#8 - 2012

NBIM DISCUSSION NOTE

This paper was part of the NBIM memo "On equity investments" (February 2012)

Capturing Systematic Risk Premia

30/03/2012

This note illustrates the empirical risk/return characteristics of the different risk premia, and how one can design scalable investment strategies to capture systematic risk premia.

Main findings

- A huge academic literature documents that exposure to systematic risk in the equity market increases the return relative to a market-capitalisation-weighted index. Our research documents that several risk premia represent attractive sources of additional returns also for a large investor requiring high investment capacity. However, capturing systematic risk premia entails risk. We argue that a fund with the Government Pension Fund Global's defining characteristics has a comparative advantage in taking those risks.

- Strategies for capturing systematic risk premia can be constructed in a number of different ways. We illustrate how different methodological approaches impact the risk and return related to the capture of different systematic risk premia, and show that there are significant differences across methodologies, regions, market segments and time. Efficient capture of systematic risk premia requires carefully designed investment strategies that are customised to each risk premium.

- We focus on five well-documented systematic risk premia, and illustrate how these risk premia may be captured by an investor requiring high investment capacity. The size of the different risk premia is significant even when we focus on the most liquid market segments. Each risk premium has, however, a substantial negative tail, and the success of a premium-harvesting strategy depends on the investor's ability to sit through periods of underperformance.

- The illiquidity premium is one of the risk premia that are suitable for a large investor with a long investment horizon to capture. Due to the challenges related to defining and isolating this premium, we suggest that this premium should be captured indirectly through other risk premia such as size, value and volatility.

- The correlation between the different risk premia and their correlation with the market are generally low. Hence, the diversification benefits from approaching the risk premia within a unified framework are substantial. We show that a combined capture of a wide set of risk premia not only leads to lower volatility, but more importantly also to significantly reduced tail risk.

- Each risk premium has indirect exposures to other sources of systematic risk. Moreover, these indirect risk exposures vary over time. A strategy seeking to capture systematic risk premia therefore needs to actively manage both direct and indirect risk exposures, possibly also taking systematic risk exposures emanating from other parts of the Fund into account. As a result, the capturing of systematic risk premia should be implemented as a part of the overall risk allocation process within the Fund rather than through a passive tilt in the benchmark.

Defining systematic risk premia

Whereas exposure to the equity risk premium can easily be obtained through a simple buy-and-hold strategy in the broad equity market, capturing other risk premia in the equity market is usually more demanding. Equity risk premia such as the value premium and the momentum premium can only be harvested through dynamic trading strategies where the portfolio of stocks one would like to hold changes over time as the attributes of the different stocks change. There are many ways to design and implement strategies seeking to capture risk premia in the equity market. Some of the most important decisions that have to be made are:

- Which indicators for the underlying risk premium should be used?
- What should the underlying universe of stocks be?
- How should the factor-mimicking portfolio be constructed?
- How often should the factor-mimicking portfolio be rebalanced?
- How should risk, turnover and trading costs be monitored and managed?

In this and the next section, we will go through how academic researchers and commercial index providers have addressed these questions. We then analyse the internally developed investment strategies for capturing various risk premia in isolation, as well as a combined strategy seeking to exploit multiple sources of systematic risk. The purpose of this study is to illustrate the empirical risk/return characteristics of the different risk premia, and how one can design scalable investment strategies to capture systematic risk premia. There are a number of issues that are relevant if the investment strategies are to be implemented in the markets that are not addressed here. These issues include e.g. how the overall risk budget should be distributed across the different risk premia, and how the portfolio should be optimized with respect to minimizing turnover and market impact.

The value and small-cap premia

There are a large number of studies of the value premium in academic research. Different researchers have used different specifications of the value premium, but most specifications relate the price of a stock to earnings, cash flow, dividends, book assets or some other measure of fundamental company value. Two of the most well-known papers on the value premium are Fama and French (1993) and Lakonishok, Shleifer and Vishny (1994). These two studies use fairly similar specifications of the value premium, including price to book (P/B), price to cash flow (P/CF) and price to earnings (P/E). They also broadly agree on the empirical evidence of the value premium, but propose different interpretations of it. Thanks partly to easy data availability, the Fama-French value premium specification using P/B has become the benchmark value series in academic finance.

The Fama-French methodology for calculating factor-mimicking portfolios is fairly straightforward. Each quarter, the universe of US stocks, with some exceptions, is sorted based on each stock's size and P/B. Based on whether the market cap of a company is below or above the median market cap, a small-cap and a large-cap universe are identified. A long-short portfolio mimicking the value premium is constructed by going long the cheapest 30 percent of the stocks in both the small-cap and large-cap universes, and short the 30 percent of the most expensive stocks in the respective universes. The stocks that go into the long basket and the short basket are equally weighted. The return on this portfolio is then interpreted as the value premium. Next quarter, the factor-mimicking portfolio is rebalanced using the same procedure.

A small-cap premium is calculated along the same lines: the ranking based on P/B defines a value and a growth universe consisting of the cheapest and richest 30 percent of the stocks, respectively. The remaining 40 percent are called neutral stocks. A long-short portfolio mimicking the small-cap premium is constructed by overweighting the small-cap stocks within each of the three valuation segments, and underweighting the big companies within the same segments. The stocks that go into the long basket and the short basket are equally weighted. The return on this portfolio is interpreted as the small-cap premium.

The purpose of the Fama-French double-sorting procedure is to reduce any small-cap bias in the value factor, and to reduce any value bias in the small-cap factor. If, for example, there is a strong correlation between P/B and market capitalisation, a portfolio based on a sort on P/B alone would also have an exposure to market capitalisation. The double-sorting procedure reduces such biases, and as a result creates "pure" value returns and "pure" small-cap returns. On the other hand, this approach means that the small-cap segment weighs 50 percent in the calculation of the value return, whereas the market-cap weight of the small caps is much less than 50 percent. This means that an investor trying to capture the value premium through the Fama-French procedure will have to invest a bigger percentage in the illiquid small-cap segment than in the liquid large-cap segment. The Fama-French value return series is therefore considered as the return on a hypothetical equity portfolio that is close to impossible to implement in practice, as the methodology does not take into account transaction costs, market impact and trading restrictions normally faced by investment managers (Cremers, Petajisto and Zitzewitz 2010); Huij and Verbeek 2007).

An additional concern is related to the industry tilts in the factor-mimicking portfolios, which potentially can be large. The Fama-French sorting procedure does not take into account what industry the companies belong to. Hence, this methodology allows large and persistent industry concentrations in the factor-mimicking portfolios if there are big valuation differences across sectors. The value return could therefore potentially be impacted by industry-specific effects that are unrelated to the value premium. Risk management of the factor-mimicking portfolio, where exposures to industries and other risk factors are monitored and managed, is another example of practical considerations that are important for an investor seeking to capture the value premium.

The benchmark value series from Fama-French is calculated based on US companies only. However, there is also a value return series for international markets calculated by Fama and French. The international value aggregate is calculated using a somewhat different methodology: for each country, the companies are sorted into value and growth portfolios based on four valuation ratios: price to book, price to earnings, price to cash earnings, and dividend yield. The international version of the Fama-French value return is calculated by weighting together each country's factor return using the market capitalisation of the different countries as weights. Hence, the methodology differs from the standard Fama-French procedure both because it uses several different valuation indicators instead of just one, and because there is no adjustment for size in the construction of the country portfolios.

The momentum premium

Jegadeesh and Titman (1993) documented that strategies which buy stocks that have performed well in the past and sell stocks that have performed poorly in the past generate significant positive returns over holding periods of 3-12 months. Based on these results, Carhart (1997) used a specification where the momentum effect was defined as the equally-weighted average of firms with the highest 30 percent 11-month returns lagged one month, minus the equally-weighted average of firms with the lowest 30 percent 11-month returns lagged one month. The momentum portfolios are rebalanced on a monthly basis. This specification has become a standard in the academic literature.

Fama and French have also developed their own specification of the momentum effect, where they apply the same double-sorting procedure as for the value and size effect. Each month, six value-weighted portfolios are formed on size and prior 12-month return, where the most recent month is skipped. The portfolios used to construct the momentum return are the intersections of two portfolios formed on size, defined as market capitalisation, and three portfolios formed on prior 12-month return. The monthly size breakpoint is the median NYSE market capitalisation, whereas the monthly prior 12-month return breakpoints are the 30th and 70th NYSE percentiles. The Fama-French momentum factor is then specified as the average return on the two high prior return portfolios minus the average return on the two low prior return portfolios. Hence, the Fama-French momentum factor can be interpreted as a size-adjusted momentum factor.

Practitioners in the asset management industry often employ a wider set of momentum factors, seeking to capture momentum effects over other horizons as well as over the 12-month horizon which has become the standard in academic research. Moreover, factors capturing momentum effects in returns are often also augmented with factors capturing market sentiment towards the different stocks. An example of a momentum factor not based on prior returns is earnings revisions, where sell-side analysts' upgrades and downgrades of earnings estimates are used as an indicator of market sentiment (Chan, Jegadeesh and Lakonishok 1996).

