure 10 do not completely represent risk. High tracking errors observed during the financial crisis reflect a relatively high risk of obtaining returns different from the Fund's benchmark. But, given the tremendous

misvaluation opportunities available in markets, the true risk for a long-term investor is not deviating from the benchmark to take advantage of them at that time.

The behavior of the benchmark deviations is consistent with our discussion of the active returns in Sections II.C and II.E. Benchmark deviation volatility at the Fund level was 50 basis points or lower until 2007, then increased to above 1.5% by 2009, and has since then again dropped to below 50 basis points, which is well below the current limit of 1%. Moreover, comparing the volatilities of benchmark deviations for fixed income to that for equities, it is clear that the elevated Fund-level benchmark deviation volatility through the crisis was caused primarily by increases in risk emanating from the fixed income portion of the Fund.

Despite the relatively high monthly autocorrelation of the Fund's active returns, we do not adjust our tracking error estimates. We acknowledge that such adjustments are possible, but given only 36 monthly observations, the autocorrelation adjustment would introduce considerable noise. Moreover, much of the positive autocorrelation in active returns was caused by the 2008-2009 period, and we expect the autocorrelation of active returns going forward to be more in line with the pre-crisis experience. That said, we can ballpark the magnitude of the downward bias in our calculations due to ignoring the autocorrelation of active returns. If the data follow a simple first-order autoregressive structure (AR(1)) with a monthly autocorrelation of 0.6, in line with the results for the post-crisis sample period, a measured volatility of benchmark deviations of 50 basis points implies a true volatility of benchmark deviations of 52.1 basis points taking into account the autocorrelation structure, or a 4.2% relative downward bias. If, however, the autocorrelation is only 0.2, which more consistent with the pre-crisis experience and in our view also more realistic going forward, the implied true benchmark deviation volatility is 50.3 basis points, or a relative downward bias of only 0.7%. We conclude that the autocorrelation adjustment, while theoretically justified, would not meaningfully change our analysis, and only increase the noise in our estimates due to the short estimation period.

II.G.2: Sources of Benchmark Deviation Risk

To put the current 50 basis point benchmark deviation volatility into perspective, recall from Section I that there are several ways a fund manager can deviate from a market benchmark to generate value. In particular, under the newly implemented ORP framework, the active returns we measure can be decomposed into the difference between the ORP and the benchmark provided by the Ministry of Finance plus the difference between the actual returns on the Fund

and the ORP. That is,

$$\text{Active return} = \underbrace{r - r_{bmk}}_{\text{Benchmark deviation}} = \underbrace{r_{ORP} - r_{bmk}}_{\text{Factor tilts, improved diversification}} + \underbrace{r - r_{ORP}}_{\text{Security selection}}, \quad (3)$$

where r is the Fund return, r_{bmk} is the benchmark, and r_{ORP} is the return of the ORP portfolio.

The first component of this decomposition, in turn, could in principle be further decomposed into the various stated objectives of the ORP, including better diversification and risk factor tilts. The difference between the Fund returns and the ORP is the risk budget for security selection. Without more detailed data on the Fund's holdings and, more importantly, the composition of the ORP, it is impossible for us to quantitatively decompose the active return in this fashion. Moreover, given the ongoing development of the ORP framework, such an historical decomposition is not that useful for our forward-looking commentary. Instead, we consider a conceptual experiment of how superior diversification alone can induce a significant amount of deviation from benchmark. As we demonstrate, the benchmark deviation arising solely from increased diversification can be as large as 50 basis points for an equity portfolio, suggesting that the current active risk budget as well as the amount of active risk currently taken by NBIM may be insufficient to accomplish the diversification objective of the ORP—which is also the first goal of a long-term investor (see Section I.A).

NBIM can move closer to the true world market portfolio and thereby increase diversification in two ways. First, it can add certain emerging market or frontier countries that are excluded from the benchmark, but NBIM considers them investable. Second, noting that different index providers make different free float adjustments, NBIM has developed its own methodology to compute the appropriate weighting of the index constituents. We attempt to approximate the effects of these two considerations by comparing the benchmark deviations between different publically available indices. We collect monthly returns for our sample period on three MSCI indices: Developed, Emerging, and Frontier. We then compute the benchmark deviations of Developed vs. (Developed + Emerging + Frontier) (i.e., the World) and of (Developed + Emerging) vs. (Developed + Emerging + Frontier), where, naturally, we capitalization weight the index series. We plot the volatility of the benchmark deviations of (Developed + Emerging) vs. (Developed + Emerging + Frontier) in Figure 11.

We estimate that the annual volatility of deviations from excluding both Emerging and Frontier markets from a developed-markets only benchmark is large at about 2% per year. The Fund's equity benchmark includes most emerging markets, so this calculation is too extreme. More relevant for the Fund, the inclusion or exclusion of Frontier markets creates a 28 basis points volatility of benchmark deviations over the whole sample period, or a 24 basis points

benchmark deviation volatility since 2009. To gauge the magnitude of benchmark deviations that different index weighting methodologies can create, we compute the standard deviation of the return difference between the MSCI and FTSE All World indices. For the entire sample, the volatility of benchmark deviations between the two indices is almost 1%. For the most recent five-year period, it is 47 basis points. These results indicate that seeking greater diversification could easily consume a significant amount of deviations from benchmark. The results also show that the selection of the passive market benchmark is not inconsequential.

As Section I discusses, holding a diversified portfolio increases an investor's risk-return trade-off. The ORP accomplishes this and moves the Fund closer to the unobserved world market portfolio. In the equities portion of the Fund, the ORP includes certain markets that are not included in the Ministry of Finance's benchmark and NBIM also computes its own float-adjusted index weights which allow it to give greater weight to issuers with more illiquid securities. Based on the two admittedly crude calculations, each of these adjustments taken alone could induce 25 to 50 basis points of tracking error, leaving little room for factor tilts and ultimately security selection under the Fund's current tracking error limit. We take up this issue again in Section IV.

II.G.3: Tail Risk

The volatility of benchmark deviations is a symmetric measure of active risk that treats positive and negative deviations from the benchmark as equally important. From the stakeholders' perspective, it is reasonable to think that large losses of active management are more important to avoid compared to capitalizing on equally large gains. Moreover, volatility is a risk measure which focuses largely on the central mass of the distribution and therefore has relatively little to say about tail risks—unless it is combined with a distributional assumption such as normality. In this section, we examine to what extent the distribution of benchmark deviations is skewed or fat tailed. If there are significant differences from normality, benchmark deviations are better constrained by also taking into account measures that incorporate higher moments of the whole distribution. Some commonly used left tail risk measures in risk management fall under the rubric of “Value-at-Risk” (or “Expected Shortfall” and “Conditional Value-at-Risk”) risk methodologies.

Panel A of Figure 12 shows in the blue line the empirical distribution of active returns constructed using a standard non-parametric density estimator applied to the full sample of active returns.²⁹ In orange we overlay a normal distribution with matching mean and volatility. It is obvious that the distribution of benchmark deviations is non-normal. First, and most

²⁹ We use a Gaussian kernel smoother with a bandwidth that is 25% larger than the asymptotically optimal choice under normality for our sample size to accommodate the significantly fatter tails.

strikingly, the empirical distribution is fat tailed (resulting in more central mass and longer tails). Second, somewhat more subtle to discern in the plot, it is negatively skewed. This is more clearly seen in Panel B of Figure 12, which shows the corresponding cumulative density functions (or CDFs). Compared to a normal CDF, the empirical CDF departs substantially more from normality in the left tail than the right.

The empirical distribution or CDF can be used to calibrate the probability of a loss exceeding a certain threshold given the amount of tracking error allowed. However, such calculations can be misleading under non-normality. For example, in Panel B the 90% relative Value-at-Risk of the Fund, defined as the loss in terms of percent return that should not be exceeded more than 10% of the time, is -25 basis points under normality and only -16 basis points with the empirical distribution. Does this mean the fund is actually less risky than it would be under normality? Unfortunately not—due to the fat tails of the empirical distribution, the 10% most extreme losses are likely to be much larger in magnitude compared to a normal distribution. This is a well-known shortcoming of standard Value-at-Risk analysis.

Panel C of Figure 12 plots the corresponding Expected Shortfall which is explicitly designed to address this shortcoming. The Expected Shortfall is defined as the expected return *conditional* on the return falling beyond a certain critical threshold in the tail of the distribution. For example, the 10% Expected Shortfall, corresponding to the 90% Value-at-Risk calculations above, is the expected return for the 10% worst outcomes. Under normality this expected loss is -35 basis points but under the empirical distribution it is -38 basis points. The differences between the two distributions become even larger further in the tails.

To illustrate how sensitive the tail risk calculations are to sample selection, we show in Figure 13 the same three plots except we exclude the financial crises along with the unusual recovery in 2009. The empirical distribution during “normal times” is still fat tailed but much less asymmetric. Again, this analysis is not meant as a definitive answer to how large the tail risk of the Fund is—we do not believe we have enough and the right kind of data to answer this question—but rather as an illustration of how to constrain benchmark deviations beyond simple “tracking error” limits.

II.G.4: Summary of Benchmark Deviation Risk

Overall, our assessment is that the amount of active risk taken by the Fund has been significantly reduced since 2008, and NBIM is more conscious of how that risk budget is broken down into systematic factor strategies and security selection that goes beyond factor exposures. We build on these positive developments in Section IV.

II.H: Summary

Our results on the Fund's active management since 2009 can be summarized as follows:

1. The Fund has become more cautious

In its investment strategies, NBIM has reduced its use of derivatives, does not employ leverage to increase the Fund's exposure to risky assets, and has reduced its reliance on external managers. The risk taking as measured by the volatility of benchmark deviations, or "tracking error," was formally reduced by the Ministry of Finance in 2009 to 1.0% per annum, and NBIM has taken approximately half of this limit.

2. Only a very small amount of the Fund's overall return comes from active management

Benchmark returns account for 99.3% of the Fund's return over the full sample, January 1998 to June 2013, and even more of the Fund's return, at 99.6%, post the financial crisis. Benchmark returns also dominate in the fixed income and equity return components of the Fund.

3. Active management has added value

The Fund returns in excess of benchmark are significantly positive over the full sample, and over the pre- and post-financial crisis periods. Since January 2009, active management has contributed 0.10% per month.

4. Active management is exposed to factor risks

Over 60% of the Fund's active returns can be explained by exposure to systematic risk factors. The magnitude of the factor exposures have changed over time, and fixed income and equity active returns are exposed to different factors.

5. The development of the ORP enables the Fund to obtain superior diversification and efficient harvest factor risk premiums

Introduced in 2011 and under internal control of NBIM, the ORP has been used to seek greater diversification opportunities than available in publicly available indexes and is a systematic approach to harvesting Liquidity, Term, ValGrth, and Size factors in scale. NBIM also uses the ORP to track the efficient implementation of index replication. The ORP is a significant step forward in the Fund's investment management framework.

[This page is intentionally left blank]

Section III: Active Mandate

In this section, we review how delegation can improve the Fund's expected return and risk relative to the current strategic benchmark. The Fund's comparative advantages come from both structural characteristics stemming from the nature of the Fund's underlying cash flows and developed advantages. The latter arise from the development of the Fund's investment management organization, governance structure and reporting, and management and personnel, rather than the attributes of the Fund's cash flows and lack of immediate liabilities. We couch our review of the investment frameworks of comparable funds in the context of investments in real estate, private equity, and infrastructure—areas where measurement, benchmarking, and governance problems are most challenging. The opportunity cost funding model of active management, as pioneered by Canada Pension Plan Investment Board (CPPIB), is a viable, and attractive, alternative to the current active mandate of the Fund.

III.A: Comparative Advantages and Disadvantages

The Fund's developed comparative advantages, which have risen endogenously in its governance framework, are as valuable as its structural advantages emanating from its cash flows and lack of immediate liabilities.

The Fund's *structural* advantages and disadvantages of its long-term horizon, its large size and economies of scale, and its ability to hold illiquid assets have been covered in publications issued by the Ministry of Finance, NBIM, and also studies like Ang, Goetzmann and Schaefer (2009). These advantages arise because of the underlying characteristics of the Fund itself. While these are by no means exclusive to the Fund, they do constitute a valuable basis for strategic differentiation from the average investor in the market. NBIM has created investment processes that incorporate and exploit some or all of the Fund's structural comparative advantages. Although we have not examined the real estate program due to its short history, we would assume that it also incorporates some of these natural comparative advantages.

In this section, we focus on *developed* advantages or capabilities, which are important complements to the Fund's structural advantages. The developed comparative advantages arise because of the organization, management, and culture of the Fund. Indeed, the ORP created by NBIM is an example of a program built around these advantages (see Section II.F). Developed comparative advantages include the investment strategy and professionalism of the Ministry of Finance and NBIM, and the knowledge and oversight of the ultimate owners of the Fund—the

Norwegian people. As shorthand, the difference between structural and developed advantages is similar to behavioral traits arising from “nature versus nurture.”

The developed comparative advantages we discuss are not exhaustive; we concentrate on four which we feel have particular implications for the investment process and upon which our recommendations in Section IV build: having a single, dedicated fund manager, transparency, a mandate based on investment criteria, and the degree of professionalism. Other developed characteristics include the public mission of the fund manager (agent), the well-informed stakeholders of the Fund including Parliament (principal), the close alignment between the interests of the principal and agent, and the importance of the Fund in the context of the national economy.

III.A.1: Dedicated Funds Manager

NBIM is based in Norway and has developed as a distinct investment institution within a central bank. The roles of the fund manager acting as agent, NBIM, and that of the Ministry of Finance, which acts as principal, are distinct. There is also clear accountability between the two levels. While there are some differences between the principal and agent, both have the same broad objective: to obtain the best risk-adjusted return of the Fund within a set of constraints.

The Ministry of Finance delegates responsibility for management to only one fund manager. In our view, NBIM shares all of the same goals as the Ministry of Finance—to manage the Fund prudently and efficiently in line with the preferences of the Norwegian people. NBIM can take a long-term horizon in its strategic planning knowing that it is not subject to potential short-term runs of capital. The Ministry of Finance also knows that NBIM is giving full attention to the management of the Fund, rather than having its focus split across multiple clients. There are other advantages of having a dedicated funds manager including economies of scale, not least with respect to back-office costs. NBIM has managed to attract and retain talented and experienced investment professionals (as we further comment below), and it can be difficult to build such an investment culture in public service institutions. Thus, many principal-agent misalignments which often dog traditional delegated investment setups are mitigated by the Fund’s governance structure.³⁰

A contrary perspective is that since NBIM does not have to work to attract funds in the market means, there is no “market correction” for the possibility of over-paying for funds management. Another potential disadvantage is that it may be more difficult for the Fund to change capabilities quickly as the fund manager must develop new expertise. The Fund has been run extremely efficiently; all stakeholders are conscious of delivering high-quality investment management services at lowest cost. The careful expansion of the Fund into real

³⁰ For a summary of some of these issues see van Binsbergen, Brandt and Koijen (2012).

estate reflects NBIM's desire to maintain its professional culture as it ventures into non-public markets (see also below). As an example of how Norway benefits from a single fund manager, consider the case of Sweden: in that country, state pension assets are managed by five different fund managers. A review in 2012 ("The Buffer Fund Inquiry") recommended consolidating the number of funds to take advantage of the economies of scale of in-house asset management and corresponding lower costs. There have also been tensions arising from this issue in the Dutch pension system. ABP, a large Dutch pension plan, originally had a dedicated funds manager, APG. In 2009, APG started managing pension funds for other clients. In the early years, this created "strategic tensions" between ABP, which desired a dedicated manager, and APG, which had to balance the needs of multiple clients although ABP was its largest, and most important, client.³¹

III.A.2: Transparency

Transparency helps ensure broad public support for the fund and enables strong legitimacy. Indeed, to our knowledge there is no comparable fund in the world that combines the extent of transparency, significance to the nation's economy, and degree of public interest and scrutiny of its activities. Transparency works hand-in-hand with the governance structure of the Fund, which separates the roles of the Ministry of Finance and NBIM.

One finding of the Ang, Goetzmann and Schaefer (2009) report was that had NBIM informed the public of the style of its active management, or the factor exposures exhibited in its deviations from benchmark, then the losses during 2008-2009 might not have been surprising when viewed in the context of the historical records associated with the types of active management styles pursued by NBIM at that time (see Section I.D). Stating clearly the expected verification horizons and potential drawdowns of different investment decisions helps all stakeholders to properly anticipate losses and manage risk.

While transparency is undoubtedly a hallmark and strength of the Fund, it also merits consideration in the context of potential large and very visible investments in operating companies. This issue is currently small, but as the Fund continues to expand in real estate, and potentially other non-public markets in the future, the disadvantages of transparency will grow. The Fund now has some, but limited, experience in this regard with some of its large real estate holdings, although there is a much higher potential for public controversy in the case of private equity and infrastructure. The risk of negative publicity of investment decisions, or headline risk, also increases with transparency.

³¹ See presentation by Bart Le Blanc, trustee of ABP, "Fiduciary Pension Fund Management: Lessons from the ABP/APG Experience (or How to Manage the Managers' Manager)," presentation to the June 2013 Rotman ICPM Discussion Forum.

The Fund's rebalancing framework is a superb example where transparency is achieved, but allows NBIM leeway to implement the required transactions without facing undue adverse selection. There are no ad-hoc decisions to rebalance, which are often taken by many funds at periodic intervals; the adoption of a rebalancing rule ensures that rebalancing is done automatically. Rebalancing allowed the Fund to maintain its exposure to equities in late 2008 and early 2009 when equity prices were low. The buying at low prices certainly turned out to an astute investment decision. It was the transparency of the rebalancing regime that permitted the Fund to buy assets which were severely depreciating in value at the time when many other funds failed to do so.³² As Section I emphasizes, rebalancing adds value over time, but such counter-cyclical strategies go against investors' behavioral inclinations. The rebalancing framework is a tremendous comparative advantage which minimizes the risk of time-inconsistent actions where the Fund minimizes the temptation to over-react and change its asset allocation in response to short-term noise.

III.A.3: Investment Mandate

The Fund's mandate states that "The Bank shall seek to achieve the highest possible return, net of costs" subject to some risk appetite constraints. The Fund's mandate concerning responsible investment practice states first that, "The management of the investment portfolio shall be based on the goal of achieving the highest possible return" before describing the requirements of Norges Bank's active management regarding responsible investing, active ownership, contribution to international standards regarding responsible investment practice, and environment-related investments. Thus, environmental, ethical, and other considerations play important roles in the investment style taken by Norges Bank, but it is within the context of an investment mandate.

Excluding companies on the basis of *any* criteria, even if they are well-motivated from the preferences of the asset owners, must result in lost investment opportunities, except in the cases where the exclusion constraint is not binding because the fund manager already chooses not to hold any of those stocks. That is, investment restrictions must create weakly inferior portfolios. In the Fund's context, a non-investment mandate may impair the welfare of future generations at the cost of meeting non-investment objectives imposed by the current generation.

As the Fund grows in size, there will likely be greater temptation to allocate more capital to non-investment activities. CalPERS is a good example not to emulate: a Columbia Business School case study, "California Dreamin': The Mess at CalPERS," points out that "California

³² CalPERS, for example, did not put in a formal rebalancing policy until 2009. It reduced its equity weight from 56% at June 30, 2008 to 49% at June 30, 2010. See "California Dreamin': The Mess at CalPERS," Columbia CaseWorks case study, ID#120306.

politicians have historically been more concerned with maintaining a voter-friendly image of CalPERS rather than letting it focus solely on achieving high risk-adjusted returns.” The risk is more than the ad-hoc and inconsistent changes in investment policy induced by political interference. Well-intentioned mandates directing capital to specific regions, or to spur economic development within local jurisdictions, tend to perform worse, on average, than unconstrained benchmarks. Hochberg and Rauh (2013) show that U.S. public pension funds over-weight in-state investments, which perform lower by two to four percentage points than similar out-of-state investments and similar in-state investments made by out-of-state investors. Bernstein, Lerner and Schoar (2013) report that, “sovereign wealth funds with greater involvement of political leaders in fund management are associated with investment strategies that seem to favor short-term economic policy goals in their respective countries at the expense of longer-term maximization of returns.”

The Fund has so far fairly successfully resisted the temptation to direct active management with a non-investment mandate. The exception is section 2-4 of the current mandate of the Fund which directs the Fund to establish environment-related mandates, which “shall normally be in the range of 20-30 billion kroner.” While very small relative to the total size of the Fund, this is still a restriction on the Fund’s investments. A recent white paper by NBIM, “Experience with Environment-Related Mandates,” reports that these investments have underperformed since the end of 2009 by 40%. The environment-related investments also carry costs eight times that of the entire Fund. Removing this explicit constraint can potentially improve the Fund’s investment opportunity set.

III.A.4: Professionalism

Talent and experience are essential requirements for success in any active investment program, and they are especially important in private market investing where extensive interaction and negotiation with external parties is an integral part of the investment process. Experienced professionals are highly sought after by all types of investors and intermediaries in both public and private markets. Most institutional investors have realized that they need to have in place a sufficiently strong and compelling employment value proposition to succeed in the competition for talent.

The Fund’s size, stability, and relative budget allow NBIM to potentially hire some of the best management talent in the world. The organization’s public service mission resonates with many recruits, it has global reach and impressive scale, possesses an established and well-known reputation, and has a “pure” investment focus without the asset gathering and client servicing requirements of some other investment organizations. However, while the Fund is able to pay more than a traditional government bureaucracy, it is constrained by social norms from replicating compensation practices at some investment banks, hedge funds, and private equity

groups. Some of these constraints are consistent, however, with the desire of the Fund to deliver active management in a cost-efficient manner. The senior staff of NBIM appears to be committed to the success of the Fund in a way that extends beyond their direct compensation, which is often lower than commercial institutions with comparable assets under management.

Merely being able to hire the best in class is not a guarantee of generating risk-adjusted returns. Dysfunctional governance structures with excessive asymmetric information, adverse selection, and the inability to monitor, all work against management taking advantage of investment talent. Being able to accurately measure skill is a precondition of efficiently managing it, and transparency enhances the Fund's ability in this regard.

As the Fund moves further into private real estate markets, it would be important to assess whether NBIM has the willingness and ability to implement compensation plans that are sufficiently competitive to attract and hold investment professionals with the desired experience and profile. Without more information, we do not explore this issue. A corollary consideration is the willingness of the organization to accommodate and embrace groups that have cultures and ways of operating that will be distinct from the public market programs which have historically been the mainstay within NBIM. The practical experience gained by NBIM in building out its real estate program over the next few years will be instructive in this regard.

The following quote from a spokesman for CalSTRS, one of the largest public pension plans in the U.S., is relevant: "The reason our policy prohibits (solo) direct investment is that, as a public entity, CalSTRS does not have the level of staffing nor the compensation levels required to adequately lead, review and monitor deals at the level required in the direct-invest environment."³³ While some might express regret that a fund as large as CalSTRS is not able to undertake direct investments, it is commendable that they have assessed and sensibly concluded that the constraints of their human resource policies preclude them from building the in-house capabilities required to succeed.

Finally, we note that NBIM's network of international offices outside Norway provides a strong benefit for any active management programs that might be undertaken. NBIM's already established offices in Oslo, London, New York, Singapore, and Shanghai confers home-ground advantages, especially for any new private market investment programs undertaken in those geographies.

³³ Quoted by Jacobius, A., " 'Solo' Investing Makes for Good Returns, Pensions & Investments, Nov. 7, 2012.

III.B: Frameworks for Delegating Investment Management

A large degree of delegation requires skill by all parties—the asset owner, the oversight body of the investment manager, and the investment manager itself.

We remark on the general structure for delegation in asset management before describing the specific framework at several funds, especially with regard to investments outside traditional equities and bonds.

There are generally three parties involved in a typical investment management framework: the asset owner/fund sponsor, the oversight body of the investment manager (we call it the “board”), and the investment manager itself. Ultimately the degree of delegation is determined by the asset owner and depends on the investment objectives and expertise of the asset owner, the capabilities of the board and investment manager, the degree of trust between the asset owner and board, and the desired mix of active and index management, among other factors. For example, a risk-averse corporate pension plan sponsor with little or no internal investment expertise might quite appropriately choose a passive asset mix for the pension fund and only delegate execution responsibilities to an oversight committee. In contrast, there are many examples of smaller endowments where the asset owner specifies a real return expectation and all investment decisions are made by a highly engaged and qualified investment committee.

Most large pools of assets including pension plans, sovereign wealth funds, and large endowments have in place some version of the three-party delegation model. It is appropriate and common practice for the asset owner to define return expectations and risk tolerances for the management of the pool of assets. Where the asset owner has the expertise, the task of translating these return objectives and risk tolerances to an investable portfolio is sometimes retained at that level, as is the case of the Fund and several of its peers. In contrast, many large public sector pension plans delegate responsibility for defining the investable portfolio to the board of the asset manager.

We note some key considerations for the asset owner in deciding the investment management decision rights to retain or to delegate to the board and investment manager:

1. Active versus index management

A passive, index approach requires little delegation other than execution responsibility; the decision to engage in active management in any form necessitates delegation. The amount of delegation depends on the relationship between the asset owner and board, and the experience and reputation of the investment organization. Even for pure “passive” approaches, active decisions need to be taken by the asset owner or board on

benchmarks, rebalancing policies, and possibly factor exposures that can be captured by public indexes (see Section I).

2. Capabilities of the three parties

It is more the exception (with Norway and Singapore as prime instances) than the rule that the asset owner has the necessary internal expertise to make key investment strategy decisions and oversee active investment programs. For many pools of assets, boards contain many individuals chosen for representative purposes rather than investment expertise and there are severe constraints on the ability to staff the investment manager with highly experienced and skilled professionals; for such funds there may be limited ability to delegate responsibilities. In contrast, large U.S. university endowments and some of the large Canadian pension plans have well-qualified and experienced people serving on boards and skilled professionals within the investment manager; for these institutions both the opportunity, and desirability, for optimal delegation of investment decision making to these parties is considerably greater.

A larger degree of delegation, with more freedom to deviate from benchmark, requires high levels of skill by all parties, all else equal: the asset owner to understand potential downsides and map strategy taking into account the uncertainty of possible active losses, the board to monitor and verify, and the investment manager to implement.

3. Governance

Stakeholder considerations are important determinants of the delegation framework. Generally speaking, the more independence granted the management of the asset pool, the greater the opportunity for delegation across the three parties.³⁴

4. Decision processes

In cases where the asset owner has the capabilities and wants to fix a rigid asset allocation and sets associated benchmarks for the fund, then delegation to the board and asset manager centers mainly on oversight and monitoring the execution of the corresponding investment programs. An important special case of this framework is the “endowment model” which relies on specifying separate benchmarks for each asset class and advocates large allocations to private and illiquid assets (see Swensen (2000)). These initial allocation decisions by the asset owner, however, can take considerable time (witness the Fund’s process to approve an allocation to real estate and the gradual increase in the equity proportion), and are sometimes not easily altered. If there is an

³⁴ Dyck, Morse and Pomorski (2013) report that funds with more independent investment managers generate higher performance. Note that this is not a causal relationship.

acceptance of the concept of time-varying risk premiums and a belief that a more dynamic approach to investment decision making and portfolio construction can add value, then it is logical to involve the parties closest to market developments, the board and asset manager, in those activities. When the investment manager is very skilled, a large degree of delegation is optimal.³⁵

Any specific delegation framework needs to incorporate and balance these different factors. We now discuss a specific delegation framework, the “Opportunity Cost model,” with these factors in mind.

III.C: Investment Framework at Canada Pension Plan Investment Board

The Opportunity Cost Model does not specify asset classes, and instead puts the onus on the fund manager to decide how to outperform the benchmark Reference Portfolio.

One way of improving the Fund’s expected return and risk relative to the current strategic benchmark is to enlarge the number of asset class benchmarks. This was done in 2010 with the inclusion of an allocation of up to 5% for real estate, coming from the fixed income allocation of the Fund. This approach retains decision rights and responsibility for asset allocation with the Storting and the Ministry of Finance.

An alternative approach, and one that is used for the governance of some other single purpose pools of assets similar to the Fund, is to give additional flexibility to the Fund’s manager, NBIM, to decide whether and how to expand the strategic asset allocation. In these models of delegation, as discussed in Section III.B, one of the primary roles of the asset owner is to establish clear return objectives and risk appetite parameters for the management of the fund.

III.C.1: Reference Portfolio

One way to do this is by specifying a Reference Portfolio (a similar concept to the strategic benchmark for the Fund prior to the inclusion of the Real Estate allocation), which reflects in a clear and simple way the desired or necessary level of systematic risk required to achieve the Fund’s objectives. The asset owner, which in the context of Norway would be represented by the Ministry of Finance, specifies an “active risk” appetite, which is the degree to which the fund can deviate from the composition of the Reference Portfolio, through means such as limits on the deviations that can be taken from benchmark (conventional “tracking error”),

³⁵ See He and Xiong (2013).

rebalancing requirements, and other risk measures. Responsibility to determine asset allocation and to approve investment programs within the defined risk appetite constraints is then the responsibility of NBIM. In practice, the fund manager would likely define a Policy Portfolio which would reflect its asset allocation decisions and corresponding approvals for investment programs. NBIM's ORP might serve as a basis for the Policy Portfolio under this framework.

The Reference Portfolio as first adopted at Canada Pension Plan Investment Board (CPPIB) in 2006 represents a passive investment strategy that, with the contributions to the fund, is reasonably expected to meet the liabilities of the fund. The Reference Portfolio is established by the board of directors at CPPIB and represents "the minimum level of systematic risk" required to meet the objectives of the fund. Since it represents a truly viable, passive management alternative, it serves as a good benchmark for active management. Any deviations taken from the Reference Portfolio, taken with the view to generate superior risk-adjusted returns, are owned by the fund manager. Consequently, the active decisions must generate value, on average, in excess of the Reference Portfolio.

The concept of the Reference Portfolio has been outlined in CPPIB's 2010 and 2012 annual reports, and in a Columbia Business School case study published in 2012, "Factor Investing: The Reference Portfolio and Canada Pension Plan Investment Board." The main elements of this governance model, with some variations reflecting specific circumstances has been adopted by GIC Private Limited (GIC), one of the sovereign wealth funds of Singapore, the New Zealand Superannuation Fund, and a number of other funds. An excellent description of the various responsibilities and accountabilities for the parties involved is provided in the feature article titled "GIC's New Investment Framework" in the "Report on the Management of the GIC Portfolio for the Year 2012/2013." We refer to this approach as the "Opportunity Cost Model" for active management.

A hallmark of the Opportunity Cost Model is that the fund manager is free to invest in assets not included in the Reference Portfolio, but all investments are benchmarked against it. The Reference Portfolio could be implemented by a small team of 10-12 people, at almost zero cost. The fund manager is deemed to have added value if it covers its own costs and beats the Reference Portfolio. CPPIB's Reference Portfolio consists of 65% equities and 35% fixed income, with further breakdowns for domestic and international assets. GIC's Reference portfolio is also a 65% equities and 35% fixed income portfolio. These portfolios are simple, highly scalable, and can be implemented very cheaply.

III.C.2: Total Portfolio Approach

In the Opportunity Cost Model, the fund manager looks beyond asset class labels and considers each investment in the context of its exposure to the underlying factor drivers captured in the Reference Portfolio. In this way, all investments are characterized in terms of their funding costs based on securities in the Reference Portfolio. For example, instead of directly specifying a private equity benchmark, CPPIB recognizes that private equity investments embody some elements of equities and bonds. In particular, private equity investments are often levered—economically this is equivalent to a short position in fixed income assets. Similarly CPPIB does not consider real estate a separate asset class, but funds real estate investments with a mixture of equities and bonds.

CPPIB refers to this decomposition of asset returns into Reference Portfolio factors as the “Total Portfolio Approach.” Effectively, CPPIB funds the investment in alternative investments by transfers from its Reference Portfolio. Appendix B gives two examples of how the Reference Portfolio is used in this approach. CPPIB also reports on its balance sheet the economic exposures of each alternative asset class in terms of its Reference Portfolio factors. Appendix C goes through a stylized example of decomposing direct real estate into a mixture of Reference Portfolio factors of equities and bonds. In practice, the funding mix can be sharpened by matching the alternative asset by public securities by geography and risk characteristics. The more precise the matching, the more complex the decomposition problem becomes. The asset manager has a good incentive to find a good funding mix because deviations from the benchmark are capped (through the risk limits imposed by the asset owner) and benchmark deviations are exacerbated, over the long run, by poor funding decompositions. There is a trade-off, however, between precision in the funding mix and complexity; the level of skills and infrastructure necessary to implement the funding decomposition increases in the level of granularity of the matching procedure.

One of the investment principles underlying the Total Portfolio Approach is that reliance on asset labels and a singular benchmark masks the heterogeneous nature of assets, especially in private markets. A fully leased and well-located core office building represents different risk and return attributes for a portfolio than a retail shopping center in a secondary market undergoing necessary refurbishment, notwithstanding that both are labeled as real estate. This would be equally true of two private companies with similar business models but significantly different capital structures.

Another principle is that the risk and return attributes of individual assets can be represented by some distinct combination of the factors in the Reference Portfolio (equity and bond exposures). The nature and degree of certainty of the cash flows as well as the capital structure of each asset determines the mix of asset exposures used to fund that asset. Investments such

as real estate, private equity, high yield bonds, and infrastructure can have idiosyncratic characteristics not necessarily captured in the funding assets that in turn may provide beneficial risk-adjusted returns for the portfolio.

The economic theory behind the Total Portfolio Approach is essentially a restatement of the CAPM, and the Arbitrage Pricing Theory of Ross (1976) for multiple factors. Assume the CAPM holds and asset i has a beta of 1.3. Then, the expected return of asset i can be represented by

$$E(r_i) - r_f = 1.3 \times E(r_m - r_f),$$

where $E(r_i)$ is the expected return of asset i , $E(r_m)$ is the expected return on the market portfolio which we will proxy as the equity market, and r_f is a risk-free bond return. We can equivalently write:

$$\underbrace{E(r_i)}_{\$1} = \underbrace{-0.3 \times r_f + 1.3 \times E(r_m)}_{\$1}, \quad (4)$$

which states that the opportunity cost of investing \$1 in the asset is a replicating portfolio which shorts 30 cents in bonds and takes a leveraged position of \$1.30 in equities.³⁶ If the CAPM does not hold, then the investor may choose the asset rather than the risk-adjusted benchmark of equity and bonds because being exposed to the idiosyncratic risk of the asset is a good investment opportunity; the idiosyncratic returns of the assets not attributable to stock and bond indexes offer a non-zero expected return.

A mathematical interpretation of the Total Portfolio Approach based on equation (4) is that the Reference Portfolio contains well-understood sources of risk premiums (stocks and bonds) that can be harvested cheaply. An investment in a private market is projected onto those factors. The factor loadings in the projection (the beta of 1.3 in the example of equation (4)) allow an investor to better understand the underlying risk factors of the investment. The orthogonal, or unexplained, component of the investment return represents the part that cannot be replicated (“spanned”) by a portfolio of stocks and bonds. If the investor has a competitive edge in information processing and trading (see Section I.E), then this is precisely the component of the asset’s return that is most attractive. In fact, the orthogonal return components are the *raison d’être* for investing in alternative assets like real estate, private equity, and infrastructure. A focal area for investment research and risk modeling under the Opportunity Cost Model is to elucidate the nature of these orthogonal risks, including factor exposures not

³⁶ For an interpretation of factor regressions used in performance evaluation along these lines, see chapter 10 of Ang (2014).

available in public markets and idiosyncratic risks which represent opportunities for security selection (see Section III.D).

The replicating factor benchmark has tremendous benefits in benchmarking private market investments. The Total Portfolio Approach forces an investment manager to articulate the marginal benefits of these investments in excess of zero-cost equity and bond index portfolios. The fund manager must cover its own expenses and beat the risk-adjusted factor benchmark, which is not easy given the complexities of these investments. Any new private market asset is evaluated on its opportunity cost of easy-to-implement public market investments. There are also benefits for public market investments. For example, an asset class like high yield bonds has characteristics between equity and safer traditional sovereign and investment-grade bonds. A concentrated position in a public market company is benchmarked against the opportunity cost of investing in a more diversified public portfolio.

Analyzing assets according to their equity and bond return equivalents allows for much greater comparability of investment opportunities and more refined decision making. This approach also results in an explicit risk factor-based approach to portfolio construction rather than simply using asset labels. One of the lessons of the global financial crisis was that equity-like risk was present in many investment portfolios under the guise of different names and structures; the total portfolio approach seeks to identify this ex ante instead of realizing it after the fact.

III.C.3: Governance

Perhaps the most attractive feature of the Opportunity Cost Model is a clear delineation of accountabilities. The asset owner makes the most important decisions about return objectives and the required level of systematic risk by approving the composition of the Reference Portfolio and specifying other risk appetite parameters. The board of the fund manager ensures that the manager has the requisite capabilities in place before commencing investment programs, monitors the execution and results of those programs, and approves appropriate internal benchmarks. The fund manager then has the latitude to make investment decisions and shape the actual composition of the portfolio subject to an active risk limit determined by the asset owner. Arguably, this places appropriate accountability with the party best able and positioned to make informed decisions. The onus is always on the fund manager to justify costs of active management and to outperform the Reference Portfolio. Thus, all alternative assets use the opportunity cost of the Reference Portfolio as their benchmarks, rather than having to specify separate benchmarks for each alternative asset class. This also provides more flexibility for the fund manager to take into account the time varying nature of risk premiums by not having fixed allocations to asset classes. The manager has appropriate incentives to make the investments with the best marginal contribution to risk and return for the overall portfolio rather than the best available investment within each asset class.

While conceptually simple, the Opportunity Cost Model is challenging to operationalize. The determination of an optimal Reference Portfolio requires extensive modeling of portfolio performance under different return assumptions, economic scenarios, liability behaviors, and constraint assumptions to determine the most robust portfolio composition to meet the plan's objectives. At first glance, having the fund manager decide on an appropriate funding mix of a private market investment, or a concentrated public markets one, may appear to give an incentive for the fund manager to "game" the factor benchmark. This is countered by the fact that the investment deviates from the fund's benchmark, and that any deviation is costly in terms of consuming part of the active risk budget. The fund manager wants to find an appropriate risk-adjusted funding mix that minimizes benchmark deviation volatility. However, computing the funding mix is not trivial and requires the developed comparative advantages of skill and professionalism.

The Opportunity Cost Model is only suitable for true, long-horizon investors. Inevitably there will be performance differences in any short time period between a real estate property, for instance, and the mix of equity and bonds used to fund it. Indeed, it could be as long as 7-10 years in the case of real estate for the Total Portfolio investment thesis to be realized (see Appendix C); needless to say, this is well beyond the time frames of the vast majority of investors and only possible with the discipline of long-horizon investment and measurement processes.

The traditional approach of defining a policy portfolio determines fixed, or slowly-moving, exposures to asset classes. Once done, subsequent investment decisions are taken with a view to filling those predetermined allocations with suitable portfolio managers or investments. This is quite different from the Opportunity Cost Model where every investment actually requires two decisions—first the asset to add to the portfolio (the "long" decision) and of comparable importance, the mix of Reference Portfolio components to sell which funds the asset purchase (the "short" decision). (Note that the short decision may involve leverage, as is the case in the simple example in equation (4)). Thus, the Opportunity Cost Model recognizes the fact that traditional "asset classes" contain stock and bond risks, whereas institutions can "double count" these exposures with alternative asset classes. The funding decision also provides a risk-adjusted benchmark to measure value-added of the long decision. In the Opportunity Cost Model, the fund manager is not forced to adhere to fixed asset class positions and can change these allocations based on a view of whether they are cheap or expensive. The two legs of every investment decision in the Opportunity Cost Model, however, are a much more time and data intensive process than traditional asset class investing. Not only does it require highly skilled investment teams to identify and structure attractive investment opportunities, it also requires a highly experienced and capable portfolio construction group.

Such a division should operate independently of the investment teams to determine the funding basket for each investment opportunity. The group needs to be familiar with each investment program and have sufficient resources to operate within the time constraints inherent in transaction-oriented situations. There should be a culture of collaboration between the groups so that information flows freely between them to enable the best long and short decisions for the total portfolio. This structure creates an inherent tension between the portfolio group and the investment teams—a tension that is healthy as long as there is a balance of capabilities and respect across the groups, and there is active support for the portfolio construction function from the senior management team.

A final governance challenge is that the board or similar governing body of the investment manager must have the skills to oversee, evaluate, and monitor the systems created by the investment manager to implement the Total Portfolio Approach. The board must also be conversant with the construction of, and the concepts behind, the Reference Portfolio.

III.D: Challenges and Approaches of Investments in Alternative Assets

Developed advantages explain the success of the most successful institutional investors in private markets.

The mandate of this report asks us to specifically review investment strategies “outside of GPFG’s current investment universe.” We focus on private equity and unlisted infrastructure as these are areas where the Fund has yet to commit capital. In its Report to the Storting 15, the Ministry of Finance decided not to open up the Fund to investments in private equity and infrastructure, but indicated it would revisit those types of investments at a future date.

The Ministry of Finance has asked us to consider “comparable investors’ use of strategies to exploit these investment opportunities.” Our intent, following the mandate, is not to explicitly recommend whether investments in any specific alternative asset classes should be undertaken; rather, we describe the best practices the Fund could emulate, were it to choose to enter these areas. We concentrate on comparable funds’ approaches, including benchmarking and governance, rather than describing risk-return properties. It is our opinion that because of the heterogeneity of both investors and the investments within each asset class, such reports on “average” returns are not very relevant for the Fund. We interpret the Fund’s set of “comparable investors” narrowly—although many funds have alternative investment programs, many of them lack one or more of the Fund’s natural advantages (like the absence of immediate liabilities), or developed advantages (like professionalism), or both (see Section III.A). The small peer group steers our investigation towards qualitative analysis.

Quantitative analysis is difficult even for a broad set of institutional investors' experiences in alternative investments because of deficient data, as described in the Report to the Storting 15, and in separate reports by Ludovic Phalippou ("An Evaluation of the Potential for GPFG to Achieve Above Average Returns from Investments In Private Equity and Recommendations Regarding Benchmarking") and de Jong and Driessen ("Capturing Illiquidity Premiums") which were commissioned by the Ministry of Finance. These studies also describe the challenges of measuring returns in illiquid markets (see also Appendix C). Two recent papers provide some observations about private market investing by large institutional investors: Dyck and Pomorski (2011) and Fang, Ivashina and Lerner (2012). Neither is definitive or is meant to represent comprehensive examinations of the experience of institutional investors in private equity and infrastructure investing over long periods of time.