The volatility premium

The low-volatility premium refers to the observation that low-volatility stocks and low-beta stocks tend to outperform high-volatility and high-beta stocks, which runs counter to the fundamental economic principle that risk is compensated with higher expected return. Although this pattern was pointed out decades ago (Black, Jensen and Scholes 1972; Black 1972, 1993; Haugen and Baker 1991), this topic has received renewed interest in more recent years (Ang et al. 2006, 2009; Baker, Bradley and Wurgler 2011; Frazzini and Pedersen 2010). There are various explanations for this premium, but perhaps the most important ones focus on leverage constraints and 'lottery ticket' preferences among investors. Most active money is delegated to asset managers whose performance is measured relative to a broad equity market index. Furthermore, these managers are more often than not limited in the amounts of leverage that they can use in their portfolios. For these managers, low-beta stocks will look risky as they are not allowed to leverage up, and offer little in return to compensate. Instead, they tend to focus on the high-volatility segment, where the stocks that potentially may give an extraordinary return without using leverage are located. As a result, there is little demand for low-volatility stocks among these investors, effectively depressing their price and increasing their return to investors willing to hold them, and at the same time bidding up the prices on high-volatility stocks. If these managers have significant assets under management, the impact could be substantial. Baker, Bradley and Wurgler (2011), Frazzini and Pedersen (2010) and Barberis and Huang (2008) discuss this mechanism in more detail.

Frazzini and Pedersen (2010) propose a methodology for constructing factor-mimicking portfolios capturing the low-beta premium. Their BAB ("Betting Against Beta") factor is constructed both on a country-by-country basis and for all stocks taken together. Given a universe of stocks, all securities are ranked in ascending order on the basis of their estimated beta. The betas are calculated for each security relative to the corresponding MSCI local market index using a rolling window with one year of daily data or three years of monthly data, depending on data availability. The ranked stocks are assigned to one of two portfolios: low beta and high beta. Securities are weighted by the ranked betas and the portfolios are rebalanced every calendar month. The low-beta portfolio is leveraged to a beta of 1, and the high-beta portfolio is de-leveraged to a beta of 1, such that both portfolios are rescaled to have a beta of 1 at portfolio formation. Hence, the BAB factor is the zero-cost, zero-beta portfolio that is long the low-beta portfolio and short the high-beta portfolio.

Consistent with the model, Frazzini and Pedersen find that the BAB factor produces significant risk-adjusted returns. However, these profits are not riskless, as leveraged portfolios tend to underperform when funding constraints tighten. The authors show that the BAB portfolios tend to lose money when the TED spread widens, presumably because tightening funding conditions lead to de-leveraging of the low-beta, leveraged portfolios. Thus the BAB strategy resembles a strategy harvesting illiquidity premia in that at least part of the profitability reflects liquidity risk.

Searching for exploitable risk premia

The last 30 years have seen the development of a number of different indices seeking to capture different risk premia in the equity market. Traditional style indices developed by several major index providers have a number of advantages when it comes to investment capacity and robustness compared to the Fama-French factors. As a result of this, traditional style indices give a better representation of how much of the value premium can realistically be captured by an investor.

Traditional style indices have several drawbacks, however, such as potentially large unintended tilts towards risk factors other than the targeted factor. As a result of this, a new generation of indices has been developed. These new indices are sometimes called investment process indices because they aim at emulating the activities of active investment managers seeking to capture various systematic risk premia. In other words, these indices represent passive, rules-based investment strategies where, to varying degrees, practical considerations regarding liquidity, turnover, rebalancing and the management of indirect exposures to other sources of systematic risk are taken into consideration.

These new strategy indices also cover risk premia such as momentum, which is characterised by a high frequency of rebalancing and high turnover, features considered undesirable from the perspective of traditional passive, low-turnover, low-cost indices.

Traditional style indices

During the 1980s, the providers of core equity indices moved into the style index arena. A key driver of this development was the increased popularity of equity style management. Sharpe (1988, 1992) established the "returns-based style analysis" methodology, where regression analysis was used to determine the exposure of various portfolios, such as mutual funds, to different style factors. Sharpe focused on large cap, small cap, value and growth as style choices for his analysis, which was a key step in establishing the popular equity styles seen today in the marketplace.

Early style indices were broadly constructed by sorting the overall universe of stocks according to a single factor such as the price-to-book ratio, and then dividing the universe in half. The low P/B half became the value index, and the other half became the growth index. This methodology was applied to various broad large-cap and small-cap equity indices.

Today's equity style indices are constructed using a somewhat more sophisticated methodology. Instead of defining growth as the opposite of value, modern style indices define growth more explicitly, for example in terms of realised growth in earnings or revenues and/or projected growth. In addition, many index providers now establish both value and growth classifications using multiple indicators of value and growth, not just one single factor such as P/B. As a result of this, most modern style indices use a two-dimensional process for dividing the core universe into growth and value, using one set of indicators to identify value stocks and a different set of indicators to define growth stocks. As a consequence of this, stocks generally are not necessarily either value or growth, but can have characteristics of both styles. Still, in most cases the two style halves continue to be defined in terms of each other in the sense that the two styles together cover fully and exactly the market capitalisation and constituents of the broader core index.

Compared to the theoretical benchmark indices calculated by Fama and French, the traditional style indices have significantly higher investment capacity. This is due to the fact that traditional style indices are broader in the sense that they cover a larger part of the universe of stocks, rather than focusing only on the cheapest and richest tails of the stock universe. In addition, the traditional style indices are constructed based on investable market capitalisation, rather than equally weighting the stocks as in the Fama-French methodology. In addition to increasing investment capacity, these differences in methodology may also have consequences for the risk and return related to the different risk premia.

FTSE style indices

FTSE is one of the providers of style indices. FTSE creates a value index and a growth index based on the same universe of constituents that go into the construction of the broad FTSE All-World Index.

FTSE uses four different value measures and five different growth measures, listed in Table 1. After a normalising process where the different measures are standardised and outliers are removed, a composite value score and a composite growth score are assigned to each company. The different value measures and growth measures are equally weighted when calculating the composite value and growth scores. As a result of this process, each company has a value score, ranging from 0 (high value) to 100 (low value), and a growth score, ranging from 100 (high growth) to 0 (low growth). Finally, an overall style ranking (OSR) is calculated for each company by averaging the composite value and growth scores.

Table 1: Factors used by different style index providers

	Value measures	Growth measures
FTSE	Book to price	3-year historical sales growth
	Sales to price	3-year historical EPS growth
	Dividend yield	2-year forward sales estimates
	Cash flow to price	2-year forward EPS estimates
		Return on equity*(1-payout ratio)
MSCI	Book to price	Long-term forward EPS growth rate
	12-month forward earnings to price	Short-term forward EPS growth rate
	Dividend yield	Current internal growth rate
		Long-term historical EPS growth trend
		Long-term historical sales-per-share growth trend
Russell	Book to price	2-year forward EPS growth rate
		5-year historical sales-per-share growth rate

The style indices or style portfolios are created on a country-by-country basis. For the purposes of index construction, the euro zone is treated as one country. The companies in each country are ranked according to their OSR, with growth at the top (high OSR) of the ranking and value at the bottom (low OSR) of the ranking. Constituents showing high growth characteristics (OSR tending towards 100) covering the top 35 percent of the investible market capitalisation of the country are allocated to the growth index at a weight of 100 percent. Constituents showing high value characteristics (OSR tending towards 0) covering the bottom 35 percent of the investible market capitalisation of the country are allocated to the value index at a weight of 100 percent. The constituents covering the middle 30 percent of the investible market capitalisation of the country are allocated across both the value and growth indices according to the bands in Table 2.

Table 2: Investible market-cap style weighting in country

Between 55% and 64.99%	75% growth /25% value
Between 45% and 54.99%	50% growth /50% value
Between 35.01% and 44.99%	25% growth /75% value

By using this methodology, the sum of the investible market capitalisations of the value and growth benchmarks will equal the investible market capitalisations of the underlying parent benchmark. Also, the relative weights of each country will be the same in the style indices as in the underlying benchmark. This process is then repeated for each relevant country in the FTSE All-World Index. The style indices are rebalanced semi-annually in June and December.

MSCI style indices

MSCI is another major index provider which is also producing style indices. Initially, the value and growth indices were constructed based on a single-dimensional framework that allocated securities in an MSCI Standard Country Index to either value or growth based on their price-to-book ratios. As views on style definition and segmentation continued to develop, MSCI's methodology also evolved into a two-dimensional framework for style segmentation in which value and growth securities are categorised using different attributes. In addition, multiple factors are used to identify value and growth characteristics. The value and growth investment style characteristics for index construction are defined using the variables listed in Table 1.

The MSCI Global Investable Market Value and Growth Indices are constructed from the constituents of the MSCI Global Investable Market Indices on a country-by-country basis for each of the countries included in the MSCI All Country World Index, except for developed markets in Europe where the style indices are constructed on a regional basis.

MSCI's construction of the value and growth indices for each country index involves the following steps:

- Standardise the value and growth measures by calculating z-scores for each variable listed in Table 1 used to specify value and growth.

- For each security, calculate an overall value z-score and an overall growth z-score by taking the equally-weighted average of the underlying value and growth z-scores. The value z-score and the growth z-score of a security define its overall style characteristics and its positioning within the value and growth style space.

- Based on the value and growth z-scores, assign initial style inclusion factors for each security. Each security has two style inclusion factors, one for value, called the Value Inclusion Factor (VIF), and the other for growth, called the Growth Inclusion Factor (GIF), and they represent the proportion of a security's free-float-adjusted market capitalisation that should be allocated to the value and/or growth indices. The sum of the VIF and the GIF is always equal to 1. There are five possible values for the style inclusion factors: 1, 0.65, 0.5, 0.35 and 0, depending on how the different companies score with respect to value and growth.

- Allocate securities to the value and growth indices. The value and growth indices target a 50 percent free-float-adjusted market capitalisation representation for each of the value and growth indices in each MSCI Market Index. In order to achieve the 50 percent target, the style allocation process involves several steps that are not discussed in detail here.