Our focus for this and the next two sections is to describe the strategies in private equity and infrastructure taken by GIC, PGGM (a large Dutch pension fund), Ontario Teachers' Pension Plan (OTPP), the Ontario Municipal Employees Retirement System (OMERS), and CPPIB. We include some specific indications of scale and performance where available. These institutions share the same structural advantages as the Fund (see Section III.A) and attempt to incorporate them within their investment strategies. In addition to their broad public markets investment activities, these organizations also have extensive experience and successful track records in private equity, real estate, and infrastructure investing, both through funds and as direct investors. Each of these organizations also exhibits a set of essential attributes which, especially in the case of direct private market investing, have been key enablers of their success. That is, these funds have succeeded in cultivating developed advantages in their alternative investments programs.

III.D.1: Rationale for Alternative Asset Investments

We first articulate why an investor might move beyond publicly traded, liquid investments using the framework of Section I:

1. Diversification

Basic theories of diversification recommend holding a substantial portion of an investment portfolio in assets that constitute the majority of world wealth: real property and unlisted equity. Private equity markets are larger, especially in emerging markets, than public equity markets. As governments transfer assets from the public to the private sector, investments in infrastructure will increase in importance. Viewing private assets as part of the "market" portfolio does not require a belief that they outperform public markets (although the Opportunity Cost Model described in Section III.C provides one way of benchmarking them against equities and bonds) for investors

to benefit: as long as there are some common movements in illiquid asset returns with other assets, then factor model theory predicts that investors can construct portfolios with superior risk-return trade-offs (see Ross (1976)). This arises because idiosyncratic movements of returns are diversified away. The resulting factors also have higher Reward-to-Risk ratios—even for the market portfolio factor. While there is always theoretical benefit, the practical benefit of additional diversification may be small if the alternative assets are largely “spanned” by returns available in public markets.

2. Rebalancing

The inability to trade when desired is a handicap in holding illiquid investments. In equilibrium, there should be a premium for bearing illiquidity risk, but it may not necessarily be large. (Ang, Goetzmann and Schaefer (2011), for example, report that there is little evidence for higher risk-adjusted returns in illiquid asset classes.) Since the payoffs of illiquid investments may occur many years in the future, short-term investors with immediate liquidity needs cannot hold them. Consequently, a long horizon is a necessity for investing in illiquid assets, but not all long-horizon investors should hold them. The capacity to manage illiquidity risk and constraints differs widely across investors, and thus institutions should assess their developed advantages before embarking on illiquid investment programs.

3. Factor exposures

Most alternative investment returns are exposed to the same macro environments as public equity and fixed income markets, and thus most illiquid assets exhibit significant co-movements with stock and bonds. It is important to reiterate point (1) that even if no new factor exposures exist in alternative asset markets, they still provide diversification benefits. There are, however, some common components in alternative asset classes that are independent of stocks and bonds. Appendix C, for example, shows that although there is a large component of returns specific to real estate markets after accounting for equity and bond factors. The proportion of returns in private equity investments that cannot be replicated by passive equity and bond indexes is much smaller than real estate investments, but Ang et al. (2013) document that that it is statistically significant and exhibits unique properties.

4. Security selection

Just as in public markets, many studies are skeptical that active management in real estate or private equity markets, or that the average returns in these alternative asset

classes outperform public indexes.³⁷ By themselves, these are not compelling reasons to avoid investments in illiquid assets. The literature also documents wide dispersion of returns by managers in illiquid markets, much wider than the dispersion generated by active management in public markets. Furthermore, the degree of excess performance is highly persistent. Combined with the large asymmetric information, adverse selection, and non-public transaction prices in private markets, many investors simply do not have skills to find and screen managers delivering value-added returns and can be locked into underperforming investments for many years. The flipside of these inefficient markets is that there is great opportunity for skilled, long-horizon investors to generate value by careful security selection.

In summary, the benefits of investing in alternative assets are (1) increased diversification, (3) potential exposures to new factors not available in public markets, (4) and the opportunity for security selection. These benefits do not come without costs, one of the most important being (2) the hindrance to rebalancing. The Opportunity Cost Model (see equation (4) in Section III.C) explicitly captures (3) and (4) by measuring the component of alternative asset returns that is orthogonal to stocks and bonds. The benefit that informed and skilled investors can reap from successfully accessing new factors or generating value in security selection is greatly facilitated by using a framework like the Opportunity Cost Model which requires them to elucidate the nature of these orthogonal risks.

III.D.2: Characteristics of Successful Large Investors in Alternative Assets

As we see it, there are several common characteristics of successful institutional investors in private markets:

1. Rigorous and consistent investment processes

Each direct investing opportunity whether in real estate, private equity, or infrastructure tends to have its own distinct attributes and risk factors. As these programs grow in size and geographic reach, there is a need to create a number of different internal investment teams to assess the different opportunities. Successful institutions have learned the importance of having highly disciplined and repeatable investment processes that permit meaningful comparisons across opportunities and allow for well-informed investment decisions. While achieving this consistency is no small task, it is also a defining characteristic of successful institutional investors.

³⁷ For a summary of active management in real estate, see Ang, Goetzmann and Schaefer (2011). Recent summaries of private equity performance are Ang and Sorensen (2012) and Harris, Jenkinson and Kaplan (2012).

2. Responsiveness

To achieve success in private market investing, institutional investors have to conform their decision making and approval processes, including approvals by their governing boards where required, to the timetable dictated by the market. A culture of responsiveness (or its absence) is readily visible to external parties and is essential to be taken seriously as a partner or counterparty.

3. Ongoing monitoring and asset management capabilities

One of the associated obligations of a large ownership stake in a private company, infrastructure asset, or property is governance. Owners have a responsibility to ensure that proper oversight mechanisms are in place and function well; indeed, some institutional investors view having a direct and ongoing role in governance as an integral part of the overall investment process. When investments are made alongside a specialist investment partner, as is sometimes the case with private equity and real estate, some institutional investors prefer to rely on their partner to handle governance responsibilities. This is not usually feasible, however, for infrastructure investments as well as for many direct private equity investments. In these instances, the institutional investor needs to have the internal capabilities to fulfill the required governance and asset management roles. These kinds of skills are often different from those of the transaction teams; associated resource requirements can be considerable and grow in direct proportion to the number of investments undertaken. One way that some institutions have met the need for governance resources is by creating a network of qualified external parties who can serve as nominee directors for their portfolio companies.

4. Boots on the ground

Private markets are characterized by the inability of all investors to have the same information—unlike public markets where regulation dictates that investors should have the same set of information, but have different skills to process or respond to that information. Both OTPP and CPPIB, for example, now have offices in London and Hong Kong, and each would confirm that having a local presence in those geographies has brought significant benefits to their real estate, private equity, and infrastructure programs in terms of information flow, relationship building, awareness of local market developments, and ability to monitor existing investments. The presence of multiple offices around the world for NBIM could be used in a similar capacity.

With these rationales and developed comparative attributes in mind, we now turn to private equity and infrastructure markets.

III.E: Private Equity

A distinguishing characteristic of the most successful investors in private equity is that they undertake direct investing as well as investing through funds.

A recent report by Bain & Company estimates the total invested and committed capital globally within private equity at \$3 trillion—an amount which represents a considerable market opportunity for investors.³⁸ A number of institutions who have been active in private equity for a number of years have amassed sizable portfolios. CalPERS and CPPIB have recently disclosed total private equity investments of \$34.2 billion and C\$32.6 billion respectively.³⁹ Although GIC does not disclose amounts, its target allocation to private equity is 11-15%, which likely represents a considerable dollar amount.⁴⁰ AlInvest, which until its sale to Carlyle Partners in 2011 was the private equity investment arm for PGGM and APG, has disclosed total investments and commitments of €52 billion.⁴¹ Consequently, there is sufficient evidence that, even though the Fund is large, meaningful scale could be achieved over a reasonable period of time through private equity investing.

It would be challenging to think that a large institutional investor, especially one the size of the Fund, could achieve sufficient scale in private equity without incorporating some element of fund investing. Even the most active direct investors acknowledge an important role for private equity funds as evidenced by the following statement by a spokesperson from OTPP, “If we do not have expertise or relationships in a geographic market or industry, we mitigate our risks with funds and partners who do.”⁴² While the vast majority of pension and sovereign wealth funds invest in private equity exclusively through funds, a distinguishing feature of OTPP, OMERS, CPPIB, GIC, and PGGM is that they undertake direct investing in addition to their fund programs.

Phalippou (2010), in his report to the Ministry of Finance on private equity, observed that, “An investor without special selection skills ... is unlikely to earn an alpha from its private equity investment in the long run.” We agree with this assessment and also observe that the selection process for many investors is superficial at best and heavily reliant on the published performance numbers (with the inherent flaws noted by Phalippou) furnished by private equity managers.

³⁸ Bain & Company’s Global Private Equity Report 2013.

³⁹ Numbers from CalPERS Comprehensive Annual Financial Report 2012 and CPPIB’s Annual Report 2013.

⁴⁰ See <http://www.gic.com.sg/en/report/report-on-2012-2013/features#.UpPFEOLBYl>

⁴¹ AlInvest Annual Report 2012.

⁴² Quoted by Jacobius, A., “ ‘Solo’ Investing Makes for Good Returns, Pensions & Investments, Nov. 7, 2012.

III.E.1: Characteristics of Successful Investors

A distinguishing feature of funds like GIC, CPPIB, and others is that they have used their scale advantages to develop “special selection skills.” These are a special type of developed comparative advantage (see Section III.A). By virtue of their size and importance as existing or potential investors, these institutions can gain access to very detailed information about private equity funds that is not available to others. That access allows them to deconstruct performance for each individual investment in a fund to identify if, where and how value was created. Some of the relevant information gained through this process includes:

- The number of deals sourced through proprietary origination versus auction,
- How the original purchase price relates to market comparables,
- Public market equivalent (PME) performance calculations for each investment using actual cash flows;⁴³ calculations based upon the same sector and geography of the investment are especially useful,
- Decomposing realized performance to identify sources of returns such as top line growth, margin improvement, use of leverage, multiple expansion, and so forth,
- The track records of the lead private equity partners and various teams.

The effort required for this analysis and the system requirements to track the cash flows for each individual investment and its market equivalent are considerable; these can only be justified by the considerable scale of the private equity holdings of the large institutions. The asymmetry of information gained through this analysis compared to the average investor constitutes a comparative advantage. This information certainly does not guarantee the ability to select private equity funds that deliver alpha by outperforming the public market alternative net of all costs, but it does enable more informed decisions than the typical “picking past winners” approach. The information is also valuable for pricing opportunities in the secondary market for private equity funds and for deciding when to partner with private equity managers on specific direct investing opportunities. While the funds in the peer list we have selected do not separately disclose their private equity fund results, these institutions would confirm that they have historically achieved very good returns relative to their benchmarks.

Another characteristic of the most successful and largest institutional investors in private equity and infrastructure is that they have succeeded in bringing down costs relative to the typical investor. This comes about through being viewed as attractive partners through their

⁴³ See Kaplan and Schoar (2005).

developed advantages, increasing economies of scale, and the ability to bring in house many activities (see further comments below).

III.E.2: Direct or Indirect?

The direct private equity investing programs of institutions active in this area include investments that are sourced solely by the institution itself (or co-sourced with other institutions), as well as investments that are co-sponsored in conjunction with a private equity manager. Both of these are distinct from making syndicated co-investments offered by private equity managers after the close of a transaction. While large institutions may occasionally accept co-investment opportunities, these are certainly not the cornerstone of their direct programs.

In the co-sponsor model, the institutional investor relies on the private equity manager to originate the transaction, but then participates actively throughout the process and assumes its proportionate responsibility for diligence activities and costs, deal structuring, pricing decisions and ultimately commits capital on the same terms as the private equity manager. In the case of sole or co-sourced transactions, the institutions themselves originate the opportunity and then conduct and control all aspects of the transaction process. In their study of direct private equity investing, Fang, Ivashina and Lerner (2013) note that the sole or co-sourced model is most often successfully used in those markets where the institutions have either a “home” advantage or significant familiarity, as well as for investments in relatively low complexity businesses. Conversely, the co-sponsor approach is more common in other geographies where the institutions place reliance on the origination capabilities of the private equity manager and as well for higher complexity industries such as technology.

One advantage of the sole or co-sourced approach is that the institutional investors have full control over the hold period of their investments. An example is OTPP’s investment in Maple Leaf Sports and Entertainment which it purchased in 1994, held for 18 years and then sold in 2012 for C\$1.32 billion thereby realizing a substantial gain.⁴⁴ In the co-sponsor model, the hold period is influenced by the need of the private equity manager to realize proceeds from investments during the life its fund structure. That said, there are a number of examples where co-sponsoring institutions have provided full or partial liquidity to the private equity manager so that they can continue to hold attractive investments beyond the normal dictates of private equity funds.

Regardless of the form, direct private equity investing clearly places the onus of decision making on the institutional investor. This, in turn, means that the institution has to have a high degree of conviction about the merits of the investment and the opportunities for value

⁴⁴ See <http://www.forbes.com/sites/nathanvardi/2011/12/09/the-greatest-sports-deal-in-history/>

creation. Experience has shown that investors in private equity achieve value in a variety of ways. Maintaining strong discipline around acquisition pricing is especially critical and, in this respect, proprietary origination of transactions through proactive sourcing as compared to participating in competitive auctions makes a notable difference. Implementing an efficient capital structure and putting in place a clear operating plan which incorporates the key elements of the investment thesis are also important. Other examples of actions undertaken by private equity investors to achieve value include business line expansion, follow on investments, or divestitures of non-core divisions, ensuring that the management team has the right capabilities and their incentives are tied to value creation goals, and leveraging relationships with other industry participants. While the specific measures taken differ for each investee company, one consistent characteristic of successful private equity investors is that their approach to governance for the entities they own is considerably more active and involved than is typically the case for companies whose stock is widely held.

A necessary condition for successful direct private equity investing is that the institution has already established and ingrained the set of developed capabilities that were described in Section III.A. Having those capabilities entails higher internal costs for the institution compared to investing solely through funds. Fang, Ivashina and Lerner (2013) estimate the internal costs associated with direct investing programs at 0.91% of committed capital versus a cost (excluding fees) of 0.11% for investing through private equity funds. That cost differential is more than offset by the fact that direct investments do not attract the management and performance fees that are charged by private equity funds which, as Phalippou (2013) notes, have a substantial impact on net performance. Indeed, elimination of these fees, along with achieving greater control over the investment process, are the primary reasons that institutions undertake direct private equity investing. Avoiding fees, however, would be of small comfort if their direct investments did not achieve improved net performance relative to private equity funds. The large institutions in Fang, Ivashina and Lerner's study did in fact outperform various applicable benchmarks for private equity funds. Although not explicitly part of their public disclosure, the large institutions would confirm this conclusion for their own direct investing programs as well.

Two of the most experienced institutional investors in private equity are OTPP and OMERS, both of whose fund and direct programs date back to the early 1990s. OTPP notes that its private equity program has generated an internal rate of return of 19.2% since its inception in 1991, and a 12.5% annualized return over the last four years compared to its benchmark return of 7.8%.⁴⁵ OMERS has most recently disclosed a three-year annualized private equity return of

⁴⁵ See <http://www.avcj.com/avcj/news/2258784/otpp-private-capital-unit-realizes-usd5b-reinvests-usd3b-in-2012>

16.0% versus its benchmark return of 13.9%.⁴⁶ These few data points are certainly not representative of all institutional investors, but they do confirm that some institutional investors who have well established private equity fund and direct programs have achieved good results.

III.E.3: Benchmarking

There is no universal approach among institutional investors when it comes to benchmarking private equity returns. That said, virtually all of them use some mix of public equity indices as their key comparators. It is also common practice to add a required illiquidity premium to the index returns when measuring relative performance. Less common is the use of a beta adjustment to incorporate the impact of typically higher amounts of debt within the capital structures of private equity companies versus their public company comparables.

As one example, the benchmark used by OMERS for its private equity holdings is “the aggregate of Global Russell indices adjusted for geographic and sector exposures plus an illiquidity and leverage premium.” CPPIB uses a similar approach by choosing a sector and geographic index that most closely matches each of its private equity investments. CPPIB also typically assumes a market beta of 1.3 for private equity and incorporates this in its benchmarks—this entails funding private equity investments by selling a leveraged position in equities and going long a bond position from its Reference Portfolio (see Section III.C and Appendix B). More specific beta adjustments may be appropriate for large direct investments where the capital structure is known and can change over time. The OMERS and CPPIB approaches are consistent with Phalippou’s (2013) rationale and recommendation for benchmarking. These methodologies admittedly entail additional measurement complexity, but arguably also provide a better fit with the actual composition of their private equity holdings and a good basis for determining risk-adjusted, value-added returns.

III.F: Infrastructure

Infrastructure’s large scale offers attractive opportunities, but as investing is usually direct, in-house skill is essential—especially as there is no uniform index for benchmarking.

The OECD “Infrastructure to 2030” report estimates that, over this period, average annual expenditures will approximate \$3 trillion and the total value of infrastructure assets will grow to \$71 trillion. The OECD notes that there is a growing shortage of available equity to finance that growth, a situation that could be exacerbated by regulatory changes that may constrain

⁴⁶ OMERS Annual Report 2012.

traditional providers of capital such as insurance companies. A sizable portion of that infrastructure total will not likely be open to private ownership; nonetheless, infrastructure still represents a very sizable investment opportunity for institutional investors.

The OECD and other observers have noted that pension plans and sovereign wealth funds are in many ways ideal owners of infrastructure assets. They point out that the typical lengthy useful lives of these assets are a match for the long-duration liabilities of these funds. Pension plans and sovereign wealth funds have an ethos of public trust arising from their missions which aligns well with stewardship of what are often essential services to populations. They are also perceived to be less likely than corporate owners to focus on “short-termism profitability” at the expense of investing to properly maintain these assets over the long term.

A key source of the attraction infrastructure assets hold for institutional investors is that they can provide a stream of reasonably stable and predictable cash flows, often with some element of inflation protection, over long periods of time. However, as with real estate, this is far from a homogeneous “asset class” and so investors often categorize assets according to some common underlying characteristics. One category includes regulated assets such as water services and electrical or gas distribution systems. By virtue of the fact that these assets provide essential services and operate largely as monopolies, they are usually subject to some regulatory regime that dictates operating standards and caps a real return on invested capital. A second category would be volume-based infrastructure assets such as toll roads, and air and water ports. Revenues here are driven by usage and are more sensitive to macroeconomic factors than the essential services category. Another group is assets such as power generation facilities or pipelines that come with long term off-take or utilization contracts.

The predictability and risk profile of future cash flows is clearly different across each of the three categories. As well, there is a considerable amount of idiosyncratic risk associated with individual assets. Regulatory regimes can sometimes be subject to political influence; the value of long term contracts depends in part on jurisdictional standards for the rule of law, and public opinion can affect the operations of assets. The ability to recognize, mitigate, and price these kinds of risks are foundations of successful infrastructure investing.

III.F.1: Experiences of Peer Funds

The Canadian pension plans have been pioneers in infrastructure investing and remain among the most active in this sector. The most recently reported infrastructure holdings for OTPP, OMERS, and CPPIB were C\$9.6 billion, C\$11.6 billion, and C\$11.2 billion, respectively, placing them amongst the largest infrastructure investors globally.⁴⁷ In addition, all other large

⁴⁷ Numbers from OTPP Annual Report 2012, OMERS Annual Report 2012, and CPPIB Annual Report 2013.

Canadian pension plans have active infrastructure programs—as do many other funds including PGGM, GIC, New Zealand Superannuation Fund, and the Australia Future Fund.

Unlike private equity, large institutions tend not to invest in infrastructure extensively through external managers and fund structures. Transactions for these assets are usually done through very transparent processes so there is little benefit that intermediaries provide in deal origination. Infrastructure assets tend to have a low complexity profile and so are relatively straightforward to assess and operate. Moreover, the relatively short defined hold periods, and the relatively large fee structures, for funds are problematic given the nature and revenue streams of infrastructure assets. Accordingly, infrastructure investing for large institutions tends to be primarily a direct investment and ownership program. This requires internal teams with expertise and experience in this area (once again, the list of developed capabilities noted in Section III.A should be considered necessary but not sufficient conditions for success). Collaborative relationships between institutions active in infrastructure have arisen over the years, and it is now common to see several large funds form a consortium to jointly bid for an infrastructure asset. Often these consortia intentionally include a “local” institution, which can help to mitigate some of the political and regulatory risks.

To date, most of the infrastructure assets owned by institutions can be categorized as “brownfield” or well-established entities. These have obvious appeal as it is certainly easier to model risks and returns for assets with an operating history. The supply of brownfield assets is by no means exhausted, but the competition for them has heightened in recent years with a corresponding impact on pricing. Given that much of the additional future supply of infrastructure assets forecasted by the OECD will be in “greenfield” development, some experienced infrastructure investors are considering expanding their programs to take on the development and adoption risks in greenfield opportunities.

Greenfield investing is more time intensive and requires additional internal expertise. However, this is in many ways similar to the evolution within some of large institutional funds’ real estate programs, which often now include a “develop to own” element. Properly priced, the development premium can compensate for the additional risk, and can be an effective way in acquiring assets to build a sizable portfolio. Another possible future evolution for infrastructure investing could be the acquisition of operating companies to use as platforms for their investing and developing activities, as many have already done in the real estate sector. Indeed, the possibility of acquiring a platform company to pursue toll road opportunities globally was one of the stated objectives behind OTPP and CPPIB’s attempted privatization of the Transurban Group in 2010.