In this two-dimensional framework, non-value does not necessarily mean growth, and vice versa. Additionally, some securities can exhibit both value and growth characteristics, while others may exhibit neither. MSCI constructs and maintains the value and growth indices by allocating securities and their free-float-adjusted market capitalisations to the appropriate value and growth indices during the semiannual style index reviews that take place each May and November.

Russell style indices

Russell Investments was one of the first index providers to create style indices. The methodology for constructing these style indices has changed over time. Beginning with reconstitution in 2011, Russell uses three variables in the determination of growth and value, see Table 1.

Russell creates style indices based on several different parent indices. The Russell Global Index represents the investable global equity market and includes the top 98 percent of US market capitalisation (the Russell 3000 Index) and the top 98 percent of the rest of the world's market capitalisation.

The process for assigning growth and value weights is applied separately to large-cap and small-cap stocks in the Russell Global Index. Treating large-cap and small-cap stocks separately prevents the possible distortion to relative valuations that may occur if the global index is used as the base index. For each base index, stocks are ranked by their value and growth measures. These rankings are converted to standardised units and combined to produce a composite value score (CVS). Stocks are then ranked by their CVS, and an algorithm is applied to the CVS distribution to assign growth and value weights to each stock. In general, a stock with a lower CVS is considered growth; a stock with a higher CVS is considered value. A stock with a CVS in the middle range is considered to have both growth and value characteristics, and is weighted proportionately in the growth and value index. Stocks are always fully represented by the combination of their growth and value weights: for example, a stock that is given a 20 percent weight in a Russell value index will have an 80 percent weight in the same Russell growth index. Roughly 70 percent are classified as all value or all growth and 30 percent are weighted proportionately to both value and growth.

In an effort to mitigate unnecessary turnover, Russell implements a banding methodology at the CVS level of the style algorithm. If a company's CVS change from the previous year is less than +/-0.15 and the company remains in the same core index, then the CVS remains unchanged during the next reconstitution process. The banding methodology has proven to reduce turnover caused by smaller, less meaningful movements while continuing to allow the larger, more meaningful changes to occur, signalling a true change in a company's relation to the market.

Investment process indices

Traditional style indices have a number of advantages when it comes to tradability and robustness compared to the Fama-French factors. As a result of this, traditional style indices give a better representation of how much of the value premium can realistically be captured by an investor. A strategy for capturing the value premium can, in principle, be implemented by investing part of the fund in the value style index instead of the market portfolio or the growth portfolio. However, this approach has several potential drawbacks:

- Using a traditional style index based on screening on different metrics typically leads to portfolios that, in addition to the targeted factor, also have significant exposures to other sources of systematic risk. Melas et al. (2010) show that the MSCI World Value Index has significant exposures to size, momentum and volatility in addition to the targeted value factor. In addition, the MSCI World Value Index may also have significant industry concentration relative to the MSCI World. For example, as at July 2009, the traditional MSCI Europe Value Index overweighted the financial sector by 16.4 percent and underweighted the consumer staples sector by 8.6 percent relative to the standard MSCI Europe Index. This means that the performance of a strategy that implements a value tilt through the MSCI World Value Index may be driven by unintentional industry bets.

- Traditional style indices are constructed based on a complete coverage philosophy, forcing stocks that may have little or no real value or growth qualities into one camp or the other. While this method does account for all stocks in the underlying universe, it is not the most efficient way to obtain the returns of a core style. Fullness of coverage nevertheless continues to override purity of capture in the most popular style indices. Capturing the value premium through overweighting the value index and underweighting the growth index will mean that we have active positions in all stocks in the universe. Although this will give us a broad, diversified factor portfolio, the return potential could potentially be enhanced by implementing a pure value portfolio involving a smaller number of stocks. Limiting the number of stocks that are used in capturing a risk premium is also advantageous from the practical standpoint of an investment manager.

- A common belief of index providers is that indices need to be as stable, and have as little turnover, as possible. As a result of this, rebalancing of equity style indices usually occurs semiannually or annually. The primary motivation for this rebalancing frequency is to control turnover and the transaction costs related to turnover. A style such as value is a slow-moving, low-frequency investment strategy. Therefore, semiannual or annual rebalancing of the style indices does not necessarily interfere with effective capture of the value premium. However, effective capture of other sources of systematic risk premia, such as momentum, may require more frequent rebalancing. Whereas traditional style indices are constructed to keep turnover at a minimum, investment process indices do not focus on turnover per se, but rather aim at maximising after-cost returns. This means that one can allow higher turnover and higher transaction costs if the efficacy of the factor warrants a higher rebalancing frequency.

Against this backdrop, a new generation of indices has been developed. These new indices are sometimes called investment process indices because they aim at emulating the activities of active investment managers seeking to capture various systematic risk premia. In other words, these indices represent passive, rules-based investment strategies where, to varying degrees, practical considerations regarding liquidity, trading costs, rebalancing and the management of indirect exposures to other sources of systematic risk are taken into consideration.

MSCI-Barra Factor Indices

MSCI has developed a family of factor indices that aim to capture some important style factors in an index. The style indices currently available include value, momentum, volatility and leverage.

MSCI's long-short factor indices are constructed by optimising a parent MSCI Index to achieve:

- a specified high level of exposure to a particular style factor
- very low exposure to all other style, industry and country factors, and
- low tracking error relative to the corresponding MSCI benchmark index.

In addition, a number of optimisation constraints are employed in an effort to minimise stock-specific risk as well as controlling tradability and turnover. The optimisation constraints are listed in Table 3.

For investors with restrictions on shorting, MSCI has developed long-only factor indices for some of the style factors. The MSCI Europe Value Tilt Index is constructed from the MSCI Europe Index and is designed to have high exposure to the Barra Value Factor. The Barra Value Factor aims to capture the extent to which a company's ongoing business is inexpensively priced in the marketplace and is based on three metrics: book to price, earnings to price and sales to price. The MSCI Europe Momentum Tilt Index is constructed from the MSCI Europe Index and is designed to have high exposure to the Barra Momentum Factor. The Barra Momentum Factor aims to identify stocks that have been recently successful based on historical alpha (five-year beta regression) and price behaviour in the market, measured by 12-month returns.

All style indices are rebalanced on a monthly basis.

Table 3: Optimisation constraints

	Long-short factor indices	Long-only factor indices
Portfolio leverage	MSCI Europe Barra Momentum Index fixed at 130/30 MSCI Europe Barra Value Index fixed at 130/30 MSCI Europe Barra Low Volatility Index fixed at 150/50	N/A
Net equity exposure	100%	100%
Shorting cost	<100 bp cut-off for security to be included as short position <133 bp cut-off for security currently held as short position	N/A
Max. number of stocks	400	200
Max. stock weight	+/- 2% relative to weight in benchmark	+/- 2% relative to weight in benchmark
Monthly turnover	< 5%	< 5%
Trade limit	< 10% of one-month average daily volume	< 10% of one-month average daily volume

FTSE ActiveBeta Indices

The FTSE ActiveBeta Index Series is designed to offer efficient, high-capacity capture of two systematic risk premia: value and momentum. The FTSE Active Beta Momentum and Value Indices are based on all the constituent securities in the FTSE Global Equity Index Series weighted on the basis of free-float market capitalisation. In addition to indices for each risk premium, a combined index combining both sources of systematic return is calculated. The negative correlation of value and momentum allows a combined capture to provide greater consistency and stability of returns over time, compared to the independent capture of either momentum or value. Hence, the family of FTSE Active Beta Indices for any given universe comprises three indices: 1) a FTSE Active Beta Momentum Index, 2) a FTSE Active Beta Value Index, and 3) a combined FTSE Active Beta Momentum and Value Index.

These indices are created using the following three-step process:

- Specify a selection universe from which the FTSE Active Beta Indices will be created. The reference universes for the public indices will be the commonly-used FTSE market or regional indices. Each stock in the given selection universe is then ranked based on the momentum and value signals independently. Momentum is defined as past 12-month total return, whereas value is defined as a composite signal consisting of price to book value, price to sales and price to cash flow (or price to earnings, where appropriate). The composite signal is an equally-weighted average of the three valuation ratios.

- Create independent FTSE Active Beta Momentum and FTSE Active Beta Value Indices. These indices are created using buy and sell thresholds, which is referred to as buffer-based portfolio construction. A universe is first ranked by the momentum and value signals independently. A stock is included in a style index if its style rank places it within the top third of the universe market capitalisation, and is excluded from the index when its style rank places it in the bottom third of the universe market capitalisation. This buffer methodology limits turnover and results in about 50 percent market capitalisation coverage of the underlying universe in each style index. The selected securities are then weighted according to their relative float-adjusted market capitalisation.

- Combine the FTSE Active Beta Momentum Index and FTSE Active Beta Value Index to create the FTSE Active Beta MVI. In the creation of the FTSE Active Beta MVI, independent security-level positions from each style index are added with a 50 percent weight.

Although the methodology is fairly similar to the traditional style indices, the FTSE ActiveBeta Index Series applies a buffer-based construction process in order to limit turnover. Momentum strategies are characterised by a high frequency of rebalancing and high turnover, features considered undesirable from the perspective of passive, low-turnover, low-cost indices. The buffer-based index construction methodology, however, allows efficient capture of momentum strategies by limiting turnover to a reasonable level.

NBIM systematic risk factors

In the following, we calculate our own factor returns, partly because we would like to have a more detailed insight into how different assumptions regarding the design of the factor-mimicking portfolio impact the risk and return characteristics, and partly because there are no readily available factor return series for some of the risk premia we are addressing. A short description of our methodology is given below.

We apply a simple and transparent approach, using a standard procedure where stocks are sorted into ten groups, based on the risk factor under investigation. We then form ten decile portfolios by equally weighting the stocks within each group. A long and a short portfolio are then constructed based on the ten decile portfolios. The way we do this varies depending on the factor under investigation. There is a trade-off between investment capacity and the pureness of the factor return. A narrow specification, where we only use the most extreme decile portfolios as our long and short portfolios would often give a higher factor return, but will also have higher risk and lower investment capacity compared to a broader approach where we use several decile portfolios to construct the long and short portfolio.

For example, when constructing a factor-mimicking portfolio for the small-cap factor, the investment capacity in the decile portfolios with the smallest stocks will, by definition, be less liquid relative to the decile portfolio with the largest stocks. Hence, an investment strategy where we overweight the decile portfolio with the smallest stocks and underweight the decile portfolio with the largest stocks will have limited investment capacity. In order to take this skewed investability into account, we instead construct a long portfolio consisting of the five decile portfolios with the smallest stocks, and a short portfolio consisting of the decile portfolio with the largest stocks.

Whereas the factor-mimicking portfolio for the small-cap factor will be asymmetric, we use a symmetric approach for the value and momentum factors, where the long and short portfolios consist of the three decile portfolios with the highest and lowest factor scores, respectively. However, since there are reasons to believe that the factor return is more concentrated in the extreme decile portfolios, we weigh the three decile portfolios with weights of 3/6, 2/6 and 1/6 respectively, putting more emphasis on the extreme portfolios. This approach seeks to balance the need for high investment capacity with an attempt to capture the factor returns in the tails of the distribution. The factor return is then calculated as the return on the long portfolio minus the return on the short portfolio.

For the value premium, we use a composite factor consisting of earnings to price, book to price, cash flow to price, and dividend yield. The four factors are equally weighted. The small-cap factor is simply based on the market capitalisation of each stock, whereas the momentum factor is calculated as the prior 12-month return, where the most recent month is skipped. These specifications are standard specifications of the three risk premia.

For the volatility factor, we use an approach similar to Frazzini and Pedersen (2010), but use volatility calculated over a rolling window of two years using monthly data instead of the beta. We form ten decile portfolios based on each stock's volatility, where the stocks in each decile portfolio are equally weighted. Our long portfolio consists of the three decile portfolios with the lowest volatility, and our short portfolio consists of the three decile portfolios with the highest volatility. We calculate an ex-ante beta estimate for both the long and short portfolio based on the realised beta of the two portfolios, following the procedure in Frazzini and Pedersen (2010). The long and short portfolios are then scaled in order to have an ex-ante beta equal to 1 for both the long and the short portfolio. Since the low-volatility portfolio will typically have a beta lower than 1, this means that we are leveraging up in this portfolio. Analogously, the high-volatility portfolio typically has a beta higher than 1, meaning that we have to de-leverage the exposure in this portfolio in order to obtain a beta equal to 1. The net leverage is funded using the one-month US LIBOR.

For all factors, we use the FTSE All-World large/mid-cap index excluding emerging markets as our underlying universe. Excluding small caps and emerging markets will increase the investment capacity in the calculated factor returns.

Empirical characteristics of risk premia in the equity market

In this section, we look more closely at the risk and return characteristics of different strategies seeking to capture systematic risk premia in the equity market.

Table 4 summarises the return and risk for the four most common risk premia in the equity market, based on the theoretical benchmark factor returns for the US market calculated by Fama and French. Exposing a portfolio to the value, small-cap and momentum risk factors over the sample period starting in 1926 would have added significant value, as can be seen from Table 4. However, exposing the portfolio to the different factors also involves additional risk and long periods with negative excess returns, see Table 4 and Chart 1.

Table 4 compares the three systematic risk factors with the equity risk premium, as calculated for the US market by Fama and French. We see that all three factors have had both lower volatility and smaller drawdowns than the equity risk premium. In Table 4, "shortfall" denotes the average performance in the 5 percent worst-performing months for each factor, whereas "resilience" measures the average performance for each factor in the 5 percent worst months in terms of equity market performance. Although both value and small cap have had a positive beta tilt on average over the entire sample period, the underperformance in periods with large market drawdowns is modest. This is consistent with the findings in Ilmanen (2011). The momentum factor has a large negative tail, but this factor has had a negative beta tilt over the sample period. In the 5 percent worst months in terms of equity market performance, the momentum factor has, on average, had a significant positive excess return. In other words, the large negative tail of the momentum factor does not coincide with market drawdowns.

The beta tilts in the different factor returns have not been constant through time, see Chart 2. Interestingly, there has been a downward trend in the beta of the value factor over several decades, as pointed out by Franzoni (2006). Since 1960, the value factor has, on average, had a negative beta tilt. The correlation between the systematic risk factors is fairly low, see Table 5 and Chart 3. Perhaps the most important feature is the fact that the momentum premium has a negative correlation with the other equity risk premia.

As a result of 1) positive average returns, 2) low correlations both with each other and with the market, and 3) attractive tail risk characteristics, tilting a market-cap-weighted portfolio towards a mix of the three risk premia would have improved the risk-return trade-off. However, although the theoretical Fama-French factors have impressive risk-return properties, the question is how much of these effects are actually available to investors seeking to harvest these premia in practice, and how the risk-return properties will change once we impose requirements with respect to investment capacity and risk management. This will be addressed in the sections below.

Table 4: Fama-French factors, US equity market, 1926-2011

	Ret	Vol	Sharpe	Beta	Shortfall	Resilience
Equity	7.6 %	18.7 %	0.40	1.00	-12.2 %	-12.2 %
Value	4.0 %	12.8 %	0.31	0.16	-7.2 %	-1.3 %
Small cap	3.0 %	11.2 %	0.27	0.20	-5.9 %	-2.0 %
Momentum	8.4 %	16.6 %	0.50	-0.30	-12.1 %	3.9 %

Table 5: Correlation between systematic risk factors in the US market, 1926-2011

	Equity	Value	Small cap	Momentum
Equity	1.00			
Value	0.24	1.00		
Small cap	0.34	0.12	1.00	
Momentum	-0.34	-0.53	-0.20	1.00

Chart 1: Fama-French value and small-cap factors in the US market, cumulated returns