The infrastructure portfolios of large institutional investors have many examples of soundly underwritten and structured investments that have performed well to date. This is certainly not

a riskless area, however, as there are also instances where permitted returns for some assets have been unexpectedly reduced through regulatory revisions—which in some cases have resulted in valuation write downs. That said, our conversations with large institutions indicate that those which are active in this area would like to add to their infrastructure holdings where possible. We are not aware of any composite study of infrastructure investment returns for institutional investors and not all institutions separately disclose their actual and benchmark results. The reported returns for OMERS' infrastructure portfolio over the past three years are 12.8%, 8.8%, and 10.1% versus benchmark returns of 8.6%, 8.0%, and 8.5%, respectively.⁴⁸ These reflect the relative stability that comes with a mature asset profile.

III.F.2: Benchmarking

As with private equity, there is no uniform benchmark used by institutional investors for their infrastructure holdings. The most common approach is to adopt some form of absolute return benchmark which captures the desired attributes of the cash flows generated by those holdings. OTPP's benchmark is representative of this approach: Local CPI + 4% + a Country Risk Premium.⁴⁹ This benchmark reflects a desire for inflation protection, incorporates an illiquidity premium, and captures some of the jurisdictional and idiosyncratic risks through a portion of the 4% spread and the country premium. A drawback to having a single uniform benchmark for all assets is that it is not calibrated to underlying differences in risk profiles. Even if both were operating within the same country, the predictability and reliability of cash flows generated by a well-regulated water services company would likely be different from a volume-related asset such as an airport.

A single benchmark also does not take into account the capital structures of the assets; an otherwise low or moderate risk asset can be turned into a high risk investment through excessive use of leverage. A possible modification of the OTPP approach would be to adjust the 4% component of the benchmark according to the specific characteristics of the asset. Another benchmarking approach used by CPPIB is a composite of bonds and equity that match the characteristics of the asset. The bonds and equity assets are essentially factors in the Reference Portfolio which match the characteristic risks of the infrastructure investment. (This is similar to the opportunity cost captured in the PME for private equity.⁵⁰) Greater cash flow risk, or more aggressive capital structures, require higher proportions of equity and therefore entail higher required returns in this methodology. The Opportunity Cost Model used by CPPIB certainly involves more variability in benchmark returns over the short term, but can be an effective way to measure performance over the long hold periods associated with infrastructure.

⁴⁸ OMERS Annual Report 2012.

⁴⁹ See <http://www.otpp.com/investments/essentials/benchmarks>

⁵⁰ See Appendixes B and C for a decomposition of real estate into bond and equity components.

III.G: Summary

We summarize our investigation on the framework for active delegation as follows:

1. As important as the natural advantages of a long horizon, large scale, and the ability to hold illiquid assets are developed advantages of professionalism, transparency, and a dedicated fund manager.
2. We observe that organizations such as large U.S. endowments and many of the large Canadian pension plans that undertake active management in sources of risk premiums not available in standard indexes and which potentially vary over time, and who have experienced and highly qualified boards and investment managers, have high levels of delegation from the asset owner to the board and asset manager.
3. The “Opportunity Cost Model” takes advantage of a long investment horizon (natural comparative advantage) and skill (developed comparative advantage). This model is extensively used by CPPIB and GIC and represents a compelling alternative to traditional asset class portfolio construction and investing. In this model, all active decisions to deviate from benchmark are evaluated by the foregone opportunities to cheaply invest in the passive benchmark portfolio. While conceptually simple, it is challenging to operationalize.
4. The Fund’s peer group of global institutional investors has extensive and well-established private market investment programs including private equity and infrastructure in addition to real estate. There are challenges in data, measurement, and benchmarking in these alternative asset classes. There may be an opportunity for the Fund to achieve meaningful NOK or currency basket value-added returns by increasing its exposure to private markets. Developed advantages are important characteristics of the most successful institutional investors in private markets.

Section IV: Recommendations

Section I showed that long-term investors create value by choosing a well-diversified benchmark, rebalancing periodically to that benchmark, taking on factor risks, and security selection. The first three tend to be “systematic” or rules-based investment strategies with verification horizons stretching potentially to decades, while the last may involve more discretion and usually has a shorter verification horizon. Section II showed that although the amount of value added by active management is small because of the tight risk limits of the fund, there is a significant exposure to factor risks in the active returns of the Fund. The ORP is a new development which seeks to capitalize on the Fund’s advantages of scale, the ability to take illiquid positions, and the capacity to take positions reflecting factor tilts which are diversified across geographies. Section III emphasizes that professionalism and skill are important attributes that an organization endogenously develops, and these comparative advantages are just as important as the natural ones arising from the nature of a fund’s cash flows. Developed capabilities are important in implementing dynamic factor strategies, which are not as cheap as market capitalization-based index funds. Developed comparative advantages are even more essential as the Fund increases investments in illiquid asset markets—markets where passive management is not possible—yet, the lack of observable returns in these markets hampers the measurement of active management.

Based on these findings, we recommend more transparency of each stage in the investment process—first from the adoption of a market-weighted passive benchmark, to rebalancing, to adopting factor exposures and other non-market capitalization tilts, and finally security selection. As each stage is separately reported, the Fund can measure the value derived at each step and better manage risk. Importantly, such transparency can highlight the role of factors in driving the total returns of the Fund. Second, we recommend the Fund adopt the “Opportunity Cost Model.” This model provides an internally consistent benchmark and is a cheap, viable alternative to active management. It is straightforward to extend this framework to many alternative asset classes. The framework holds the fund manager responsible to cover its own costs in taking deviations away from benchmark, and consequently to justify its ability to actively manage the portfolio. We also recommend that permitted risk taking be increased. A higher risk budget accommodates the factor risks taken in the ORP and helps ensure NBIM is not penalized for poor performance due to factors, gives NBIM a larger buffer, and recognizes that some benchmark deviations are taken for what would be defined in a traditional sense as non-active management purposes, like rebalancing the portfolio or improved diversification. Ensuring robust governance is essential to successfully implementing all of these recommendations.

IV.A: Report Each Stage of Value Added

Transparency of the investment process allows all stakeholders to measure the value added at, and better manage risk arising from, each discrete phase of the investment process.

We recommend that returns be reported at each stage of the investment process as outlined in Figure 14. The figure describes the current management of the Fund including the ORP within NBIM. It breaks down the different parts for which the Ministry of Finance and NBIM are responsible. For some stages currently in the purview of NBIM, we also recommend that risk budgets be reported. Transparency at each stage of the investment process is consistent with the public disclosure mandate set by Parliament; the Fund is recognized to be the world leader in the sovereign wealth fund community in terms of disclosure practices. Section III.A describes how transparency is a developed comparative advantage of the Fund.

Our recommendations for more transparency encompass all stages of the Fund’s investment management. The mandate for this report concerns the “delegation to Norges Bank” taking as given the benchmarks selected by the Ministry of Finance. We do not make any comments concerning the current strategic benchmarks of the Fund, but we are moving a little outside the mandate in recommending greater transparency at all levels where strategic investments decisions are made. As Section I emphasizes, there are a series of active decisions in asset management—it is too simplistic to think of investing in terms of purely “passive” and “active” categories. NBIM’s decisions are not taken in a vacuum, quantifying the objectives behind the benchmark that is given to it allows more clarity in weighing the types of active management strategies which should be pursued. In particular, it allows NBIM to further the same strategic aims. Because of the continuum of active choices, some of the decisions regarding diversification, rebalancing, and factor choices (see Section I), could migrate between the Ministry of Finance and NBIM in the future; measuring the contributions at all levels helps to decide at which level different strategic decisions should be made. This is especially true for the factor exposure decision, which we comment on further below.

IV.A.1: Passive Benchmark

Following the framework in Section I, Figure 14 starts with the concept of the true market portfolio, which is only theoretical, but we include it to emphasize that the first aim of an investor is to hold a diversified portfolio.

Measurability starts by choosing a passive benchmark. It is a publicly-stated index from an independent index provider. This provides transparency, accountability, and verifiability. The advantages of starting with a market capitalization benchmark include that (a) it represents the typical, or average, investor; (b) represents investment capacity, and (c) there usually exist

index products, internally or externally managed, that can implement these benchmarks at effectively zero cost. As Section I emphasizes, the adoption of any passive benchmark is an active choice of the investor, and different passive benchmarks have different index constructions, liquidity requirements, and security screens. By making public these benchmarks, the Ministry of Finance already makes this stage fully transparent.

IV.A.2: Weighted Passive Benchmark

The Fund is different from a typical investor. Capitalizing on its natural comparative advantages and characteristics (see Section III.A), it moves from the passive market-weighted benchmark to non-market capitalization positions. In the taxonomy of Section I.C, these positions reflect a more optimal exposure for the Fund to risk premiums in liquid markets.

The Ministry of Finance currently reports some of the effects of these deviations from passive, market-capitalization weights. The ethical and socially responsible investing criteria remove some sectors, like tobacco, and some companies, like those deemed to violate human rights or which degrade the environment, from the index universe. The Ministry of Finance reports the benchmark return without the ethical screening exclusions compared to the full benchmark. We recommend that similar deviations from the market-capitalization benchmarks be reported for equity and bond portfolios, and for the Fund. In the sovereign bond portfolio, for example, there are country weights that are specified on GDP criteria. The differences for these holdings compared to a pure market-weighted position should be reported.

Because these deliberate tilts away from market weights implement long-run, optimal investment strategy, the market weight deviations should be interpreted carefully. We recommend reporting over long time periods and, reflecting the slow moving strategic investment changes of the Fund, reporting at infrequent periodic intervals. Examining the difference between the strategic weights and market weights can be construed as part of “benchmarking the benchmark” of the Fund. We suggest an appropriate time would be in adopting or revising the Reference Portfolio, should the Fund choose to implement the Opportunity Cost Model as we recommend in Section IV.B. Any revisions of the Reference Portfolio, and thus comparisons of the strategic benchmark to market weights, would be presumably done very infrequently.

Since the average investor holds market-weighted positions, reporting the return difference from this step shows how the Fund is actively choosing a benchmark different from the average investor. Some of these deviations may not result in higher risk-return trade-offs, but reflect the preferences of the ultimate owners of the Fund. On the other hand, if market prices of sovereign debt do not solely reflect investment risk because certain sovereign debt issues fulfill important liquidity and reserve requirements, then holding non-market weighted positions

based on fundamentals may result in excess returns. All of these decisions are conscious choices of an investor away from a market-weighted portfolio and should therefore be measured.

One disadvantage in assessing the market weight deviations taken by the Ministry of Finance at a high frequency is that it might lead to an attitude of “short termism” where excessive comparisons create a tendency to revise the optimal strategy too frequently, and a reluctance to deviate from market weights.⁵¹ We believe transparency, done at the appropriate frequency and with long samples, leads to more robustness when combined with regular communication of the purpose of deviating from the market and appropriate management of expectations and risk. The Fund is not an average investor: when greater transparency is combined with clear statements why the Fund has deviated from the market in the first place, the Fund should have greater conviction to stay the course during difficult periods. Ex post, greater transparency facilitates understanding the under- or over-performance relative to the market.

IV.A.3: Rebalanced Weighted Benchmark

This is the final stage under the responsibility of the Ministry of Finance. Rebalancing enables an investor to maintain optimal expected utility in an environment where returns are not predictable. A rebalancing premium accrues to an investor rebalancing to fixed portfolio positions. In certain cases, the rebalancing premium allows expected wealth to increase faster than the case when no rebalancing occurs because a dynamic portfolio’s variance is reduced by regular rebalancing. We recommend that the difference in returns between a fully passive portfolio, which is not rebalanced, and the rebalanced portfolio be reported (similar to Figures 1 and 2).

One of the Fund’s major developed comparative advantages is its discipline in following a rebalancing rule. The benefits of rebalancing can be made even clearer by separately showing the value-added to rebalancing. Indeed, rebalancing does not always add value—rebalancing can underperform passive buy-and-hold portfolios for several years, but the rebalancing premium will be positive over a long enough horizon. Transparency of the gains or losses due to rebalancing can allow a long-term investor to better anticipate short-term losses.

Like the reporting of the optimal strategic weights vs. market weights, the rebalancing gains should be reported over long samples. We recommend detailed examinations at the adoption or revision of the Reference Portfolio (see below), or at similar infrequent intervals. Harvard, the largest university endowment, compares the returns of its fund relative to a 60%

⁵¹ This is similar to the behavior of retail traders who tend to trade too often, and lose money. See Barber and Odean (2000).

equity/40% bond mix, representing a typical institutional investor.⁵² It also compares its policy portfolio returns, which are the counterpart to the equity-bond benchmark of the Fund, to the 60% equity/40% bond portfolio. Yale compares its performance to other endowments in terms of asset allocation and value-added within each asset class.⁵³ We believe it is inappropriate to compare the Fund's returns to other sovereign wealth funds or large pension funds because of different aims, liabilities, and governance structures. It is important, however, to compare the Fund to the average investor—represented by passive, non-rebalanced market weights—and in doing so anchor the Fund's rationale for taking on strategies that are different from average.

IV.A.4: Transparency of NBIM's Active Management Activities

We now turn to the investment processes within the fund manager. The rebalanced, diversified benchmark (with non-market capitalization weights) set by the Ministry of Finance is not the portfolio directly implemented by NBIM. The Ministry of Finance sets a set of risk constraints, which includes a tracking error limit and the set of investable securities. NBIM takes responsibility for deviating from the benchmark staying within the defined risk parameters.

There are two important differences between the “active” investment decisions made by the Ministry of Finance and NBIM. The first is verification horizon and the frequency at which decisions are revised. Strategic benchmarks are adopted with horizons stretching over decades and should be revised infrequently. The second difference is cost. An ideal benchmark able to be implemented at zero cost serves as a realistic alternative to relatively expensive delegated active fund management (see Section I.C).

Risk—in the form of deviations from benchmark—should be taken because the fund manager believes it can add value and cover its costs. As Section II reviews, some of this risk is taken to harvest factor risk premiums which are not contained in the Fund's benchmark. Other value-added comes from security selection. Both of these areas have shorter verification horizons than the Fund's benchmark; horizons for most (scalable) factor strategies are five to ten years and horizons for most security selection programs are measured in quarters (or at even shorter intervals).

Both factor and security selection strategies are costly. All of the factors covered in Section II require skill to design and trade. Because of the Fund's scale, even harvesting a small cap premium cannot be done at zero cost. All of the dynamic factors we consider, with the exception of SmLg, do not have published benchmarks with all the ideal criteria laid out in Section I.C. It is possible that because of the required higher turnover of factor strategies, none can be implemented as cheaply (at essentially zero cost) as classic market capitalization index

⁵² See Harvard Management Company Endowment Report 2013.

⁵³ Yale Endowment Update 2012.

strategies. (To our knowledge, there are currently no third-party factor strategies that cost below five basis points.) If they do become viable alternatives to using an active fund manager, then the Ministry of Finance might consider placing them in the Fund's strategic benchmark, assuming they meet the objectives of the Fund.

Our recommendation is that the formal mandate of transparency be stated in broad terms. The fund manager should be required to articulate its principles and broad sources for active management, and report its ex-ante risk budget and ex-post returns emanating from each source. In practice, we suggest that the risk and returns of the active return, or deviation from benchmark, be broken down into the ORP and the remainder, security selection. We also recommend that the different components of risk and returns within the ORP be reported.

IV.A.5: Operational Reference Portfolio

We recommend that returns, and risk budgets, of two main components of the ORP be reported (we reviewed the ORP in Section II.F):

1. Diversifying beyond the benchmark given by the Ministry of Finance, and
2. Taking on exposure to systematic factor risk.

The verification horizons for these strategies can stretch up to a decade, as discussed in Sections I and II. Separately reporting these sources of returns allows them to be appropriately benchmarked, and is consistent with the principle that different sources of risk should be measured appropriately over different horizons. Breaking out the returns to these components reinforces that these sources add value over time to the Fund's stakeholders. But, there may be short-term losses—in fact, for factor risk premiums, superior returns over the long run represent compensation for bearing potential losses which other investors abhor. By making transparent these further diversification and factor breakdowns, expectations of risk can be better managed. In addition, both components are costly to implement.

Improved Diversification

Because the choice of a market diversified benchmark is arbitrary, component (1) can add value by enlarging the investment universe beyond that which index providers choose. The Fund's long horizon and tolerance for illiquidity allow it to take more illiquid positions than the average investor. Thus, the Fund need not be bound by liquidity and investability constraints in the public indexes. However, the Fund's large size means that these investment opportunities may be limited, and implementation costs may outweigh the benefits of increased market completion. Whereas the public indexes are designed to facilitate rebalancing at low cost, very illiquid positions may be problematic to rebalance. More illiquid securities may suffer losses

when liquidity evaporates in markets. Consequently, enlarging the investment universe beyond the benchmark is not without risk—it is a costly, active decision.

Factor Risk

A large academic literature shows that some of the factor tilts in component (2), like value and size for equities and duration for bond portfolios, are associated with long-run average returns in excess of market-capitalization benchmarks. These factors are chosen by NBIM; presumably it will not choose to obtain factor exposure if the costs of implementing these factors are larger than their expected risk premiums. Judicious choices of factors should lead to higher long-run returns as the factor premiums are equilibrium compensation for bearing risk. NBIM should continue its practice of publishing white papers justifying the risk factor exposures taken. This provides a solid basis for investing in these factors.

We believe that the Ministry of Finance should not mandate the amount of risk devoted to taking factor exposures or determine which factor risk premiums should be harvested. The amount of factor risk taken should be a function of the degree of conviction of the manager in staying the course—especially during times when factor exposures are responsible for generating losses relative to a market benchmark, the ability of the manager to hire talented personnel to implement the factor exposures, and the trading capabilities of the organization. Certain factor risks are correlated, and these correlations change over time (see Section II.D). Expertise in collecting one factor may not translate in the ability to successfully invest in another. During certain periods, the rewards for increasing factor risk exposures can be very attractive, and the fund manager should have flexibility in timing factor risks. Factors like liquidity can be harvested either as a separate stand-alone factor or be incorporated in the implementation of rebalancing of the overall Fund, the equity or bond portions of the Fund, or in the construction of other factors.

We recommend not adopting factor benchmarks at the level of the Ministry of Finance because the Fund is very large, and optimal implementation of factor strategies involves more trading and execution skills than market-capitalization based index funds. Without expertise, much of the factor risk premiums can potentially be eaten up by transaction costs. Our opinion is that the factor choices should be the purview of NBIM and not the Ministry of Finance until such time as factor index funds become essentially zero-cost, viable alternatives for large-scale active management.

At the same time, rebalancing factors will add value over the long run, just as rebalancing equity and bond positions in the benchmark portfolio. Separately reporting NBIM's risk budget devoted to factors can help the Fund's stakeholders understand why factor risks are being taken: the factor risk exposures should be rewarded in the long run. Public support for factor

strategies is most needed after factor strategies have experienced losses, and transparency helps the public to understand when and how these losses can occur. Simply put, transparency, together with proper management of expectations, helps the Fund to stay the course. It is during these usually volatile, bad times that the factor investing opportunities are most attractive. In fact, if the motivation for taking factor strategies is properly anchored in the public domain, astute commentators can confirm that the Fund is acting counter-cyclically (“buying low and selling high”) if the factor components are separately reported. Transparency can help the Fund not to disinvest at the worst times and not to chase returns by loading up on factor risks after a lucky run of good returns. At the minimum, knowledge of the risk budget devoted to factor strategies, combined with some estimates of the possible distribution of losses of these factor risks, facilitates risk management.

Transparency of the amount of risk devoted to factor investing helps NBIM not to be unfairly penalized for poor performance stemming from taking on factor risks, when these factor risk exposures will be rewarded over long horizons. Breaking out the factor returns is especially advantageous for the current mandate of taking into account the fiscal strength of countries in managing the Fund’s sovereign bond portfolio. Implementing this mandate requires taking deviations from the benchmark, since the benchmark’s weights are not computed with this consideration in mind. This is a mandated move away from market weights, but it may not be associated with long-run excess returns. If NBIM were to report a separate breakdown of each factor exposure, in terms of both realized returns and risk budgets, then the separate contributions of all factors can be measured. The tilts away from market weights can either be chosen at the level of the Ministry of Finance, as is the case with the fiscal strength consideration for sovereign bonds, or at the level of NBIM, like the size and value factors. This approach can also be extended to a socially responsible investment mandate, or a mandate involving environmental considerations. Thus, in Figure 14, we separate factor exposure into dynamic factors, which theory and empirical evidence suggest add value over time, and other factors which may not. The latter include the fiscal strength mandate and environmental, social and governance (ESG) criteria.

Not Separately Reporting Index Replication Costs

Our assessment of the ORP in Section II.F identifies that NBIM also uses the ORP to implement smart rebalancing of the benchmark. The additional diversification and factor exposure components are risk-return enhancing activities and have relatively long horizons for verification stretching over years. The goal of smart rebalancing is cost minimization, and its horizon is short: over days to months. Thus, we do not view smart rebalancing in the same category as the other value-enhancing elements of the ORP.

We recommend that NBIM consider reporting its smart rebalancing activities separately from the ORP. We are not saying that transaction costs are not important—quite the contrary. At present, the cash flows are still large relative to the size of the Fund and rebalancing can be effected mostly through inflows. As the Fund becomes larger, rebalancing may need to be accomplished by active selling and buying of the portfolio in addition to inflows. Liquidity management and minimizing transaction costs will require only more skill as the Fund increases in size. The choice of the benchmark, however, assumes it can be replicated and rebalanced at close to zero cost. While there are costs to tracking the index, there are also benefits to holding a large portfolio. Balance sheet utilization activities like securities lending and collateral enhancement offset transaction costs. These activities, while designed to provide value for the Fund, flow directly from the decisions of the choice of benchmark and the rebalancing policy. One possibility would be to aggregate the results of benchmark implementation of cash flows and smart rebalancing with fund balance sheet utilization activities. These are all short horizon activities, require trading and other skills within NBIM, and explicitly capture elements of cost and return that are often unaccounted for by other funds.