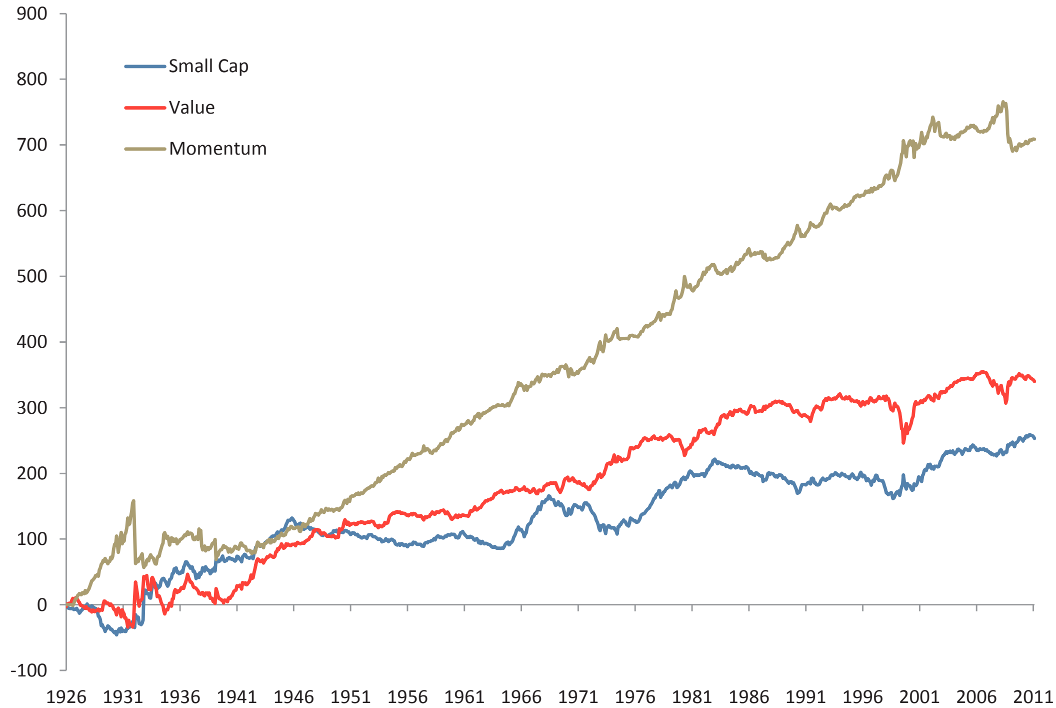

Source: Kenneth French's website, NBIM calculations

Chart 2: Beta for factor returns, rolling 24-month window

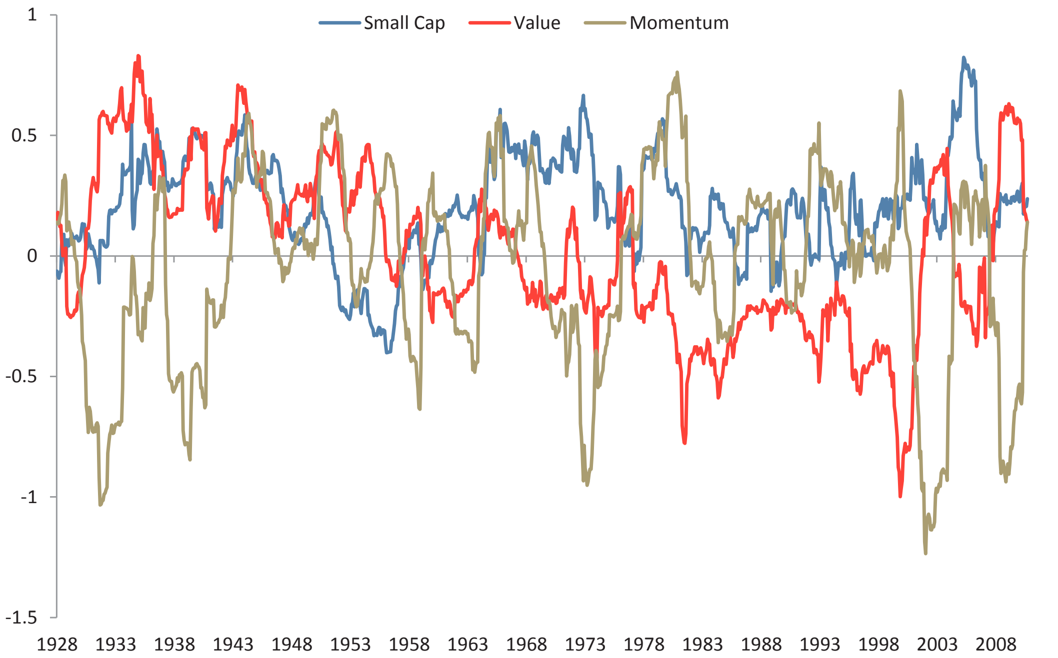

Source: Kenneth French's website, NBIM calculations

Chart 3: Rolling pairwise correlations between factor returns, rolling 24-month window

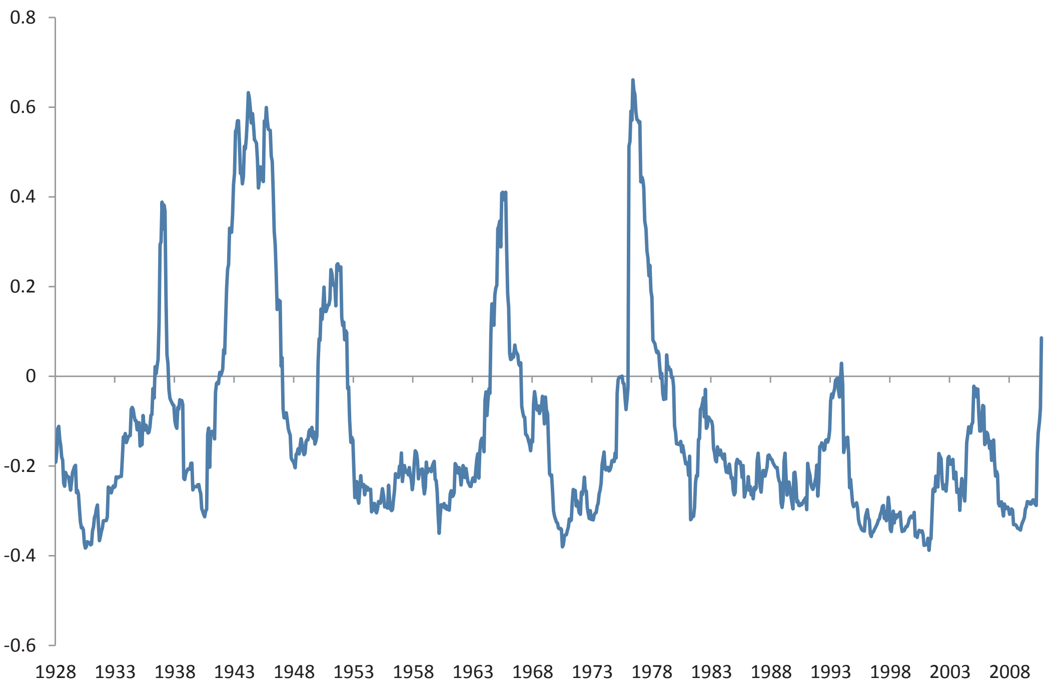

Source: Kenneth French's website, NBIM calculations

Value

Value premia in different market segments
In addition to the US market, Fama and French also calculate factor returns for international markets, but with a shorter history. There are four different factor return series representing value for non-US markets, one for each of the four valuation metrics used in the calculation of factor portfolios. Below, we use a composite value return series calculated as the equally-weighted average of the four factor return series. In order to calculate global Fama-French returns, we equally weight the US factor return and the international factor return.

Table 6 summarises the value return based on the Fama-French methodology, globally as well as for the three different regions, for the period 1975-2011. We compare the Fama-French factors with the MSCI style indices. MSCI World VmG ("Value minus Growth") represents the difference in returns between a long exposure in the MSCI World Value index and a short exposure in the MSCI World Growth index. These indices represent an improvement relative to the Fama-French returns when it comes to investability and liquidity, as explained above, but this comes at a cost in terms of performance, consistent with Houge and Loughran (2006). We see that the average returns and Sharpe ratios are consistently lower in all regions compared to the Fama-French benchmark series.

There are significant regional differences in the value premium. Chart 4 plots the cumulated value returns based on the MSCI methodology, and Chart 5 shows the average pairwise correlations between the value returns in different regions using a rolling window of 24 months. The regional differences we see for the MSCI style indices are similar to those of Fama-French returns, with the highest Sharpe in the Pacific region and the lowest Sharpe in the US. Chart 5 illustrates further the fact that the correlations in value returns in different regions are time-varying and generally fairly low.

We also report the value premium for emerging markets based on MSCI's methodology and for the Russell style indices over the period 2003-2011, broken down by size segments. The results indicate that the value premium is higher in less liquid and less efficient segments of the equity market.

Table 6: Value premia in different market segments, 1975-2011

	Ret	Vol	Sharpe	Beta	Shortfall	Resilience
Theoretical Benchmarks						
Fama-French, Global	4.6 %	8.5 %	0.55	-0.03	-5.0 %	-0.3 %
Fama-French, US	3.3 %	12.1 %	0.28	-0.09	-7.6 %	-0.3 %
Fama-French, Europe	3.7 %	7.7 %	0.48	0.06	-4.3 %	-1.0 %
Fama-French, Pacific	8.6 %	10.7 %	0.81	-0.07	-5.8 %	1.1 %
Traditional Style Indexes, Developed Mkt						
MSCI World VMG	2.3 %	6.8 %	0.34	-0.07	-4.1 %	0.1 %
MSCI US VMG	0.7 %	9.2 %	0.08	-0.13	-6.0 %	0.6 %
MSCI Europe VMG	1.7 %	6.9 %	0.25	0.05	-4.4 %	-1.0 %
MSCI Pacific VMG	4.8 %	9.9 %	0.48	-0.12	-6.3 %	1.8 %
Traditional Style Indexes, Emerging Mkt						
MSCI Emerging Markets VMG, 2003-2011	3.8 %	4.2 %	0.91	0.00	-2.2 %	0.0 %
Traditional Style Indexes, Size Segments						
Russell Global Dev Large Cap VMG, 2003-2011	-0.2 %	5.0 %	-0.05	0.08	-3.2 %	-0.9 %
Russell Global Dev Mid Cap VMG, 2003-2011	0.1 %	4.6 %	0.02	-0.02	-2.6 %	0.0 %
Russell Global Dev Small Cap VMG, 2003-2011	1.1 %	5.6 %	0.19	-0.12	-2.8 %	1.2 %

Chart 4: Value premia in different regions based on MSCI methodology. Cumulated returns

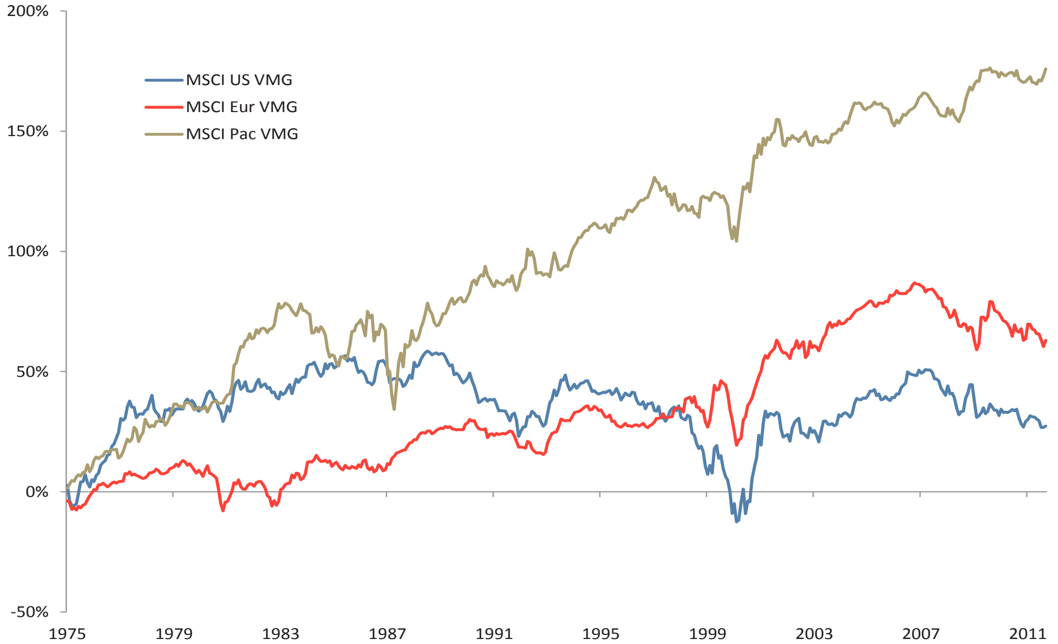

Source: MSCI, NBIM calculations

Chart 5: Average pairwise correlations between value returns in different regions

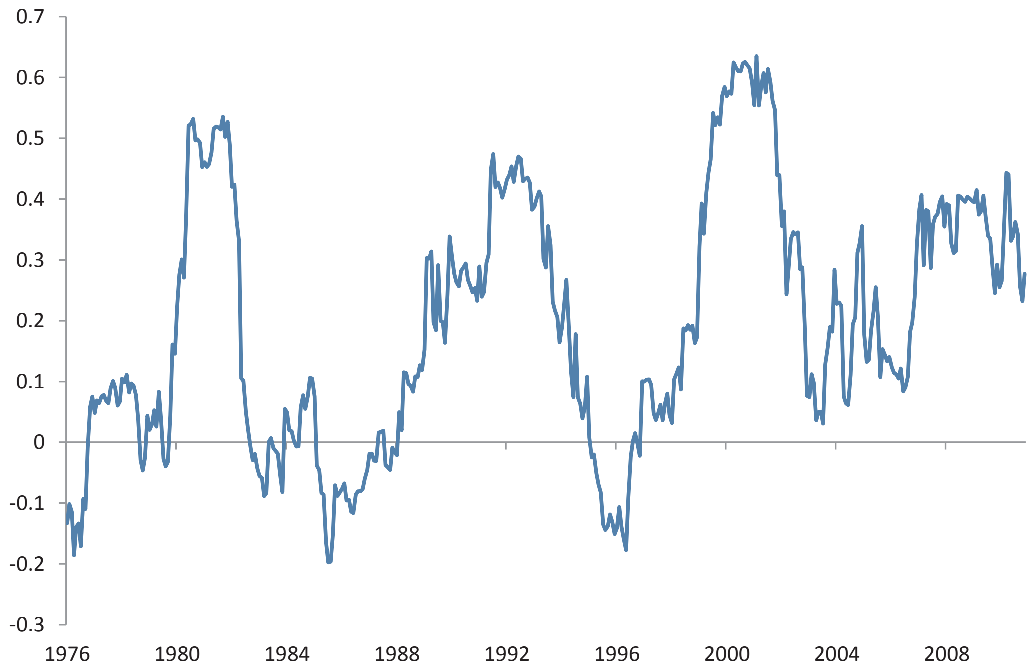

Source: MSCI, NBIM calculations

Comparing methodologies
The available return history for the investment process indices is short. Table 7 compares the risk and return of the investment process indices with the traditional style indices and the Fama-French benchmark series over the period 2003-2011. Table 7 and Charts 6-8 illustrate that:

- most specifications show a positive value premium over this (short) sample period, and
- the magnitude of the value premium and the risks related to capturing this premium vary considerably across different specifications.

Over this short time span, the traditional style indices also have lower Sharpe ratios than the Fama-French factors. The investment process indices, on the other hand, have Sharpe ratios comparable to the Fama-French factors. The high Sharpe ratios of these indices are, however, due to significantly lower volatility due to the extensive risk management process that is built into the portfolio construction process. The returns on these indices are in line with the traditional style indices and lower than the Fama-French returns.

Table 7: Value premia, 2003-2011

	Ret	Vol	Sharpe	Beta	Shortfall	Resilience
Theoretical Benchmarks						
Fama-French, Global	3.9 %	9.0 %	0.44	0.26	-4.2 %	-2.0 %
Fama-French, US	2.5 %	12.2 %	0.20	0.39	-6.7 %	-4.8 %
Fama-French, Europe	4.0 %	9.1 %	0.44	0.23	-4.2 %	-1.6 %
Fama-French, Pacific	7.3 %	6.6 %	1.10	-0.01	-2.9 %	1.7 %
Traditional Style Indexes, Developed Mkt						
MSCI World VMG	0.9 %	5.7 %	0.16	0.07	-3.5 %	-0.4 %
MSCI US VMG	0.3 %	7.3 %	0.04	0.03	-4.6 %	-1.4 %
MSCI Europe VMG	-0.2 %	6.1 %	-0.04	0.12	-3.2 %	-0.7 %
MSCI Pacific VMG	3.6 %	5.8 %	0.62	-0.12	-2.8 %	2.4 %
FTSE World VMG	1.5 %	4.6 %	0.32	0.02	-2.3 %	0.2 %
FTSE America VMG	0.8 %	2.9 %	0.26	0.00	-1.6 %	-0.3 %
FTSE Europe VMG	0.4 %	7.8 %	0.05	0.20	-4.1 %	-1.6 %
FTSE Pacific VMG	3.4 %	5.3 %	0.65	-0.11	-2.6 %	1.8 %
Russell Global Developed VMG	-0.1 %	4.9 %	-0.02	0.06	-3.1 %	-0.8 %
Investment Process Indexes						
FTSE All-World ActiveBeta Value Index	1.6 %	3.2 %	0.51	0.11	-1.7 %	-0.9 %
MSCI Europe Value Tilt Index	2.3 %	5.2 %	0.45	0.14	-3.2 %	-1.9 %
MSCI Europe BARRA Value Index	0.6 %	2.1 %	0.28	0.04	-1.5 %	-0.4 %
MSCI USA Barra Earnings Yield Index	0.6 %	2.6 %	0.24	0.01	-1.4 %	-0.1 %

Chart 6: Sharpe ratios for various specifications of the global value premium, 2003-2011

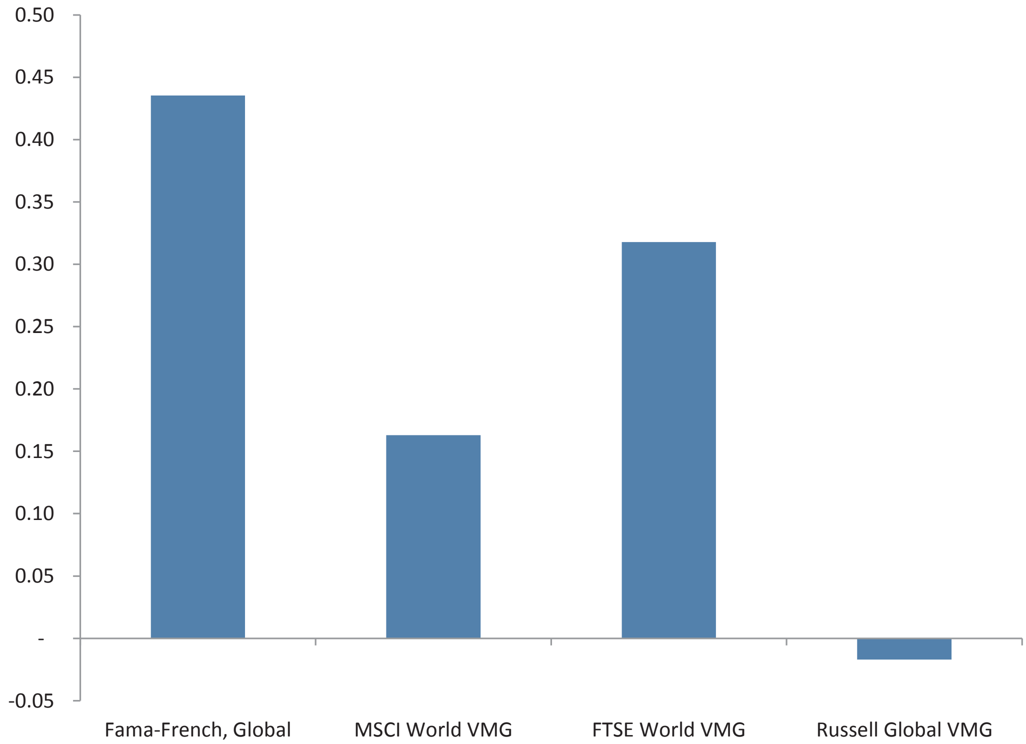

Source: MSCI, FTSE, Russell, Kenneth French's website, NBIM calculations

Chart 7: Sharpe ratios for various specifications of the European value premium, 2003-2011