Granular Reporting by Large Funds

Several large funds report breakdowns of risk budgets and return in terms of the sources of risk taken in their active management activities. CPPIB, for example, translates all of its asset allocation positions into risk exposure positions in terms of the constituent factors in its Reference Portfolio.⁵⁴ It also breaks down the total fund risk into active and Reference Portfolio components, and into “alpha” and “beta” components. ATP reports risk statistics, measured by potential tail losses, in five risk factor classes: interest rates, credit, equities, inflation, and commodities.⁵⁵ It also reports breakdowns on its risk allocations and factor exposures.

NBIM already reports some risk breakdowns stated in terms of expected relative volatilities in its annual reports: for equities, fixed income, and for the Fund.⁵⁶ It also reports some statistical properties of the ORP. The increased transparency we advocate is not far from its current practice, and we believe NBIM is well positioned to provide more detail on its risk budgets and realized returns of its active management strategy.

IV.A.6: Security Selection

The final stage involves how NBIM uses fundamental analysis to select securities. The current active return, defined as the return of the Fund minus its benchmark, does not distinguish between the value-added from security selection and the factor risk exposure in the ORP. While

⁵⁴ See CPPIB Annual Report 2013.

⁵⁵ The ATP Group Annual Report 2012.

⁵⁶ See NBIM Annual Report 2012.

the ORP is also active, from the viewpoint that it requires skill to implement and is not costless, the strategies taken in the ORP are rules-based, or systematic, while the process of security selection involves more discretion. Security selection strategies typically have a shorter holding period horizon than factor strategies in the ORP.

In addition, uncoupling the active return isolates the “true active return,” which is the return earned by NBIM in excess of the factor risks, the costs of rebalancing the portfolio, and the decisions taken to enlarge the Fund’s universe beyond the benchmark. This decomposition also raises the standard for active management.

IV.B: Adopt the “Opportunity Cost Model” for Active Management

The Opportunity Cost Model provides an internally consistent benchmark for active management which is straightforward to encompass many alternative asset classes.

We recommend that the Fund adopt the Opportunity Cost Model for active management, similar to the models adopted by CPPIB, GIC, and other funds.

IV.B.1: Why the Opportunity Cost Model?

The primary goal of the delegation model is to improve the Fund’s risk-return trade-off in a way that cannot be obtained by setting strategic benchmarks. The Opportunity Cost Model does this in several ways:

- Through the Total Portfolio Approach, active management is always evaluated against the opportunity cost of liquid, public markets as represented in the Reference Portfolio. Thus, the focus of active management becomes the component of returns that cannot be obtained in the Reference Portfolio benchmark.
- The Opportunity Cost Model plays to both the natural and developed comparative advantages of the Fund. A long horizon is necessary because there will be significant differences—even over 10 year intervals—between a private market investment and the mix of the public market equity and bonds in the Reference Portfolio representing the opportunity cost of the private investment. A developed advantage of skill and professionalism is required to both find attractive investments that are not spanned by the Reference Portfolio and compute the mix of equity and bonds to fund them.
- In deviating from the Reference Portfolio, active management must cover its own costs. The Opportunity Cost Model raises the bar and accountability for active management because since the Reference Portfolio can be passively implemented at close to zero

cost, the difference in returns between the fund and the Reference Portfolio directly represents the value-added from active management. In other delegation models relying on non-tradable benchmarks for private asset classes set by the asset owner, the difference between the fund and benchmark returns does not directly measure the worth of active management.

- The Opportunity Cost Model is scalable. It is straightforward to apply to many types of illiquid asset markets. As the Fund increases in size, it may consider expanding the universe of investments. The Opportunity Cost Model can be used over many different alternative investment classes. When implemented in its full form, the framework provides the answers to if and when a particular new alternative investment class should be introduced.

We now comment on the two core elements of this model, namely the Reference Portfolio and the Total Portfolio Approach, as described in Section III.C.

IV.B.2: Reference Portfolio

The Fund should consider adopting a rebalanced weighted benchmark comprising only tradable securities in liquid markets: a benchmark of equities and bonds similar to the Fund's previous benchmark that did not include real estate. At the same time, NBIM should be allowed to make its own decisions on how best to beat this passive benchmark. NBIM can choose liquid, public investments or investments in illiquid, private markets.

In the context of Figure 14, the rebalanced weighted benchmark constitutes a passive portfolio which can be cheaply implemented in index funds. This benchmark is entirely owned by the Ministry of Finance and the Storting. If the Fund were to adopt the Opportunity Cost Model, then an appropriate Reference Portfolio might consist of 60% equities and 40% bonds, using the same equity and bond benchmarks that are currently set by the Ministry of Finance. This is the same as the Fund's strategic benchmark before the introduction of real estate.

IV.B.3: Total Portfolio Approach

When NBIM deviates from the Reference Portfolio benchmark, it must undertake a cost-benefit analysis. In pursuing active management—whether in public or private markets—it must recoup its costs and obtain a higher return than the Reference Portfolio. Specifically, we recommend that NBIM adopt the Total Portfolio Approach described in Section III.B, and under this delegation model there will be no need for the Ministry of Finance to set a specific real estate benchmark. Under the Total Portfolio Approach, all investments taken by NBIM away from the Reference Portfolio are internally benchmarked and funded. NBIM would be free to pursue investments in direct real estate, for example, but it has to believe those investments

offer a superior risk-return trade-off than publicly traded real estate (REITS, which are in the Fund's benchmark) or some other mix of bonds and stocks that would be used as a funding benchmark. In fact, exposure to real estate can be made through regular equities and bonds, REITs, direct real estate, debt related to real estate deals (like mezzanine financing), or agency bond issues which are tied to real estate markets (like mortgage-backed securities issued by Fannie Mae). All of these trade-offs would be evaluated internally within NBIM. Similarly, if the Ministry of Finance allowed such investments in the Fund's investable universe, NBIM would decide if a direct private equity investment would be superior to a levered position in public equity. In adopting the Opportunity Cost Model, responsibility for making these trade-offs lies with NBIM—the party best informed to make these investment decisions.

IV.B.4: Advantages in Implementation

A major advantage of the Opportunity Cost Model is that it will include real estate risk in the Fund's deviations from benchmark (traditional "tracking error"), which is currently not the case. While the amount of direct real estate is small, the Fund's holdings of real estate and other illiquid assets may grow over time. Under the Opportunity Cost Model, benchmark deviations encompass all private, illiquid investments, including real estate. It is in fact ideal to implement the Opportunity Cost Model while investments in real estate are still small, so that there are minimal adjustments in changing the benchmarking of real estate to public market securities. The experience in applying the Total Portfolio Approach to real estate will be valuable if the Fund moves into other alternative asset markets.

A second advantage is that NBIM is responsible for creating its own internal benchmarks. The Ministry of Finance is responsible for the Reference Portfolio, which involves the mix of equity and bond passive indices and the rebalancing decision. Deviating from the Reference Portfolio is entirely NBIM's decision and the cost of any such deviation is borne by NBIM. This simplicity and the clear demarcation of responsibilities should help create a more robust governance structure.

Third, benchmarks commonly used in real estate and other private, illiquid markets are not tradable. This is true for the tailored International Property Databank (IPD) which is the Fund's real estate benchmark. These reported returns do not represent actual returns achievable by any investor as they pool the experiences of many different institutions. Idiosyncratic risk faced by an individual investor is much higher than the reported illiquid asset index because an individual cannot obtain the same degree of diversification. Unlike public markets, there is no tradable market IPD portfolio. In the Opportunity Cost Model, the Reference Portfolio does represent an implementable return. The real estate investments, and other investments in other illiquid asset markets, are benchmarked against realizable returns. If the Fund moves into private equity or infrastructure, there is no need for the Storting or Ministry of Finance to

evaluate benchmarks for these asset classes—which have similar disadvantages as the IPD for real estate. It might still be desirable to use a benchmark such as IPD in an alternate way—to assess the skill of NBIM’s real estate investment team, for example, and for the Ministry of Finance to assess different risk scenarios—but the funding benchmark would be used to evaluate returns relative to the Reference Portfolio.

An alternative to the Opportunity Cost Model is an absolute mandate, like a real return plus a spread. This class of mandates also includes a fixed level, like 4%. We do not recommend the Fund take on an absolute mandate. Absolute mandates cannot be replicated by the fund manager or the asset owner, and thus there is no way to guarantee the absolute mandate by passive or active management. Consequently, it is almost certain that an absolute target cannot be delivered by the fund manager. If the returns are below target, then disappointment can lead to dissatisfaction with active management. If returns are above target, then spending may be reset at levels that are unsustainable. The Opportunity Cost Model represents a demanding, but implementable, alternative for active management.

IV.B.5: Challenges in Implementation

In considering the Opportunity Cost Model, we recommend that the Fund be mindful of:

1. Reference portfolio composition

The Ministry of Finance would be responsible for the most important decision in the Opportunity Cost Model—setting the Reference Portfolio. This represents the index, viable alternative to active management that has a high probability of meeting the goals of the Fund. In the context of Norway, this includes the sustainability of the spending rule, projections on the oil revenues, expectations on the risk premiums and distributions of the factors in the Reference Portfolio, how these factors might optimally differ from market capitalization weights, the rebalancing rule, and which factors can be passively harvested at close to zero cost to benchmark active management. While the Ministry of Finance undertakes many of these calculations now, its role would be enhanced in anchoring the Reference Portfolio strategy and communicating it. The Ministry of Finance would also need to analyze anticipated deviations from the Reference Portfolio undertaken by NBIM’s active management activities.

When the Opportunity Cost Model is implemented, the Ministry would need sufficient resources to ensure that it, or the governing body of the delegated investment manager, could evaluate and monitor the implementation of the Total Portfolio Approach. The Reference Portfolio and the Total Portfolio Approach are challenging to implement and are much more time and data-intensive processes than traditional asset allocation methods.

2. Specifying limits on active risk

This is similar to the current benchmark deviation volatility and the same considerations apply. As Section II notes, the volatility limit should be accompanied by other risk management tools. We recommend that the risk taking of active management be increased in adopting the Opportunity Cost Model, as we detail below, but be accompanied by higher moment risk limits.

The risk calculations become more difficult under the Opportunity Cost Model because specific asset-class benchmarks are not adopted. Various benchmarks for alternative asset classes, like the IPD currently used for real estate, can be used to estimate benchmark deviations both at the asset owner level and within NBIM. As Section III.C discusses, there can be large deviations between private market investments and the funding portfolio of public market assets in the Reference Portfolio. These differences may extend to 7-10 years. (Appendix C, for example, reports lags of several quarters for real estate investments.) Incorporating such deviations into the risk management process is challenging. Breaking out the different components of active risk taking, as we recommend in Section IV.A, facilitates these calculations.

3. Encouraging risk taking

In the traditional asset allocation model, there is a fixed allocation to real estate. In the Opportunity Cost Model, there are no fixed asset class allocations. One potential tendency in the Opportunity Cost Model is that the fund manager may shy away from holding real estate, which is detrimental for the asset owner because of its diversification benefits, because such an investment requires deviating from the benchmark.

We believe this bias goes in the asset owner's favor. The Reference Portfolio is sufficient to meet the Fund's objectives with a reasonable probability given current assumptions on inflows and spending. The Reference Portfolio captures the opportunity cost of any active deviation from benchmark in both private and public markets. Deviating from the Reference Portfolio should be done only to gain compensated risk. If the fund manager cannot find such opportunities, and cover the costs of the investments, then it should not be taking on the alternative investment.

A more relevant concern is how to align the long horizon of the asset owner with the fund manager in the Opportunity Cost Model—which requires a long horizon in order to implement well. Breaking down the various active sources of risk plays a role (see Section IV.A), as making transparent these benchmark deviations shows the risk devoted to different active management activities. When the verification horizons of

these sources of risk are communicated, the various sources of risk can be appropriately benchmarked. Another approach is to use moving averages in reporting benchmark deviations, which smooth out the effects of temporary price movements. The most important element is long-term commitment by the asset owner and the board of the investment manager to the Opportunity Cost Model (we comment further below).

4. Gradually increasing the Fund's universe of permissible holdings

In the form of the Opportunity Cost Model adopted by CPPIB, GIC, and others, fund managers face few or no restrictions on the assets that can be held. The Fund's investments are currently subject to restrictions, some of which reflect ethical and social considerations. Just as the Fund gradually moved from first all bonds, to equities, more equities, and then to direct real estate, the range of allowable investments in other illiquid asset markets can be increased over time. The optimal implementation would involve NBIM being responsible for the choices of all its assets in an unrestricted asset class universe. Reflecting the Fund's current approach to ethical and socially responsible investing, the fund manager might be restricted from choosing specific securities within each asset class associated with issuers that fail to meet certain criteria.

The full implementation of the Opportunity Cost Model does not restrict the maximum investments made in each asset class, except that the overall risk limits are satisfied. In an extreme case, this might mean that the Fund could hold very large real estate positions, much larger than the current 5% limit, if NBIM felt that the real estate market offered more attractive opportunities than equities and bonds in the Reference Portfolio. We believe that these decisions are best left to the fund manager, and the Ministry of Finance should concentrate on employing risk limits at the overall Fund level.

We recommend a gradual transition to the Opportunity Cost Model. In the transition process, the Fund may consider a temporary hybrid arrangement where the Fund's investment universe is specified and there are maximum limits on different asset classes. These maximum limits would gradually be raised.

While we did not review the Fund's experience in real estate given the short history and lack of data, we believe that as the Fund's investments in real estate are currently small, it is an ideal time to incorporate the real estate investments into the Opportunity Cost Model. It is reasonable that experience in the real estate program, with appropriate funding mixes from the Reference Portfolio, would need to be evaluated before the Fund applies the model to other asset classes. A successful record of real estate investments under the Opportunity Cost Model augurs well for the Fund to apply the

framework to other alternative asset classes. The Fund's real estate experience would need to be separately reviewed.

5. Having a dedicated fund manager invested in public and private markets

Implementing the Opportunity Cost Model requires a dedicated fund manager with full access to all markets. Indeed, at the heart of the framework is a fund manager who is able to transfer risk-taking ability from public markets, where markets are largely efficient and there are few information asymmetries, to private markets, where there are potentially greater investment opportunities for a skilled manager who has access to superior information, and superior capabilities of processing and acting on that information, than the average investor. The Reference Portfolio ensures that illiquid investments carry appropriate risk benchmarks derived from liquid, public markets. Thus, having a dedicated fund manager who has experience in all asset classes is a precondition for adopting the Opportunity Cost Model.

6. Maintaining an investment mandate

Central to the Opportunity Cost Model is a demanding rebalanced and diversified benchmark: the Reference Portfolio. Any benchmark deviation taken by the fund manager is measured relative to the Reference Portfolio, which by assumption meets the Fund's objectives and can be implemented at close to zero cost. The fund manager chooses when to take deviations from the benchmark, and does so when it is compensated by positive risk-adjusted gains relative to the Reference Portfolio. Imposing constraints, especially those involving directed investments in areas that are not taken on the basis of enhancing the Fund's risk-return trade-off, handicaps the fund manager's ability to outperform the Reference Portfolio. It also penalizes asset owners because the constraint makes it less likely the Reference Portfolio can satisfy the aims and goals of the Fund.

7. Enhancing the role of the ORP

In practice, benchmarking of investments would not take place directly with the Reference Portfolio but with an internal benchmark at NBIM. This Policy Portfolio would reflect strategic factor and asset allocation decisions, and incorporate the different decision processes in various investment programs. The ORP would likely serve as an internal Policy Portfolio under the Opportunity Cost Model. In such a role, it is helpful to differentiate the risk budgets between active investment decisions, like the factor exposures, from those involved in implementing the Reference Portfolio, like rebalancing (see Section IV.A).

The composition and use of a Policy Portfolio should be determined by the fund manager. A Policy Portfolio is, strictly speaking, not a requirement within the Opportunity Cost Model. GIC's Policy Portfolio reflects longstanding commitments to active programs such as private equity and real estate which were long established before the adoption of its Reference Portfolio. Given the still nascent nature of NBIM's real estate program, there is not a similar need or rationale for the Fund to enshrine an allocation to real estate within a Policy Portfolio.

We believe the ORP could serve as an internal Policy Portfolio under the Opportunity Cost Model. One benefit of adopting the ORP as the Policy Portfolio is that it would fix a desired level of various factor exposures within the Fund. Using the ORP this way presents an opportunity for the Fund to go beyond the incarnations of the Opportunity Cost Model practiced by CPPIB and GIC. Since the ORP incorporates factor risk, the opportunity cost benchmarks for a private market investment could be funded not only by simple market-weighted equity and bond positions but also by tilting those equity and bond positions toward relevant factor risk exposures. This would be relevant, for example, in dealing with private equity investments which often have negative exposures to a value factor and positive exposures to a size factor (see Phalippou, 2010). Effectively, the factor exposure impact of private investments on the Fund would then be offset through its public market holdings to maintain the desired fixed exposures.

8. Ensuring internal competence

In the Opportunity Cost Model, taking active risk relative to the Reference Portfolio is predicated on the ability of the fund manager to hire qualified people, build state-of-the-art systems, and to efficiently collect and analyze information—in other words, competent and world-class developed advantages (see Section III.A). Not being tasked with evaluating the internal governance structure of NBIM, we merely note that the funds which have adopted the Opportunity Cost Model all have strong governance structures with these qualifications.

IV.C: Increase Risk Taking of Active Management

The active management risk-taking limits should be increased.

Benchmark deviation limits cap how far the fund manager can deviate from the Reference Portfolio to seek additional returns. In the delegation framework reviewed in Section III.B, the optimal amount of delegation depends on the cost of active management, governance

structure, the capabilities of the parties, and the details of the decision processes between the parties. In terms of the Opportunity Cost Model, delegation permits the fund manager to find return opportunities which are not available in the fund's Reference Portfolio benchmark. However, because the Reference Portfolio is designed to meet the fund's objectives with a reasonable probability, there is little reason to actually take unrewarded risk beyond it.⁵⁷ Delegation is risky—active management may result in lost value relative to the Reference Portfolio—and thus it is appropriate to place constraints on benchmark deviations.

IV.C.1: Upside and Downside Limits on Benchmark Deviations

The most common form of limit on benchmark deviation is imposed on the range of possible outcomes. "Tracking error" is designed to measure the typical size of these deviations, and does so in terms of standard deviation. Since standard deviation is only one measure of the distribution of active risk, the Fund might consider adopting other risk constraints characterizing higher moments. Standard deviation treats positive and negative deviations from benchmark symmetrically, whereas investors typically have asymmetric preferences over positive and negative outcomes.⁵⁸ Investors care about the whole *distribution* of benchmark deviations and are especially concerned with downside losses. Thus, we recommend that the Fund adopt limits on tail risk.

A formal constraint on tail risk might take the form of a "Value-at-Risk" constraint or some similar left-hand tail measure. Some methods for computing this limit were covered in our analysis of the Fund's benchmark deviation volatility in Section II.G. An important consideration is that our calculations were backward looking under a regime of low risk-taking limits, and the risk limit to be adopted is a forward-looking measure that depends on whether the Opportunity Cost Model would be implemented, the factor exposures taken by the fund manager, the risk preferences of the ultimate owners of the Fund, and the Fund's governance structure. NBIM would presumably give guidance on these issues. The Ministry of Finance would need to do a formal evaluation to set this criterion taking into account the risks of the factors harvested by the Fund, the breakdowns of those risks in the ORP, the risk of security selection, and the total amount of benchmark deviation risk finally adopted, among other considerations. Such a study should also include analysis on any potential conflict between the standard "tracking error" limit currently used and the new tail risk constraints; one constraint might be binding more regularly rendering the other less useful, and since all risk limits are hard to measure *ex ante* and have different consequences *ex post*, the fund manager may implicitly weight one risk

⁵⁷ When the investor starts with a non-optimal benchmark, then imposing a tracking error mandate around that benchmark results in a portfolio inferior to the optimal benchmark (see Roll (1992)). Section I.C points out that an ideal benchmark delivers an optimal set of exposures to factors in liquid markets.

⁵⁸ A popular class of utility functions capturing this phenomenon is loss aversion, see Kahneman and Tversky (1979).

measure more than another in pursuing its active strategies. Our recommendation on increased transparency of risk taking mitigates some of these concerns (see Section IV.A).

One possibility is that Norges Bank be required to provide guidelines on the minimum underperformance of the Fund relative to the Reference Portfolio that would be expected in a given 12-month period with 90% probability. The Ministry of Finance would set a formal limit on left-hand tail risk of active returns in terms of “Value-at-Risk,” which is the limit where active return underperformance would be expected to be exceeded in one of ten years, on average. In terms of the language of “Value-at-Risk” used in risk management, this limit is a 10% Relative Value-at-Risk computed at the annual frequency. At many trading institutions, these Value-at-Risk limits are stated with shorter time intervals, often daily, with more infrequent occurrences, often 1% or 5% probabilities. Given that many strategies appropriate for the Fund, particularly for factor risk, operate at low frequencies, we believe the time interval should be at least annual. We recommend that the probability be 10% to correspond to business-cycle frequencies, rather than further out in the left-hand tail. Value-at-Risk criteria for smaller left-hand tail probabilities and at higher frequencies are more appropriate for trading strategies within NBIM rather than for adoption at the level of the Ministry of Finance. We reiterate that a specific left-hand tail limit is beyond the analysis of this report, but our calculations in Section II.G, with the caveats mentioned above and also noting that the analysis in that section applied to monthly and not annual active returns, indicate that a very rough risk limit could be that the active loss might be expected to be below -6% in one year out of every ten years, on average. Various Value-at-Risk limits have been adopted by the Fund’s peers. CPPIB, for example, sets a 90% Value-at-Risk measure over a one-year horizon, in line with our recommendation.