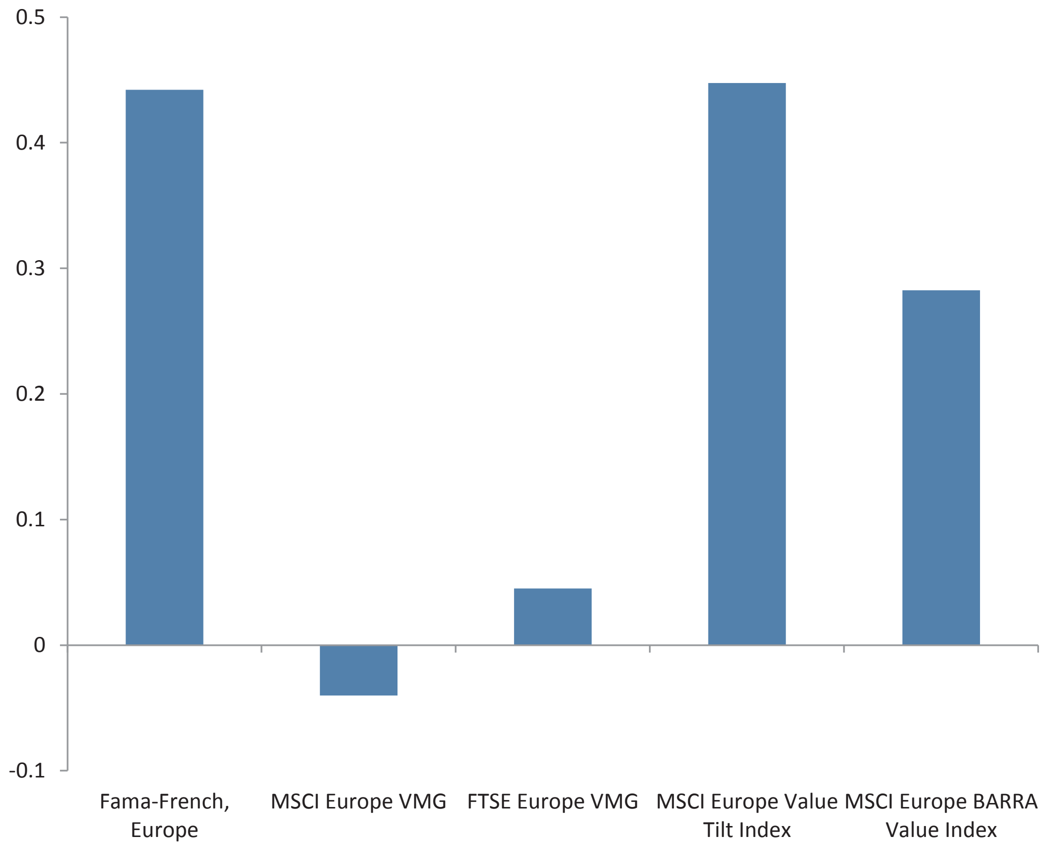

Source: MSCI, FTSE, Kenneth French's website, NBIM calculations

Chart 8: Sharpe ratios for various specifications of the US value premium, 2003-2011

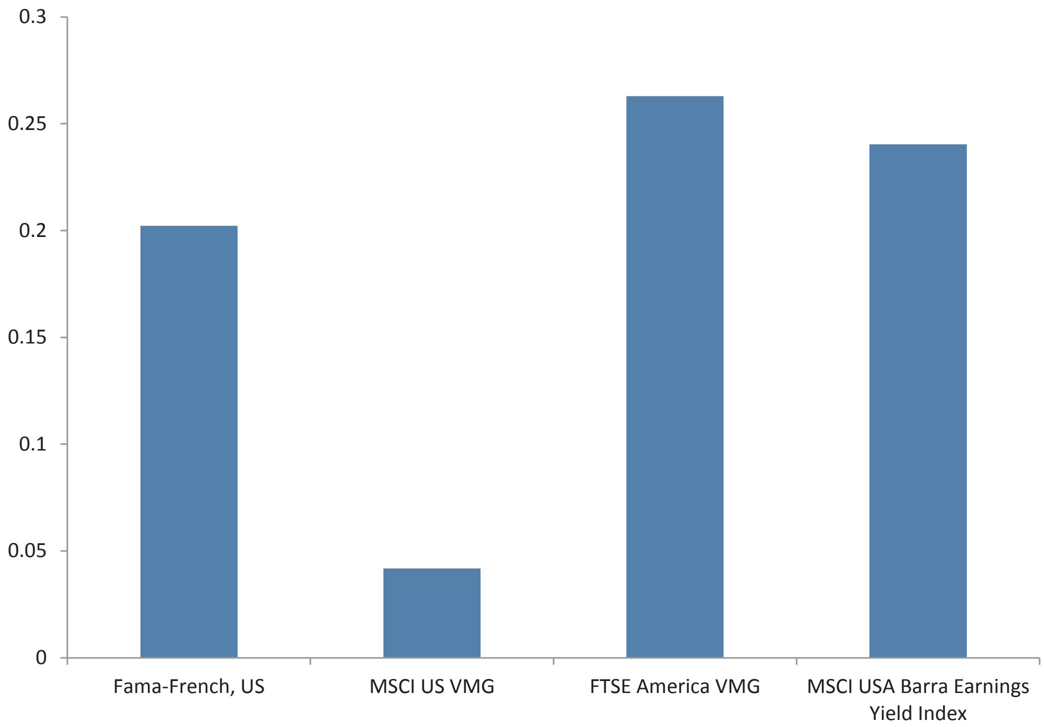

Source: MSCI, FTSE, Kenneth French's website, NBIM calculations

Imposing sector neutrality

As described above, neither Fama-French factors nor the traditional style indices adjust for sector, implicitly allowing potentially large sector tilts in the factor-mimicking portfolios. Based on our own value factor described above, we look now at the impact of neutralising any sector tilts that may arise in an unconstrained value factor.

Table 8 compares our value factor when we only adjust for region with a specification where the value factor is both region- and sector-neutral. The table shows that the sector-neutral specification gives a higher average return and lower volatility and drawdowns over the sample period. As a result, the Sharpe ratio is significantly higher for the sector-neutral specification. Chart 9 illustrates the cumulative value returns for the two specifications.

Table 8: Value premia in FTSE World large/mid-cap universe, 1994-2011

	nbimValue Region neutral	nbimValue Region & Sector neutral
Performance		
Ave Returns, Gross	6.8 %	7.3 %
Ave Returns, Net	6.2 %	6.7 %
Risk		
Volatility	11.3 %	8.1 %
Shortfall 5%	-7.1 %	-4.5 %
Max Drawdown, 6m	-35.4 %	-17.6 %
Sharpe		
Sharpe, Gross	0.60	0.90
Sharpe, Net	0.55	0.82
Beta	0.00	0.05
Resilience		
- Ave Ret when Mkt>0	0.4 %	0.6 %
- Ave Ret when Mkt<0	0.9 %	0.6 %
- Ave Ret when Mkt<VaR5%	-0.5 %	-0.4 %

Chart 9: Imposing sector neutrality. Cumulated value returns, 1994-2011

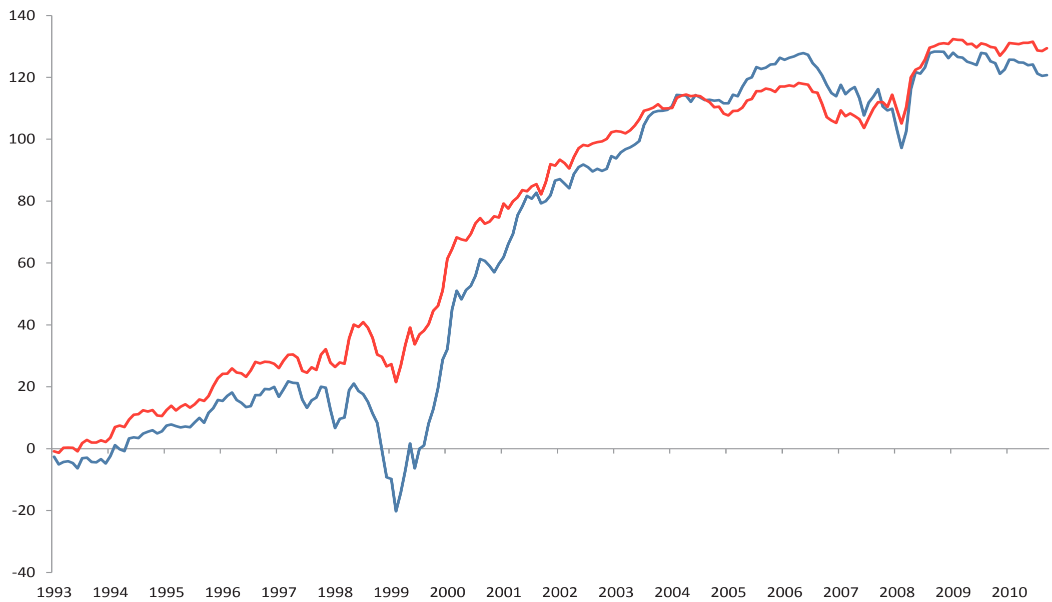

Source: FTSE, NBIM calculations

Rebalancing frequency

Table 8 reports the performance and risk of the two alternative value specifications assuming a monthly rebalancing frequency. We are also interested in the sensitivity to less frequent rebalancing, as less frequent rebalancing may have consequences both for the gross factor returns and for factor returns net of transaction costs. Table 9 reports the results.

The value factor is a slow-moving factor, in the sense that it normally takes time for a company to move from the long portfolio to the short portfolio. This means that the value factor does not require frequent rebalancing in order to capture the risk premium. This is confirmed in Table 9, as less frequent rebalancing seems to have a small impact on the Sharpe ratio net of transaction costs.