We recommend that the Fund also adopt tail risk measures on the upside. This is less commonly done by financial institutions. Academic theory suggests that large deviations on the upside are a cause for concern just as are large deviations on the downside.⁵⁹ Unusual positive returns may indicate that an investment strategy is not sustainable, or that there are risks being taken on that are not fully understood which may cause large unanticipated losses in the future. Indicating that a large, positive return is very unusual helps manage stakeholder expectations so that there are fewer demands to increase payouts which cannot be maintained in the long run. A formal mandate might specify that Norges Bank give guidance on the benchmark deviations on both the left- and right-hand tails, but that only a formal left-hand tail risk limit be assigned by the Ministry of Finance.

⁵⁹ See, for example, Admati and Pfleiderer (1997).

IV.C.2: Increased Volatility Limit of Benchmark Deviations

Selection of targets limiting the deviations from benchmark taken by active management is ultimately a question of the Fund's risk appetite, and is also influenced by stakeholders' experience in bearing losses and success in adhering to a rebalancing policy.⁶⁰ We recommend that the Fund take on more risk than its current regime, which is a 1.0% tracking error limit per year. A modest increase might be to raise the active risk taking of the Fund to correspond to a benchmark deviation volatility of 1.75% per year. The small increase is in line with the tradition of the Fund gradually adopting changes in investment strategies. It is only 0.25% higher, computed using the superior calculations of ex-ante risk adopted since the financial crisis, than the Fund's risk budget of active management prior to the financial crisis. A specific risk limit increase would need to be determined taking into account the form of the Opportunity Cost Model being adopted and the specific superior diversification and factor risk strategies being pursued in the ORP.

The analysis in Section II points out that benchmark returns constitute over 99% of the variance of total Fund returns, so active management represents a tiny fraction of the Fund's overall performance. Our recommendation pushes up this risk budget from tiny to very small. Typical tracking error mandates for actively managed fixed income portfolios are over 1% and range between 2% and 6% for equity portfolios.⁶¹ Tracking errors below 2% are typically regarded as enhanced passive, or pure passive, portfolios. Petajisto (2013) constructs a comprehensive sample of all-equity mutual funds, including both index and active funds. He finds that 94% of funds have tracking errors above 2% and 68% have tracking errors above 6%. Clearly a modest increase in the Fund's benchmark deviation volatility to 1.75% leaves it well on the lower end of typical active risk mandates. By some measures, it will still leave the Fund as an "enhanced index" fund.

As important as the total amount of active risk taken is where that active risk is being employed. The fund manager should give broad descriptions on how it seeks to generate value relative to the Reference Portfolio and the risk budgets devoted to each area. Thus, our recommendations on increased transparency of risk go hand in hand with the limits on risk (see Section IV.A).

⁶⁰ Malmendier and Nagel (2011) show that investors who lived through the Great Depression became more risk averse compared to younger investors who did not experience such large losses.

⁶¹ These ranges are from the Bank of International Settlements 2003 report, "Institutional Asset Managers: Industry Trends, Incentives and Implications for Market Efficiency." Vardharaj, Fabozzi and Jones (2004) report that typical active tracking error mandates can exceed 10%. Zephyr, a performance evaluation company, reports that typical active mandates have tracking errors of 4-7% (see <http://www.styleadvisor.com/resources/statfacts/tracking-error>).

IV.C.3: Why Increase Active Risk?

There are several reasons why the Fund should increase its active risk:

1. Improved diversification adds long-term value

As Section I describes, the foundational theories in finance suggest that diversification improves risk-return trade-offs. Active management can achieve greater diversification than available through public indexes. This improved diversification comes at the cost of deviating from the benchmark. Our calculations in Section II.G suggest that, in practice, it is reasonable to expect significant realized benchmark deviation volatility by diversifying beyond a well-established world index. If some of the increased risk taking is devoted to this purpose, as stated in the aims of the ORP, then this will add value to the Fund in the long run.

2. Risk taking dedicated to factor exposure adds long-term value

Section I outlines how a long-term investor can add value by choosing a diversified portfolio, rebalancing, taking on factor risk, and undertaking security selection or factor timing. The first two of these are captured in the Reference Portfolio. Return deviations relative to a rebalanced, passive benchmark based on market-capitalization weights do not differentiate between factor risk and security selection.

The academic literature suggests that factor risk premiums—like the value and size premiums currently being pursued by NBIM—add value over time. Factor risk premiums are included in the ORP. They are not included in the Fund’s strategic benchmark. There is opportunity for the Fund to obtain superior risk-adjusted returns than its passive traditional benchmarks if the Fund takes on dynamic factor risk. That is, if more risk is delegated to NBIM and that risk is taken in factor positions, then this should result in a higher risk-adjusted return for the Fund.

3. Sufficiently large prudent risk buffer

Although the current tracking error limit is understood to apply during normal times, and occasionally risk may rise above the limit during very volatile periods, NBIM has historically not taken its full allotment of risk. We are not privy to the exact determination of NBIM’s optimal allocation of risk, but we speculate that it allows a buffer due to the reputational risks of breaching the limit. Some of the benchmark deviations are employed by NBIM in activities that do not strictly represent active investment strategies. For example, rebalancing on dates other than the prescribed rebalancing dates when indexes are reconstituted requires deviating from benchmark,

and having an extremely low risk budget with regards to these activities would entail large transaction costs. Increasing the Fund's risk taking limit recognizes these facts.

4. Positive historical record of active management

The Fund's experience with active management is overall positive, and statistically significant, over its history, as we review in Section II. Active returns did not detract value over the whole financial crisis, and during that time the Fund rebalanced when many other investors failed to. Active returns were negative over the initial months of the financial crisis (see Section II.C and Ang, Goetzmann and Schaefer, 2009), but much of these losses were attributable to large exposures to factor risks. Much of the rebound in active returns from the worst months of the financial crisis can also be attributed to factor exposures. The Fund emerged from this period more robust—in terms of managing through a turbulent period as well as not changing the overall framework of its governance structure. If NBIM were able to achieve a monthly Reward-to-Risk ratio of 0.10 to 0.15 on an annual 1.75% benchmark deviation volatility, which we believe to be realistic in light of the active returns of the Fund not including the financial crisis and discounting the recovery from the crisis, the manager would cover its costs and deliver a meaningful amount of net NOK or currency basket value-added.

We believe that the Fund's history of successful active management is an important consideration in whether the Fund's risk taking should be scaled up. The primary benefit of an active program is to deliver value-added. Increasing the risk capacity should, based on the experience to date, enable the Fund to significantly increase its value. Put simply, the size of the active mandate can be scaled up so that it becomes more material. But with increasing materiality comes a greater need to measure the value delivered by active management more precisely. Combining a larger active risk budget with the Opportunity Cost Model, which compares the value earned by active management to a benchmark that can be implemented passively at close to zero cost, can help in this regard.

5. Limits on downside risks

While we recommend that overall risk taking of the Fund should be increased, it would be accompanied by risk limits on the risks that really matter—downside risk relative to the benchmark.

6. More transparency of the active return components

With greater transparency of each component of the investment process, the value added by true active management—the returns added by security selection in excess of

factor exposures—can be directly measured. This makes clear the role of pure active management.

We do not believe that the Ministry of Finance should specify detailed risk budgets for the ORP, within the ORP, or the return in excess of the ORP. The Ministry of Finance does not have the same information set as NBIM, and circumstances can lead to internal risk allocations between the ORP and the actual portfolio to change over time. The importance of the components within the ORP—further diversification and factor risks—can themselves change, especially if NBIM engages in factor timing. However, we would expect that in its periodic reports, NBIM give targeted risk budget ranges for the ORP and the ORP's components. These can be compared to the actual risk taken.

7. Real estate risk will be counted in the benchmark deviations

Real estate is currently not included in the Fund's computation of tracking error. If the Opportunity Cost Model were to be employed, then real estate risk can be aggregated and be counted as part of overall active risk. Benchmark deviation volatility should be increased to provide sufficient flexibility to undertake active risk programs after absorbing the real estate impact (see Appendix C). If the Fund moves into other illiquid asset classes, these investments will also be included in the Fund's benchmark deviations.

Any decision to deviate from public markets should rest with NBIM, and such a decision will entail a benchmark deviation. For real estate investments specifically, the fund manager can decide on publicly traded REITs, which are already in the benchmark and require taking only a small tilt away from benchmark weights, or direct private real estate, whose returns can significantly differ from public real estate returns in the short run. Because of this difference, even a modest amount of direct real estate exposure requires a reasonable amount of benchmark deviation relative to a Reference Portfolio. Indeed, in the Opportunity Cost Model, the decision to invest in real estate itself must be justified by providing superior investment opportunities compared to a portfolio of plain-vanilla equities and bonds.

8. Better measurement and management of risk

As discussed in Section III, there is a dearth of quantitative tools in the investment industry for evaluating the risks of investments in illiquid or private markets. A prerequisite of expanding to private markets beyond real estate, and even for increasing the risk devoted to real estate, should be the confirmation of state-of-the-art risk tools employed by NBIM in managing illiquid asset investments. Adoption of the Opportunity Cost Model, especially after a successful implementation of the framework to real

estate, goes some way to fulfilling this requirement. The Opportunity Cost Model also provides a coherent framework for managing concentrated risk in public market investments.

9. Gradual increase of risk taking

The Fund should not immediately move to a new regime of higher risk taking; we recommend that benchmark deviation volatility should be increased gradually. The Fund's size also probably precludes finding enough large-scale suitable investment opportunities at short notice to use up the entire increase in the risk budget, even accounting for a risk buffer. We recommend that NBIM give guidance on the timeline of expanding its risk taking.

10. Periodic review of risk taking

We recommend that the benchmark deviation volatility and tail risk limits be periodically reviewed, at least in the Fund's quadrennial reports. We recommend that an evaluation of risk taking continue to be a regular section of the annual reports by the Ministry of Finance to the Storting.

IV.D: Summary

We recommend:

1. Breaking out the returns to each component of the investment process starting from selecting a benchmark to represent a diversified portfolio: (1) taking systematic positions away from market capitalization weights, (2) rebalancing, (3) seeking greater diversification than available in the benchmark and taking on factor exposures in the Operational Reference Portfolio (ORP), and (4) security selection in excess of the ORP. In the cases of (3) and (4), risk budgets associated with each should also be reported. A breakdown of returns and risk budgets for separate factors can also be done.
2. Adopting the "Opportunity Cost Model" for active management. The Ministry of Finance selects a public equity and bond benchmark that is periodically rebalanced, which represents a passive alternative to active management. The fund manager, NBIM, is free to take deviations from that benchmark in public and private markets. As its name suggests, the Opportunity Cost Model consistently benchmarks any active investment by taking into account the investment opportunities foregone in investing in public equity and fixed income markets in the Fund's benchmark. NBIM justifies its use of

active management by covering its costs and by beating the equity and bond benchmark.

3. Increasing the Fund's risk taking. The Fund should in parallel adopt additional risk limits characterizing the tails of returns deviating from the benchmark, especially on the downside. Some of the new risk taking limits would be assigned to factor exposure in the ORP, and larger factor risk exposure should be rewarded, over the long run, with higher returns for a given level of risk. Although the Fund's permitted deviations away from benchmark will increase, there will also be downside risk limits. Greater transparency also allows a more precise measurement of the amount of risk taking in true active strategies in excess of factor exposures.

Appendix A: Empirical Robustness

We provide in this appendix two additional analyses of the active returns on the Fund. In Section A.1, we measure the residual return of the fund obtained from regressing the total returns on the appropriate benchmarks. Section A.2 presents and discusses simple correlations of the active returns with the risk factors, as opposed to the partial correlations presented and discussed in Section II.E.

A.1: Residual Return

As discussed in Section II.C, part of the active return on the Fund, as measured by the difference between the Fund and benchmark returns, can be attributed to tactical over- or under-exposure to the benchmark(s). While the Fund is prohibited from using leverage to increase its benchmark exposure, it can, for example, increase its exposure to the equity benchmark by either tilting its fixed income investments toward credit or its equity investments toward higher beta stocks. We consider such as tilt an active decision and therefore include it in the active return. In this appendix, we consider the alternative approach of adjusting the leverage of the benchmark along the lines of Ang, Goetzmann and Schaefer (2009) residual return regressions.

We define the residual return as the non-zero mean errors from a regression of the Fund's returns onto the benchmark. In this performance regression, which is presented in Table A.1, the mean of the error terms is often referred to as "alpha" and the coefficient on the Fund's benchmark is "beta." As the Fund's mandate is not stated in risk-adjusted terms, the alpha does not correspond to the value added by active management. A benchmark where the beta is not one is also not replicable (see Section I.C) because a levered exposure to the benchmark cannot be obtained without specifying an appropriate risk-free asset to short. Commonly chosen risk-free assets, like U.S. T-bills, are not the Fund's risk-free asset. The Fund is currently not able to take such explicit leverage.

At the overall Fund level, in Panel A of Table A.1, it appears this regression adjustment accounts for the entire average active return (the alpha estimate is zero with a p-value of 0.45), but it only partially explains the average active returns during the pre- and post-crisis periods. In the pre-crisis sample, the alpha was three basis points per month and post crisis it was six basis points, both statistically significant at the 5% confidence level. The resulting monthly Reward-to-Risk ratio was about 0.3 in both subsamples, which, referring back to Panel A of Table 1, is in line with that of the Fund.

Unfortunately, this Fund-level analysis is misleading. Considering the asset-class level results in Panel B, we observe that only a small fraction of fixed income and equity returns can each be explained by levered exposure to their respective benchmarks. The alpha for both fixed income

and equities is positive in all three samples, extremely so for fixed income. The monthly Reward-to-Risk ratios over the non-crisis periods range from 0.20 (equities pre crisis) to 0.39 (fixed income post crisis). These results are in line with the active returns discussion in Section II.C.

Why is it that the active returns at the Fund level have a loading on the benchmark that is significantly above one, with a beta of 1.06 that soaks up the entire average active return, while the active returns at the asset class levels do not exhibit a leveraged exposure to the benchmarks? This is a statistical artifact driven by high correlations between the fixed income and equities benchmarks during the crisis period as well as by the high correlation between the fixed income credit factors and the equities benchmark. During 2007 and 2008, stocks and bonds were positively correlated at 0.34, so that losses in both asset classes (particularly the credit-related fixed income losses) cumulated at a time when the Fund benchmark underperformed, artificially inflating the measured Fund-level correlation with the benchmark. The asset class results in Panel B show clearly that over-exposure to the benchmark is not a systematic source of active returns.

A.2: Active Factor Simple Correlations

In Section II.E we measured the Fund's active factor exposures through partial correlations. The advantage of partial correlations, as opposed to simple bivariate correlations, is that they account for the effect of all other potentially correlated variables. The disadvantage is that with 10 factors and only 54 observations in the post crisis sample period, we lose statistical power resulting in only few significant exposures in Panels D of Tables 1 and 2. In this case, we believe the advantage outweighs the disadvantage because, as we saw in Table 4, the correlations between factors are high (in magnitude), so bivariate computations may suggest many significant relationships that are driven by a single or few common components.


Table A.2 illustrates this issue. It reports simple bivariate correlations and their p-values. Panel A is for the Fund level and Panel B for the fixed income and equity components. Most of the simple correlations are large in magnitude and have strong statistical significance, suggesting, misleadingly, that the Fund and its components have large exposures to all risk factors. Contrasting this conclusion to the results from Panels D of Tables 1 and 2, it is clear that the high correlations with all three credit factors, for example, is driven by a common component that is most cleanly captured by CreditAa, the only full-sample significant factor in Panel D of Table 1. We use all factors in our main analysis because we wish to a priori span all active management activities NBIM could undertake in its current mandate and be consistent with Ang, Goetzmann and Schaefer (2009).

Appendix B: Total Portfolio Approach

Examples of the Total Portfolio Approach at Work

We provide two investment examples to better illustrate how we use the Total Portfolio Approach. The private equity example shows how we adjust for the higher debt levels typically present in private equity investments versus public companies. The real estate example shows how we analyze the equity/debt split and therefore recognize the underlying economic exposures that are inherent in an unlevered core real estate asset.

EXAMPLE: FUNDING A PRIVATE EQUITY TRANSACTION USING THE TOTAL PORTFOLIO APPROACH




To maintain total portfolio risk/return balance, we take two actions:

- Match the sector and geographic regions of the public equities that are sold against the private equity that is bought; and
- Fund a \$1.00 private equity purchase by selling \$1.30 of public equity and buying \$0.30 of bonds as well as the new private equity investment. The resulting higher bond content in the total portfolio offsets the higher risk inherent in private versus public equity.

Real estate entails elements of both equity- and debt-like returns and risk characteristics. Accordingly, we:

- Assign each real estate investment to a risk category – low, core or high;
- Fund new purchases by selling a mix of passive fixed income and equity holdings designed to match the risk category; and
- If the property is mortgaged, our net investment is correspondingly riskier; the equity funding component is therefore raised, and the debt component lowered, to maintain the total portfolio risk.

EXAMPLE: FUNDING A REAL ESTATE TRANSACTION USING THE TOTAL PORTFOLIO APPROACH



Source: CPPIB, 2010 Annual Report, pp. 26-27.

Appendix C: Benchmarking Real Estate

In this appendix, we undertake calculations showing that real estate investments induce large deviations relative to a benchmark of publicly traded equities and bonds. As part of the calculations, we show how a “funding mix” of real estate can be computed in terms of stocks and bonds. Viewing real estate in terms of the risk of equities and bonds, plus a specific real-estate component, is referred to as the “Total Portfolio Approach” in Section III.C.

We take direct real estate returns from the National Council of Real Estate Investment Fiduciaries (NCREIF), which constructs a property index from data reported by its members. Because real estate is only periodically re-appraised and the appraisal process does not immediately reflect prevailing market values, the raw NCREIF series exhibits volatility that is artificially low. As the NCREIF series is quarterly, all our analysis is at this frequency. From June 1978 to September 2013, the volatility of the NCREIF real estate index is 2.19%, compared to 8.11% for the S&P 500 and 5.72% for long-term U.S. corporate bonds. The two public market indexes are from Ibbotson Associates. At first glance, it appears that real estate returns as reported by NCREIF are much less volatile than equities and bonds. This is not the case because reported real estate returns are artificially too low due to infrequent trading, lagged appraisals not reflecting market values, and other biases.

C.1: Unsmoothing Reported Real Estate Returns

We begin by “unsmoothing” or “de-smoothing” the NCREIF series following Geltner (1991) and Ross and Zisler (1991). The autocorrelation of NCREIF returns is 0.78 in the sample. This high autocorrelation is partly the result that the appraised returns are functions of current and lagged true returns and do not immediately reflect true market values. We apply a de-smoothing procedure that assumes that the smoothed NCREIF returns, r_t^{obs} , follow:⁶²

$$r_t^{obs} = (1 - \phi)r_t^{true} + \phi r_{t-1}^{obs}, \quad (A.1)$$

where r_t^{true} are the true real estate returns. The parameter ϕ is the autocorrelation of the raw NCREIF returns. Equation (A.1) states that the observed NCREIF returns partially contemporaneously reflect the true, market returns, weighted by ϕ , but they do not immediately adjust to market values because they smooth the previous reported returns, with weight $1 - \phi$.

⁶² For further details on de-smoothing algorithms and other biases of infrequently traded assets, see chapter 13 of Ang (2014).

We filter the observed NCREIF series to estimate the true real estate returns by inverting equation (A.1) to obtain:

$$r_t^{true} = \frac{1}{1-\phi} r_t^{obs} - \frac{\phi}{1-\phi} r_{t-1}^{obs}. \quad (A.2)$$

In Panel A of Figure A.1, we plot the raw NCREIF series and the de-smoothed real estate returns. The de-smoothing process increases the volatility of the raw NCREIF returns from 2.19% to 6.27% per quarter, which is in between the volatilities of stocks (8.11%) and bonds (5.72%).

C.2: Funding Portfolio Mix of Real Estate

To obtain a portfolio of public market assets which can serve as a factor benchmark for real estate, we regress the de-smoothed real estate returns onto stock and bond returns. We impose the restriction that the coefficients on bonds and stocks sum to one. We obtain the following estimates:⁶³

$$\underbrace{r_t^{true}}_{\$1} = -0.50 + \underbrace{0.49r_t^{stock} + 0.51r_t^{bond}}_{\$1} + \varepsilon_t, \quad (A.3)$$

where r_t^{stock} and r_t^{bond} denote the returns of stocks and bonds, respectively. Equation (A.3) indicates that the opportunity cost of investing \$1 in real estate is a portfolio of 49 cents in stocks and 51 cents in bonds. In addition to this public markets benchmark, the investor obtains a real estate-specific return, ε_t . This can be interpreted as a regression “error term.” It also represents the mean-zero component of real estate returns that are not available in public markets. The constant term of -0.50% per quarter indicates that NCREIF real estate has underperformed the 51% equity/49% bond portfolio over the sample.

As the regression R^2 in equation (A.3) is 24%, only a minority of (de-smoothed) real estate returns can be replicated in public markets. Under the interpretation of the Total Portfolio Approach, this represents a lot of scope for active management to generate returns that cannot be obtained passively in public markets. The fact that the average real estate investment, as measured by NCREIF, has underperformed an appropriate funding mix of stocks and bonds indicates that skill and implementation of real estate investments (which are developed comparative advantages for the Fund, see Section III.A), should be prerequisites for investments in this space.

⁶³ There is an implicit assumption of the factor correlation structure in estimating the funding mix, as is true for all factor models. The same correlation structure assumed when choosing the Reference Portfolio should be consistently applied to the Total Portfolio Approach.

We have estimated a funding basket in equation (A.3) by (constrained) ordinary least squares. As Section III.B discusses, in practice an appropriate public funding mix in the Opportunity Cost Model might also adjust for geographic and industry risk, the risk stemming from factor exposures, and specific leverage ratios of different deals. The regression is a very crude way of constructing the funding mix, and in fact cannot be used for private markets with short histories of returns. Other methods of obtaining a benchmark portfolio of public market assets could involve valuation or economic models.

Panel B of Figure A.1 represents a visualization of the funding mix of stocks and bonds vs. the reported NCREIF returns. We show the raw NCREIF returns, without the de-smoothing adjustment. We overlay the returns of the 51% equity/49% bond funding portfolio. We have smoothed the public markets portfolio so that it is comparable with the raw NCREIF returns. We smooth using the current and past five quarters. The six-quarter moving average window corresponds to the period in which most of the news at a given time is incorporated into the NCREIF series. Specifically, the number of periods such that the autocorrelated NCREIF process returns to 75% of its distance from the mean is $\ln(0.25)/\ln(0.78)$, which is approximately six quarters. With this moving average, the smoothed funding mix returns have a volatility of 2.18% per quarter, which is nearly identical to the volatility of 2.19% for the raw NCREIF returns. The correlation between the two series in Panel B, Figure A.1 is 29%, similar to the R^2 of 0.24% in regression (A.3).

C.3: Real Estate Deviations from the Public Market Funding Portfolio

In Panel B of Figure A.1, the volatility of the difference between the raw NCREIF returns and the smoothed funding portfolio of stocks and bonds is 2.60% per quarter. This can be conventionally interpreted as the “tracking error” relative to a smoothed 49% equity/51% bond benchmark. Alternatively using our preferred terminology of “benchmark deviation,” it represents the volatility of reported NCREIF returns in excess of a comparable risk-adjusted investment in equities and bonds that has been averaged over time to reflect the same smoothing biases in the real estate series. To put the benchmark deviation volatility of 2.60% of the NCREIF series into perspective, the volatility of the smoothed equity-bond mix is 2.18% (the unsmoothed volatility is 5.15%). An investor contemplating funding a real estate investment from an equity-bond benchmark should be prepared to endure deviations from the benchmark of the same order of magnitude as movements in the equity-bond benchmark itself.

Another interpretation is that real estate represents an opportunity for an investor to transfer risk taking from the public markets to real estate. Consider a dollar currently invested in the public market benchmark. This investment fluctuates reflecting the risk of the benchmark. Now that dollar is taken from the public benchmark and invested in real estate. This investment is subject to the same fluctuations of the benchmark portfolio from which it is funded. In

addition, it is exposed to specific risks in real estate. The risks emanating solely from the real estate market, which cannot be accounted for by the exposure of real estate to general economic fluctuations reflected in equity and bond markets, have a volatility at least as great as the foregone investment in public markets.

The benchmark deviations of real estate can also be measured in terms of unsmoothed returns. The quarterly volatility of a 60% equity and 40% bond portfolio is 5.58%. A portfolio consisting of 60% equity, 35% bonds, and 5% real estate has a quarterly volatility of 5.47%, where real estate returns are measured after de-smoothing the NCREIF series. There is a reduction in volatility because of the diversification benefits of real estate; real estate offers returns not available in equity and bond markets. The 60% equity/35% bonds/5% real estate portfolio has a deviation volatility of 0.47% per quarter relative to the 60% equity/40% bond benchmark. This benchmark deviation volatility of the portfolio with real estate is similar to the Fund's current total tracking error limit of 0.5% per quarter and represents nearly one-tenth of the standard deviation of the 60% equity/40% bond benchmark itself.

In summary, real estate investment entails significant deviations from a risk-matched equity and bond portfolio.

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, 323-350.
- Ang, A., 2014, *Asset Management: A Systematic Approach to Factor Investing*, forthcoming Oxford University Press.
- Ang, A., W. N. Goetzmann, and S. Schaefer, 2009, Evaluation of active management of the Norwegian Government Pension Fund – Global, Report to the Norwegian Ministry of Finance.
- Ang, A., W. N. Goetzmann, and S. Schaefer, 2011, Efficient market theory and evidence: Implications for active management, *Foundations and Trends*, 5, 157-242.
- Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang, 2006, The cross section of volatility and expected returns, *Journal of Finance*, 61, 259-299.
- Ang A., R. J. Hodrick, Y. Xing, and X. Zhang, 2009, High idiosyncratic volatility and low returns: International and further U.S. evidence, *Journal of Financial Economics*, 91, 1-23.
- Ang, A., B. Chen, W. N. Goetzmann, and L. Phalippou, 2013, Estimating private equity returns from limited partner cash flows, working paper, Columbia University.
- Ang, A., and M. Sorensen, 2012, Risk, returns, and optimal holdings of private equity, *Quarterly Journal of Finance*, 2, 3, DOI: 10.1142/S2010139212500115.
- Baker, M., and J. Wurgler, 2000, The equity share in new issues and aggregate stock returns, *Journal of Finance*, 55, 2219-2257.
- Banz, R. W., 1981, The relationship between return and market value of common stocks, *Journal of Financial Economics*, 9, 3-18.
- Bao, J., J. Pan, and J. Wang, 2011, The illiquidity of corporate bonds, *Journal of Finance*, 66, 911-946.
- Barber, B. M., and T. Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance*, 55, 773-806.
- Bender, J., R. Briand, F. Nielsen, and D. Stefek, 2010, Portfolio of risk premia: A new approach to diversification, *Journal of Portfolio Management*, 36, Winter, 17-25.
- Bernstein, S., J. Lerner, and A. Schoar, 2013, The investment strategies of sovereign wealth funds, *Journal of Economic Perspectives*, 27, 219-238.

- Black, F., 1990, Equilibrium exchange rate hedging, *Journal of Finance*, 45, 899-907.
- Black, F., M. C. Jensen, and M. S. Scholes, 1972, The capital asset pricing model: Some empirical tests, in Jensen, M. C., ed., *Studies in the Theory of Capital Markets*, Praeger, New York.
- van Binsbergen, J. H., M. W. Brandt, and R. S. J. Koijen, 2012, Decentralized decision making in investment management, in Scherer, B., and K. Winston, eds., *The Oxford Handbook of Quantitative Asset Management*, Oxford University Press.
- Brandt, M. W., 2009, Portfolio choice problems, in Aït-Sahalia, Y., and L. P. Hansen, eds., *Handbook of Financial Econometrics: Volume 1 – Tools and Techniques*, Elsevier.
- Campbell, J. Y., K. Serfaty-de Medeiros, and L. M. Viceira, 2010, Global currency hedging, *Journal of Finance*, 65, 87-121.
- Campbell, J. Y., and J. H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy*, 107, 205-251.
- Campbell, J. Y., and L. M. Viceira, 1999, Consumption and portfolio decisions when expected returns are time varying, *Quarterly Journal of Economics*, 114, 433-495.
- Carhart, M. M., 1997, On persistence in mutual fund returns, *Journal of Finance*, 52, 57-82.
- Chordia, T., R. Roll, and A. Subrahmanyam, 2000, Commonality in liquidity, *Journal of Financial Economics*, 56, 3-28.
- Chien, Y., H. Cole, and H. Lustig, 2012, Is the volatility of the market price of risk due to intermittent portfolio rebalancing? *American Economic Review*, 102, 2859-2896.
- Cowles, A., 1933, Can stock market forecasters forecast? *Econometrica*, 1, 309-324.
- DeCosta, D., F. Leng, and G. Noronha, 2013, Minimum maturity rules: The cost of selling bonds before their time, *Financial Analysts Journal*, 69, May/June, 45-56.
- Dimson, E., P. Marsh, and M. Staunton, 2011, *Credit Suisse Global Investment Returns Sourcebook*.
- Dyck, A., A. Morse, and L. Pomorski, 2013, Politicization and governance in pension funds: Explaining performance, working paper, University of Toronto.
- Dyck, A., and L. Pomorski, 2011, Is bigger better? Size and performance in pension plan management, working paper, University of Toronto.

- Erb, C. E., and C. R. Harvey, 2006, The strategic and tactical value of commodity futures, *Financial Analysts Journal*, 62, 69-97.
- Evstigneev, I. V., and K. R. Schenk-Hoppé, 2002, From rags to riches: on constant proportions investment strategies, *International Journal of Theoretical and Applied Finance*, 5, 563-573.
- Fama, E. F., 1965, The behavior of stock-market prices, *Journal of Business*, 38, 34-105.
- Fama, E. F., and K. R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics*, 33, 3-56.
- Fama, E. F., and K. R. French, 2010, Luck versus skill in the cross-section of mutual fund returns, *Journal of Finance*, 65, 1915-1947.
- Fama, E. F., and K. R. French, 2012, Size, value, and momentum in international stock returns, *Journal of Financial Economics*, 105, 457-472.
- Fang, L. H., V. Ivashina, and J. Lerner, 2013, The disintermediation of financial markets: Direct investing in private equity, working paper, Harvard Business School.
- Frazzini, A., and L. H. Pedersen, 2011, Betting Against Beta, NYU working paper.
- Fung, W., and D. A. Hsieh, 2004, Hedge fund benchmarks: A risk-based approach, *Financial Analysts Journal*, 60, Sep-Oct, 65-80.
- Fung, W., D. A. Hsieh, N. Y. Naik, and T. Ramadorai, 2008, Hedge funds: Performance, risk, and capital formation, *Journal of Finance*, 63, 1777-1803.
- Geltner, D., 1991, Smoothing in appraisal-based returns, *Journal of Real Estate Finance and Economics*, 4, 327-345.
- Grossman, S. J., and J. E. Stiglitz, 1976, Information and competitive price systems, *American Economic Review*, 51, 783-810.
- Gürkaynak, R. S., B. Sack, and J. H. Wright, 2007, The U.S. Treasury yield curve: 1961 to the present, *Journal of Monetary Economics*, 54, 2291-2304.
- Harris, R. S., T. Jenkinson, and S. N. Kaplan, 2012, Private equity performance: What do we know?, NBER Working Paper 17874.
- Harvey, C. R., Y. Liu, and H. Zhu, 2013, ...And the cross section of expected returns, working paper, Duke University.

- He, Z., and W. Xiong, 2013, Delegated asset management, investment mandates, and capital immobility, *Journal of Financial Economics*, 107, 239-258.
- Hochberg, Y. V., and J. D. Rauh, 2013, Local overweighting and underperformance: Evidence from limited partner private equity investments, *Review of Financial Studies*, 26, 403-451.
- Ilmanen, A., 2011, *Expected Returns: An Investor's Guide to Harvesting Market Rewards*, Wiley.
- Ilmanen, A., and J. Kizer, 2012, The death of diversification has been greatly exaggerated, *Journal of Portfolio Management*, 38, Spring, 15-27.
- Jagannathan, R., A. Malakhov, and D. Novikov, 2010, Do hot hands persist among hedge fund managers? An empirical evaluation, *Journal of Finance*, 65, 217-255.
- Jegadeesh, N., and S. Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance*, 48, 65-91.
- de Jong, F., and J. Driessen, 2013, The Norwegian Government Pension Fund's potential for capturing illiquidity premiums, Report to the Norwegian Ministry of Finance.
- Kahneman, D., and A. Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica*, 47, 263-292.
- Kaplan, S. N., and A. Schoar, 2005, Private equity performance: Returns, persistence, and capital flows, *Journal of Finance*, 60, 1791-1823.
- Lintner, J., 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics*, 47, 13-37.
- Malmendier, U., and S. Nagel, 2011, Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics*, 126, 373-416.
- Markowitz, H., 1952, Portfolio selection, *Journal of Finance*, 7, 77-91.
- Merton, R. C., 1969, Lifetime portfolio selection under uncertainty: The continuous-time case, *Review of Economics and Statistics*, 51, 247-257.
- Merton, R. C., 1971, Optimal consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory*, 3, 373-413.
- Merton, R. C., 1973, An intertemporal capital asset pricing model, *Econometrica*, 41, 867-887.
- Mossin, J., 1966, Wages, profits and the dynamics of growth, *Quarterly Journal of Economics*, 80, 376-399.

- Pástor, L., and R. F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy*, 111, 642-685.
- Petajisto, A., 2013, Active share and mutual fund performance, *Financial Analysts Journal*, 69, July/August, 73-93.
- Phalippou, L., 2010, An evaluation of the potential for GPFG to achieve above average returns from investments in private equity and recommendations regarding benchmarking, Report to the Norwegian Ministry of Finance.
- Rapach, D. E., and G. Zhou, 2013, Forecasting stock returns, in Elliott, G. and A. Timmermann, eds., *Handbook of Economic Forecasting*, Elsevier.
- Roll, R., 1977, A critique of the asset pricing theory's tests, Part I: On past and potential testability of the theory, *Journal of Financial Economics*, 4, 129-176.
- Roll, R., 1992, A mean-variance analysis of tracking error, *Journal of Portfolio Management*, 18, summer, 13-22.
- Ross, S. A., 1976, The arbitrage theory of capital asset pricing, *Journal of Economic Theory*, 13, 341-360.
- Ross, S., and R. Zisler, 1991, Risk and return in real estate, *Journal of Real Estate Finance and Economics*, 4, 175-190.
- Samuelson, P. A., 1965, Proof that properly anticipated prices fluctuate randomly, *Industrial Management Review*, 6, 41-49.
- Samuelson, P. A., 1969, Lifetime portfolio selection by dynamic stochastic programming, *Review of Economics and Statistics*, 51, 239-246.
- Sharpe, W. F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance*, 19, 425-42.
- Sharpe, W. F., 2010, Adaptive asset allocation policies, *Financial Analysts Journal*, 66, May/June, 45-59.
- Stambaugh, R. F., 1982, On the exclusion of assets from tests of the two-parameter model, *Journal of Financial Economics*, 10, 237-268.
- Swensen, D. F., 2000, *Pioneering Portfolio Management: An Unconventional Approach to Institutional Investment*, The Free Press, New York.

Vardharaj, R., F. J. Fabozzi, and F. J. Jones, 2004, Determinants of tracking error for equity portfolios, *Journal of Investing*, 13, Summer, 37-47.

Wachter, J., 2010, Asset allocation, *Annual Review of Financial Economics*, 2, 175-206.

Tables

Table 1: Overall Fund Returns
(Monthly Gross Returns)

Panel A: Total Returns

	Mean	StdDev	Autocorr	Skew	Reward-to-Risk Ratio
Full Sample	0.45	2.22	0.25	-0.95	0.20
Pre Financial Crisis	0.52	1.53	0.08	-0.59	0.34
Post Financial Crisis	0.92	2.68	0.20	-0.13	0.34

Panel B: Variance Attribution

	Full Sample	Pre Crisis	Fin. Post Crisis	Fin. Crisis
Benchmark Return	99.3%	99.5%	99.6%	
Active Return	0.7%	0.5%	0.4%	
Total Return	100.0%	100.0%	100.0%	

Panel C: Active Returns

		Mean	StdDev	Autocorr	Skew	Reward-to-Risk Ratio
Full Sample	Coefficient	0.03	0.22	0.52	-2.22	0.12
	P-Value	0.04				
Pre Financial Crisis	Coefficient	0.04	0.11	0.22	0.60	0.35
	P-Value	0.00				
Post Financial Crisis	Coefficient	0.10	0.21	0.62	1.19	0.47
	P-Value	0.00				

Panel D: Partial Correlations of Active Returns with Systematic Factors

	Full Sample		Pre Financial Crisis		Post Financial Crisis	
	Partial Corr	P-Value	Partial Corr	P-Value	Partial Corr	P-Value
Term	0.07	0.34	-0.13	0.18	0.32	0.03
CreditAa	0.41	0.00	-0.03	0.76	0.47	0.00
CreditBaa	-0.07	0.39	0.04	0.70	0.28	0.07
CreditHY	0.08	0.29	0.00	0.97	0.24	0.11
Liquidity	0.28	0.00	0.18	0.08	0.20	0.18
ValGrth	-0.17	0.03	-0.30	0.00	0.10	0.51
Smlg	0.20	0.01	0.38	0.00	0.11	0.48
Mom	0.04	0.56	0.17	0.10	0.25	0.09
LowVol	-0.10	0.20	-0.16	0.11	-0.29	0.05
SellVol	0.30	0.00	0.15	0.15	-0.18	0.24

Note to Table 1: The table presents summary statistics for the monthly gross returns of the Fund denominated in the currency basket. In all panels, the Full Sample corresponds to January 1998 through June 2013. The Pre Financial Crisis sample ends in December 2006 and the Post Financial Crisis sample starts in January 2009. Panel A shows summary statistics of the Fund returns. Panel B shows a variance decomposition of the Fund returns into a component attributed to the benchmark and the residual attributed to active management. Panel C describes the active returns defined as the difference between the Fund returns and the benchmark. The Reward-to-Risk ratio is the monthly average gross return divided by the monthly standard deviation of returns. Panel D presents partial correlations between the active returns of the Fund and the risk factors described in Table 3.

Table 2: Fixed Income and Equity Returns
(Monthly Gross Returns)

Panel A: Total Returns

		Mean	StdDev	Autocorr	Skew	Reward-to-Risk Ratio
<i>Fixed Income</i>	Full Sample	0.43	1.01	0.20	-0.43	0.42
	Pre Fin. Crisis	0.44	0.90	0.12	-0.27	0.49
	Post Fin. Crisis	0.54	1.01	0.37	-0.16	0.53
<i>Equity</i>	Full Sample	0.50	4.55	0.21	-0.76	0.11
	Pre Fin. Crisis	0.66	4.21	0.14	-0.81	0.16
	Post Fin. Crisis	1.18	4.51	0.16	-0.14	0.26

Panel B: Variance Attribution

	<i>Fixed Income</i>			<i>Equity</i>		
	Full Sample	Pre Fin. Crisis	Post Fin. Crisis	Full Sample	Pre Fin. Crisis	Post Fin. Crisis
Benchmark Return	89.3%	99.4%	83.8%	99.7%	99.6%	100.0%
Active Return	10.7%	0.6%	16.2%	0.3%	0.4%	0.0%
Total Return	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

Panel C: Active Returns

		Mean	StdDev	Autocorr	Skew	Reward-to-Risk Ratio
<i>Fixed Income</i>						
Full Sample	Coefficient	0.02	0.33	0.71	-0.59	0.06
	P-Value	0.21				
Pre Fin. Crisis	Coefficient	0.02	0.07	-0.26	-0.83	0.29
	P-Value	0.00				
Post Fin. Crisis	Coefficient	0.17	0.41	0.76	1.75	0.41
	P-Value	0.00				
<i>Equity</i>						
Full Sample	Coefficient	0.05	0.25	0.16	-0.69	0.18
	P-Value	0.01				
Pre Fin. Crisis	Coefficient	0.06	0.26	0.14	0.65	0.22
	P-Value	0.01				
Post Fin. Crisis	Coefficient	0.05	0.12	0.11	-0.50	0.42
	P-Value	0.00				

Panel D: Partial Correlations of Active Returns with Systematic Factors

Fixed Income

	Full Sample		Pre Financial Crisis		Post Financial Crisis	
	Partial Corr	P-Value	Partial Corr	P-Value	Partial Corr	P-Value
Term	0.18	0.01	-0.01	0.92	0.43	0.00
CreditAa	0.45	0.00	0.06	0.54	0.51	0.00
CreditBaa	0.13	0.09	0.08	0.42	0.21	0.15
CreditHY	0.12	0.12	-0.08	0.41	0.29	0.04
Liquidity	0.38	0.00	-0.02	0.83	0.19	0.20

Equity

	Full Sample		Pre Financial Crisis		Post Financial Crisis	
	Partial Corr	P-Value	Partial Corr	P-Value	Partial Corr	P-Value
ValGrth	-0.26	0.00	-0.26	0.01	-0.39	0.01
Smlg	0.29	0.00	0.29	0.00	0.26	0.06
Mom	0.14	0.06	0.23	0.02	-0.08	0.59
LowVol	-0.25	0.00	-0.21	0.03	-0.33	0.02
SellVol	0.27	0.00	0.16	0.10	0.19	0.19

Note to Table 2: The table presents summary statistics for the monthly gross returns of the fixed income and equities components of the Fund, both denominated in the Fund's currency basket. In all panels, the Full Sample corresponds to January 1998 through June 2013. The Pre Financial Crisis sample ends in December 2006 and the Post Financial Crisis sample starts in January 2009. Panel A shows summary statistics of the fixed income and equities returns. Panel B shows a variance decomposition of these returns into a component attributed to the respective benchmarks and the residual attributed to active management. Panel C describes the active returns defined as the difference between the Fund returns and the benchmark. The Reward-to-Risk ratio is the monthly average gross return divided by the monthly standard deviation of returns. Panel D presents partial correlations between the fixed income and equities returns of the Fund and appropriate subsets of the risk factors described in Table 3.

Table 3: Systematic Factors

Notation	Definition	Source
Term	Difference in returns on the total return BarCap US Treasury 20+ yr index and the total return BarCap US Treasury Bill 1-3 mth index	Morningstar
CreditAa	Difference in returns on the total return BarCap US Corporate Aa Long Maturity index and the total return BarCap US Treasury 20+ yr index	Morningstar
CreditBaa	Difference in returns on the total return US Corporate Baa Long Maturity index and the total return BarCap US Corporate Aa Long Maturity index	Morningstar
CreditHY	Difference in returns on the total return BarCap US Corporate High Yield Caa index and the total return BarCap US Corporate Baa Long Maturity Baa index	Morningstar
Liquidity	The negative of innovations in the on-the-run/off-the-run spread on 10-year US Treasury bonds	Federal Reserve Board
ValGrth	Difference in returns between global "value" stocks and global "growth" stocks computed using MSCI world indices	Morningstar
SmLg	Difference in returns between global small cap stocks and global large cap stocks computed using MSCI all country indices	Morningstar
Mom	Difference in returns between US stocks with past high returns and US stocks with past low returns	Kenneth French at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html
LowVol	Difference in returns between US stocks with past low volatilities and US stocks with past high volatilities	Own computations
SellVol	Return on a variance swap between implied and realized volatility on the S&P500 in excess of LIBOR	Bloomberg (MLHFEV1E prior to Jan 2009, MSQTVPPE from Jan 2009 onwards)

Note to Table 3: The table lists the systematic, dynamic factors used in the evaluation of the Fund's active returns. All factors are zero-cost, long-short positions. All factors are returns, except for the Liquidity factor, and are expressed in the Fund's currency basket. Positive realizations of the Liquidity factor indicate deteriorations in liquidity.

Table 4: Risk Factor Correlations

Panel A: Full Sample (Jan 1998 - June 2013)

		Fixed Income Factors							Equity Factors				
		Bonds	Stocks	Term	CreditAa	CreditBaa	CreditHY	Liquidity	ValGrth	SmLg	Mom	LowVol	SellVol
Fixed Income Factors	Bonds	1.00	-0.20	0.79	-0.25	-0.24	-0.53	0.14	0.08	-0.02	-0.03	0.11	0.07
	Stocks	-0.20	1.00	-0.34	0.57	0.45	0.53	0.25	-0.02	0.14	-0.36	-0.59	0.49
	Term	0.79	-0.34	1.00	-0.60	-0.49	-0.68	0.08	0.03	-0.06	0.14	0.16	-0.09
	CreditAa	-0.25	0.57	-0.60	1.00	0.39	0.50	0.09	-0.01	0.22	-0.28	-0.33	0.31
	CreditBaa	-0.24	0.45	-0.49	0.39	1.00	0.57	0.37	-0.06	0.26	-0.24	-0.31	0.53
	CreditHY	-0.53	0.53	-0.68	0.50	0.57	1.00	0.25	0.06	0.26	-0.34	-0.37	0.31
Equity Factors	Liquidity	0.14	0.25	0.08	0.09	0.37	0.25	1.00	-0.01	0.33	-0.21	-0.20	0.55
	ValGrth	0.08	-0.02	0.03	-0.01	-0.06	0.06	-0.01	1.00	0.05	-0.37	0.28	-0.08
	SmLg	-0.02	0.14	-0.06	0.22	0.26	0.26	0.33	0.05	1.00	-0.04	-0.41	0.23
	Mom	-0.03	-0.36	0.14	-0.28	-0.24	-0.34	-0.21	-0.37	-0.04	1.00	0.26	-0.13
	LowVol	0.11	-0.59	0.16	-0.33	-0.31	-0.37	-0.20	0.28	-0.41	0.26	1.00	-0.23
	SellVol	0.07	0.49	-0.09	0.31	0.53	0.31	0.55	-0.08	0.23	-0.13	-0.23	1.00

Panel B: Pre Financial Crisis (Jan 1998 - December 2006)

		Fixed Income Factors							Equity Factors				
		Bonds	Stocks	Term	CreditAa	CreditBaa	CreditHY	Liquidity	ValGrth	SmLg	Mom	LowVol	SellVol
Fixed Income Factors	Bonds	1.00	-0.38	0.90	-0.51	-0.31	-0.53	0.17	0.03	0.02	0.04	0.14	-0.25
	Stocks	-0.38	1.00	-0.33	0.48	0.54	0.40	-0.07	-0.11	-0.07	-0.29	-0.57	0.49
	Term	0.90	-0.33	1.00	-0.62	-0.32	-0.58	0.22	-0.05	0.11	0.13	0.12	-0.24
	CreditAa	-0.51	0.48	-0.62	1.00	0.47	0.53	-0.13	0.04	0.02	-0.30	-0.39	0.39
	CreditBaa	-0.31	0.54	-0.32	0.47	1.00	0.51	0.12	0.05	0.21	-0.27	-0.50	0.46
	CreditHY	-0.53	0.40	-0.58	0.53	0.51	1.00	-0.01	0.10	0.17	-0.26	-0.31	0.22
Equity Factors	Liquidity	0.17	-0.07	0.22	-0.13	0.12	-0.01	1.00	-0.01	0.28	-0.01	-0.10	-0.17
	ValGrth	0.03	-0.11	-0.05	0.04	0.05	0.10	-0.01	1.00	0.07	-0.34	0.43	-0.07
	SmLg	0.02	-0.07	0.11	0.02	0.21	0.17	0.28	0.07	1.00	0.15	-0.35	0.05
	Mom	0.04	-0.29	0.13	-0.30	-0.27	-0.26	-0.01	-0.34	0.15	1.00	0.12	-0.11
	LowVol	0.14	-0.57	0.12	-0.39	-0.50	-0.31	-0.10	0.43	-0.35	0.12	1.00	-0.35
	SellVol	-0.25	0.49	-0.24	0.39	0.46	0.22	-0.17	-0.07	0.05	-0.11	-0.35	1.00

Panel C: Post Financial Crisis (Jan 2009 - June 2013)

		Fixed Income Factors							Equity Factors				
		Bonds	Stocks	Term	CreditAa	CreditBaa	CreditHY	Liquidity	ValGrth	SmLg	Mom	LowVol	SellVol
Fixed Income Factors	Bonds	1.00	-0.05	0.69	-0.29	-0.29	-0.47	0.18	0.25	-0.33	-0.06	0.05	0.03
	Stocks	-0.05	1.00	-0.43	0.64	0.32	0.71	0.18	0.40	0.44	-0.49	-0.75	0.53
	Term	0.69	-0.43	1.00	-0.75	-0.70	-0.76	0.14	0.16	-0.42	0.19	0.25	-0.24
	CreditAa	-0.29	0.64	-0.75	1.00	0.53	0.61	-0.06	0.01	0.38	-0.32	-0.41	0.34
	CreditBaa	-0.29	0.32	-0.70	0.53	1.00	0.54	-0.07	-0.12	0.44	-0.30	-0.24	0.15
	CreditHY	-0.47	0.71	-0.76	0.61	0.54	1.00	0.20	0.18	0.56	-0.55	-0.58	0.27
Equity Factors	Liquidity	0.18	0.18	0.14	-0.06	-0.07	0.20	1.00	0.22	0.21	-0.38	-0.39	0.12
	ValGrth	0.25	0.40	0.16	0.01	-0.12	0.18	0.22	1.00	0.01	-0.51	-0.47	0.16
	SmLg	-0.33	0.44	-0.42	0.38	0.44	0.56	0.21	0.01	1.00	-0.42	-0.62	0.25
	Mom	-0.06	-0.49	0.19	-0.32	-0.30	-0.55	-0.38	-0.51	-0.42	1.00	0.71	0.01
	LowVol	0.05	-0.75	0.25	-0.41	-0.24	-0.58	-0.39	-0.47	-0.62	0.71	1.00	-0.33
	SellVol	0.03	0.53	-0.24	0.34	0.15	0.27	0.12	0.16	0.25	0.01	-0.33	1.00

Note to Table 4: The table shows correlations of monthly returns on the bond benchmark, the equities benchmark, and the systematic risk factors described in Table 3. All factors are zero-cost, long-short positions. All factors are returns, except for the Liquidity factor, and are expressed in the Fund's currency basket. Positive realizations of the Liquidity factor indicate deteriorations in liquidity.

Table 5: Risk Factor Autocorrelations

	Fixed Income Factors							Equity Factors				
	Bonds	Stocks	Term	CreditAa	CreditBaa	CreditHY	Liquidity	ValGrth	SmLg	Mom	LowVol	SellVol
Full Sample	0.12	0.20	0.03	0.07	0.14	0.27	0.26	0.15	0.10	0.08	0.03	0.42
Pre Financial Crisis	0.12	0.12	0.00	0.08	0.28	0.17	0.06	0.13	0.07	-0.06	-0.03	0.07
Post Financial Crisis	0.13	0.15	0.12	0.01	0.06	0.16	0.18	0.14	0.08	0.36	0.31	0.04

Note to Table 5: The table shows autocorrelations of monthly returns on the bond benchmark, the stock benchmark, and the systematic risk factors described in Table 3. All factors are zero-cost, long-short positions. All factors are returns, except for the Liquidity factor, and are expressed in the Fund's currency basket. Positive realizations of the Liquidity factor indicate deteriorations in liquidity. The Full Sample corresponds to January 1998 through June 2013. The Pre Financial Crisis sample ends in December 2006 and the Post Financial Crisis sample starts in January 2009.

Table A.1: Residual Returns**Panel A: Overall Fund**

		Alpha	Beta	Skew	Reward-to-Risk Ratio
Full Sample	Coefficient	0.00	1.06	-1.35	0.01
	P-Value	0.45			
Pre Financial Crisis	Coefficient	0.03	1.02	0.40	0.28
	P-Value	0.00			
Post Financial Crisis	Coefficient	0.06	1.05	1.31	0.34
	P-Value	0.01			

Panel B: Fixed Income and Equity Components

		Alpha	Beta	Skew	Reward-to-Risk Ratio
<i>Fixed Income</i>					
Full Sample	Coefficient	0.02	1.00	-0.60	0.05
	P-Value	0.25			
Pre Fin. Crisis	Coefficient	0.02	1.00	-0.80	0.32
	P-Value	0.00			
Post Fin. Crisis	Coefficient	0.16	1.03	1.68	0.39
	P-Value	0.01			
<i>Equity</i>					
Full Sample	Coefficient	0.04	1.02	-0.44	0.16
	P-Value	0.02			
Pre Fin. Crisis	Coefficient	0.05	1.01	0.51	0.20
	P-Value	0.02			
Post Fin. Crisis	Coefficient	0.03	1.02	-0.10	0.33
	P-Value	0.01			

Note to Table A.1: The table presents summary statistics for the monthly residual returns constructed as the residual from a regression of the active returns on the appropriate benchmark returns. Panel A is for the overall Fund and Panel B is for the fixed income and equities components. Alpha is the intercept of the regression and beta is the slope coefficient. The Full Sample corresponds to January 1998 through June 2013. The Pre Financial Crisis sample ends in December 2006 and the Post Financial Crisis sample starts in January 2009. The Reward-to-Risk ratio is the monthly average gross return divided by the monthly standard deviation of returns.

Table A.2: Simple Active Factor Correlations

Panel A: Overall Fund

	Full Sample		Pre Financial Crisis		Post Financial Crisis	
	Partial Corr	P-Value	Partial Corr	P-Value	Partial Corr	P-Value
Term	-0.22	0.00	-0.08	0.43	-0.32	0.02
CreditAa	0.53	0.00	0.12	0.20	0.54	0.00
CreditBaa	0.41	0.00	0.26	0.01	0.37	0.01
CreditHY	0.39	0.00	0.16	0.10	0.53	0.00
Liquidity	0.50	0.00	0.22	0.02	0.31	0.02
ValGrth	-0.18	0.01	-0.42	0.00	0.25	0.07
SmLg	0.42	0.00	0.47	0.00	0.50	0.00
Mom	-0.15	0.05	0.25	0.01	-0.46	0.00
LowVol	-0.42	0.00	-0.52	0.00	-0.64	0.00
SellVol	0.58	0.00	0.23	0.02	0.24	0.09

Panel B: Fixed Income and Equity Components

Fixed Income

	Full Sample		Pre Financial Crisis		Post Financial Crisis	
	Partial Corr	P-Value	Partial Corr	P-Value	Partial Corr	P-Value
Term	-0.18	0.01	-0.04	0.71	-0.13	0.35
CreditAa	0.48	0.00	0.08	0.40	0.42	0.00
CreditBaa	0.40	0.00	0.08	0.41	0.21	0.12
CreditHY	0.37	0.00	-0.01	0.92	0.35	0.01
Liquidity	0.50	0.00	-0.02	0.80	0.31	0.02


Equity

	Full Sample		Pre Financial Crisis		Post Financial Crisis	
	Partial Corr	P-Value	Partial Corr	P-Value	Partial Corr	P-Value
ValGrth	-0.38	0.00	-0.44	0.00	-0.10	0.49
Smlg	0.40	0.00	0.37	0.00	0.62	0.00
Mom	0.12	0.11	0.30	0.00	-0.33	0.02
LowVol	-0.45	0.00	-0.47	0.00	-0.56	0.00
SellVol	0.35	0.00	0.22	0.02	0.31	0.02

Note to Table A.2: This table presents simple bivariate correlations between the active returns of the fund in Panel A or of the fixed income and equities components in Panel B and appropriate subsets of the risk factors described in Table 3. The Full Sample corresponds to January 1998 through June 2013. The Pre Financial Crisis sample ends in December 2006 and the Post Financial Crisis sample starts in January 2009.


Figures

Figure 1: Rebalancing Premium 1961-2013




Note to Figure 1: Panel A plots cumulated wealth starting with \$1 at the beginning of June 1961 ending in June 2013 for a strategy that rebalances over stocks and bonds in the solid black line, and one that passively holds the market weights of stocks and bonds in the dashed red line. We take market weights and total returns of Treasury bonds, bills, and notes from Ibbotson Associates, and market weights and total returns of all stocks listed on CRSP on the NYSE, Amex, and NASDAQ exchanges. Rebalancing occurs quarterly and the weights are fixed at 60% equity and 40% bonds. We scale the passive, market-weighted stock and bond portfolio to have the same volatility as the rebalanced strategy. The annualized Reward-to-Risk ratio (which does not subtract the risk-free rate) of the rebalanced strategy is 0.92 and the market-weighted passive strategy is 0.78. Panel B shows the portfolio weights in the rebalanced and passive strategies.

Figure 2: Rebalancing Premium 1994-2013




Note to Figure 2: Panel A plots cumulated wealth starting with \$1 at the beginning of February 1994 ending in September 2013 for a strategy that rebalances over global stocks and bonds in the solid black line, and one that passively holds the market weights of stocks and bonds in the dashed red line. We take market weights and total returns of the Barclays Global Aggregate bond index and market weights and total returns of stocks in the FTSE World Index. Rebalancing occurs quarterly and the weights are fixed at 60% equity and 40% bonds. We scale the passive, market-weighted stock and bond portfolio to have the same volatility as the rebalanced strategy. The annualized Reward-to-Risk ratio of the rebalanced strategy is 0.82 and the market-weighted passive strategy is 0.70. Panel B shows the portfolio weights in the rebalanced and passive strategies.

Figure 3: Overall Fund Cumulated Active Returns




Note to Figure 3: This figure plots the cumulative active return on the Fund, defined as the difference between the Fund returns and the benchmark returns. The vertical orange dotted line represents the beginning of the post-crisis subsample.

Figure 4: Overall Fund Active Returns




Note to Figure 4: This figure graphs as red dots the active returns on the Fund, defined as the difference between the Fund returns and the benchmark returns, and as the blue line the fitted active returns obtained from a full-sample regression of the active returns on the risk factors described in Table 3. The vertical orange dotted line represents the beginning of the post-crisis subsample. The legend also provides the adjusted R^2 of the risk factor regression.

Figure 5: Overall Fund Active Returns Post Financial Crisis



Note to Figure 5: This figure shows as red dots the active returns on the Fund, defined as the difference between the Fund returns and the benchmark returns, and as blue line the fitted active returns obtained from a sub-sample regression of the active returns on the risk factors described in Table 3. The sample is restricted to the post-financial crisis period starting January 2009. The legend also provides the adjusted R^2 of the risk factor regression.

Figure 6: Cumulated Active Returns by Asset Class




Note to Figure 6: This figure shows the cumulative active returns on the fixed income (Panel A) and equities (Panel B) components of the Fund, defined as the differences between the two returns and their respective benchmark returns. The vertical orange dotted line represents the beginning of the post-crisis subsample.

Figure 7: Active Returns by Asset Class




Note to Figure 7: This figure graphs as red dots the active returns on the fixed income (Panel A) and equities (Panel B) components of the Fund, defined as the differences between the two returns and their respective benchmark, and as blue line the fitted active returns obtained from full-sample regressions of the active returns on subsets of the risk factors described in Table 3. For the fixed income returns, we use Term, CreditAa, CreditBaa, CreditHY, and Liquidity as explanatory factors. For the equity returns, we use ValGrth, Smlg, Mom, LowVol, and SellVol as explanatory factors. The vertical orange dotted line represents the beginning of the post crisis subsample. The legend also provides the adjusted R^2 of the risk factor regression.


Figure 8: Active Returns by Asset Class Post Financial Crisis



Note to Figure 8: This figure graphs as red dots the active returns on the fixed income (Panel A) and equities (Panel B) components of the Fund, defined as the differences between the two returns and their respective benchmark returns, and as blue line the fitted active returns obtained from a sub-sample regressions of the active returns on subsets of the risk factors described in Table 3. The sample is restricted to the post-financial crisis period starting January 2009. For the fixed income returns, we use Term, CreditAa, CreditBaa, CreditHY, and Liquidity as explanatory factors. For the equity returns, we use ValGrth, SmLg, Mom, LowVol, and SellVol as explanatory factors. The vertical orange dotted line represents the beginning of the post crisis subsample. The legend also provides the adjusted R^2 of the risk factor regression.


Figure 9: Systematic Factors






Note to Figure 9: Panel A plots cumulated returns of the fixed income factors: Term, CreditAa, CreditBaa, and CreditHY. Panel B plots the on-the-run/off-the-run Treasury spread, so that higher values denote periods of higher illiquidity. The Liquidity factor is the negative first difference of this spread. Panel C plots cumulated returns on the equity factors: ValGrth, SmLg, Mom, LowVol, and SellVol. All factors in Panels A and C are expressed in the Fund's global currency basket. Note that Panel B is not a return series. The sample period in all panels is January 1998 to June 2013.

Figure 10: Benchmark Deviation Volatility



Note to Figure 10: We plot the standard deviation of active returns on the Fund as well as the fixed income and equities components, each defined as returns net of the respective benchmarks. The standard deviation at month t is measured over a 36-month centered window $[t-16, t, t+16]$. The vertical orange dotted line represents the beginning of the post-crisis subsample.


Figure 11: Benchmark Deviations of (Developed) vs. (Developed + Emerging + Frontier)




Note to Figure 11: We plot the difference in returns, or benchmark deviations, between the MSCI Developed and Emerging index and the MSCI Developed, Emerging, and Frontier index. In both cases, the indexes are computed by capitalization weighting the Developed, Emerging, and/or Frontier market indexes. The red dots show the monthly benchmark deviations, while the black line plots a two-sided, 12-month, two-sided moving average which is computed using the returns in the interval $[t-6, t, t+6]$.


Figure 12: Tail Risk – Full Sample

Panel A: Probability Density




Panel B: Cumulative Probability Density






Note to Figure 12: Panel A plots in the blue line the empirical distribution of active returns constructed by applying a nonparametric kernel density smoother to the full sample of active returns. We also plot as orange line a matching normal density with the same mean and standard deviation. Panel B shows the corresponding cumulative density functions. Panel C shows the implied expected shortfall, defined as the expected return conditional on the return falling beyond a certain critical threshold in the tail of the distribution. The tail threshold is shown on the vertical axis and the corresponding expected loss on the horizontal axis. All panels use the full sample of active returns.


Figure 13: Tail Risk in the Non-Crisis Sample





Note to Figure 13: The plots are the same as in Figure 12, except the sample period excludes the financial crisis. We use the period January 1998 to December 2006 and January 2010 to June 2013.


Figure 14: Management of the Fund




Note to Figure 14: The figure breaks down each stage of the investment process of the Fund, marking the parts for which the Ministry of Finance and NBIM are responsible. The true passive market is non-investable, so measurability starts with the adoption of a Passive Benchmark. The Weighted Benchmark takes into account tilts towards non-market capitalization positions, including country weights and ethical screens. The Rebalanced Weighted Benchmark specifies the rebalancing regime. The Operational Reference Portfolio (ORP) within NBIM obtains superior diversification compared to the Rebalanced Weighted Benchmark and systematically harvests factor risk premiums. Security Selection adds value in excess of the ORP.

Figure A.1: Real Estate Returns and a Funding Portfolio of Stocks and Bonds

Panel A: Unsmoothing Real Estate Returns



Panel B: NCREIF vs. Smoothed Funding Mix of Stocks and Bonds



Note to Figure A.1: Panel A shows the raw NCREIF real estate returns at the quarterly frequency in the solid red line. After de-smoothing the returns using the process in equation (A.1), we obtain the unsmoothed real estate returns shown in black dots. Panel B plots the raw NCREIF returns in the solid red line. The returns of a matching portfolio consisting of 49% stocks (S&P 500) and 51% bonds (U.S. long-term corporates) are overlaid in the dashed blue line. The returns of the funding mix portfolio have been smoothed by taking moving averages over the past six quarters.