Table 9: Value premia with different rebalancing frequencies, 1994-2011

	Turnover	Ret	Ret Net	Vol	Sharpe	Sharpe Net
Monthly						
- Region neutral	196 %	6.8 %	6.2 %	11.3 %	0.60	0.55
- Region & Sector neutral		7.3 %	6.7 %	8.1 %	0.90	0.82
Quarterly						
- Region neutral	141 %	6.6 %	5.9 %	13.3 %	0.50	0.45
- Region & Sector neutral		6.4 %	5.9 %	9.1 %	0.70	0.65
Semiannually						
- Region neutral	120 %	5.2 %	4.9 %	14.8 %	0.35	0.33
- Region & Sector neutral		5.0 %	4.6 %	9.1 %	0.55	0.51
Annually						
- Region neutral	101 %	5.3 %	5.0 %	13.6 %	0.39	0.37
- Region & Sector neutral		5.0 %	4.7 %	9.2 %	0.54	0.51

Chart 10: Sharpe ratios net of costs for value strategies with different rebalancing frequencies, 1994-2011

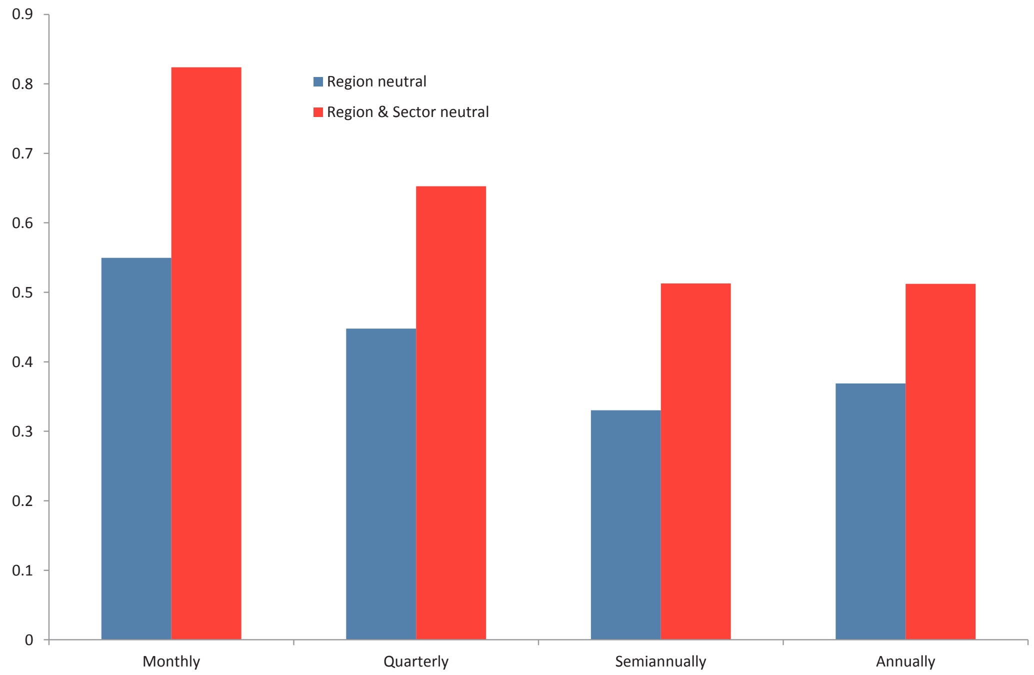

Source: FTSE, NBIM calculations

Small cap

Small-cap premia in different regions

Table 10 summarises the small-cap factor returns based on the Fama-French methodology, globally as well as for the three different regions, for the period 2003-2011 (unless otherwise stated). We compare the Fama-French factors with the MSCI and FTSE style indices. MSCI World SmL ("Small minus Large") represents the difference in returns between a long exposure in the MSCI World small-cap index and a short exposure in the MSCI World large-cap index. The FTSE SmL factor is constructed similarly.

As for the value premium, we see that the average returns and the Sharpe ratios are consistently lower in all regions compared to the Fama-French benchmark series. Moreover, Table 10 shows that there are important differences in performance and risk across different regions.

Table 10: Small-cap premia in different regions, 2003-2011

	Ret	Vol	Sharpe	Beta	Shortfall	Resilience
Theoretical Benchmark						
Fama-French, US, 1926-2011	3.0 %	11.2 %	0.27	0.20	-5.9 %	-2.0 %
Fama-French, US, 1975-2011	4.0 %	10.3 %	0.39	0.10	-5.8 %	-1.8 %
Fama-French, US	4.9 %	29.5 %	0.17	0.20	-4.0 %	-0.5 %
Traditional Style Indexes, Global						
MSCI World Small vs Large	5.3 %	24.0 %	0.22	0.18	-3.6 %	-1.4 %
FTSE World Small vs Large	3.5 %	21.2 %	0.16	0.17	-3.7 %	-1.2 %
Traditional Style Indexes, US						
MSCI US Small vs Large	3.4 %	31.5 %	0.11	0.31	-4.5 %	-2.0 %
FTSE Americas Small vs Large	3.9 %	27.3 %	0.14	0.24	-4.2 %	-1.7 %
Russell 2000 vs 1000	1.8 %	29.1 %	0.06	0.24	-4.0 %	-1.7 %
Traditional Style Indexes, Europe						
MSCI Europe Small vs Large	4.6 %	31.4 %	0.15	0.14	-5.3 %	-2.2 %
FTSE Europe Small vs Large	4.2 %	27.3 %	0.15	0.09	-4.8 %	-2.2 %
Traditional Style Indexes, Pacific						
MSCI Asia/Pacific Small vs Large	2.0 %	24.9 %	0.08	0.00	-4.1 %	0.3 %
FTSE Asia/Pacific Small vs Large	2.6 %	26.0 %	0.10	0.12	-4.1 %	-1.0 %
Traditional Style Indexes, Emerging Markets						
MSCI Emerging Markets, Small vs Large	1.7 %	27.1 %	0.06	0.09	-4.3 %	-1.3 %

Chart 11: Small-cap premia in different regions. Cumulated returns, 2001-2011

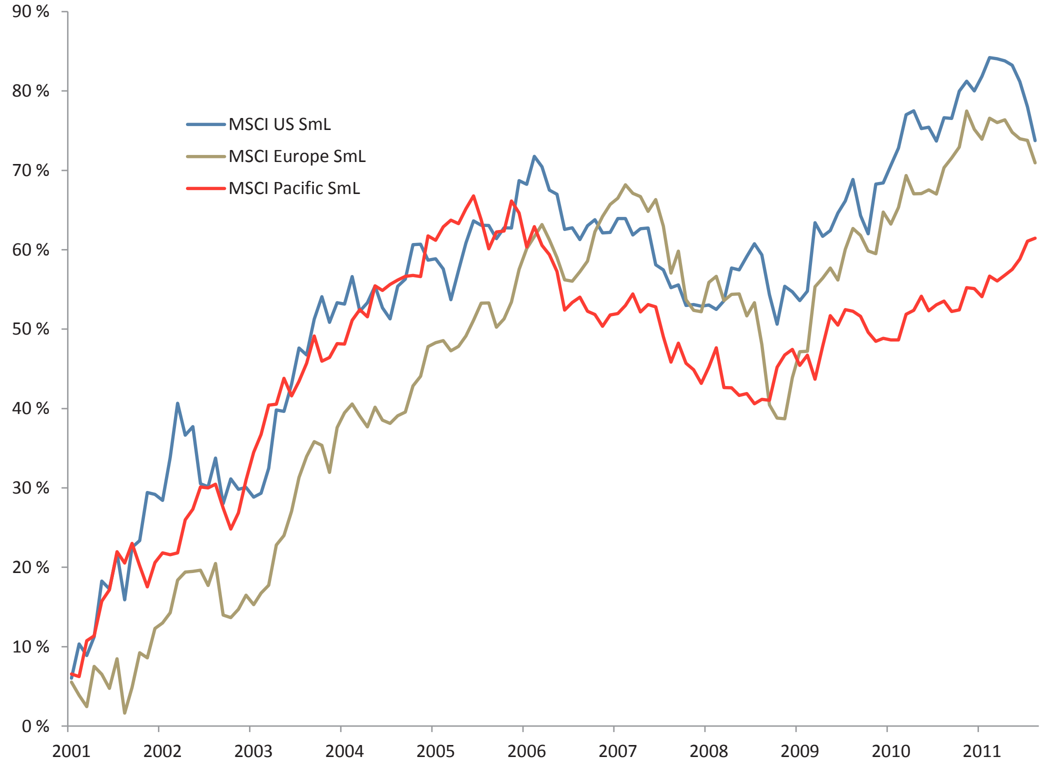

Source: MSCI, NBIM calculations

Mid-cap premia versus small-cap premia

The small-cap segment is, by definition, populated by small and illiquid companies with limited investment capacity. The question then is whether it is possible to capture a similar premium by investing in the mid-cap segment rather than the small-cap segment.

Table 11 summarises the return and risk for the mid-cap premium. As Chart 12 illustrates, the average return for the mid-cap factor is lower than for the small-cap factor. However, the volatility in the mid-cap factor is also lower. Hence, when we look at the Sharpe ratios in Chart 13, the differences are not that large. The conclusion is that the size effect is also present in the mid-cap segment.

Table 11: Mid-cap premia in different regions, 2001-2011

	Ret	Vol	Sharpe	Beta	Shortfall	Resilience
Traditional Style Indexes, Global						
MSCI World Mid vs Large	3.5 %	16.2 %	0.21	0.13	-2.9 %	-1.4 %
FTSE World Mid vs Large	3.2 %	14.8 %	0.22	0.11	-2.4 %	-1.1 %
Traditional Style Indexes, US						
MSCI US Mid vs Large	3.6 %	23.1 %	0.15	0.19	-3.9 %	-1.8 %
FTSE Americas Mid vs Large	3.5 %	19.9 %	0.18	0.17	-3.1 %	-1.5 %
Traditional Style Indexes, Europe						
MSCI Europe Mid vs Large	1.9 %	19.9 %	0.09	0.07	-3.8 %	-1.3 %
FTSE Europe Mid vs Large	2.0 %	19.5 %	0.10	0.07	-3.5 %	-1.3 %
Traditional Style Indexes, Pacific						
MSCI Asia/Pacific Mid vs Large	1.8 %	18.5 %	0.10	0.02	-2.6 %	-0.6 %
FTSE Asia/Pacific Mid vs Large	1.4 %	16.8 %	0.09	0.01	-2.4 %	-0.2 %
Traditional Style Indexes, Emerging Markets						
MSCI Emerging Markets, Mid vs Large	2.2 %	20.6 %	0.11	0.08	-2.9 %	-1.3 %

Chart 12: Mid-cap premia in different regions. Cumulated returns, 2001-2011

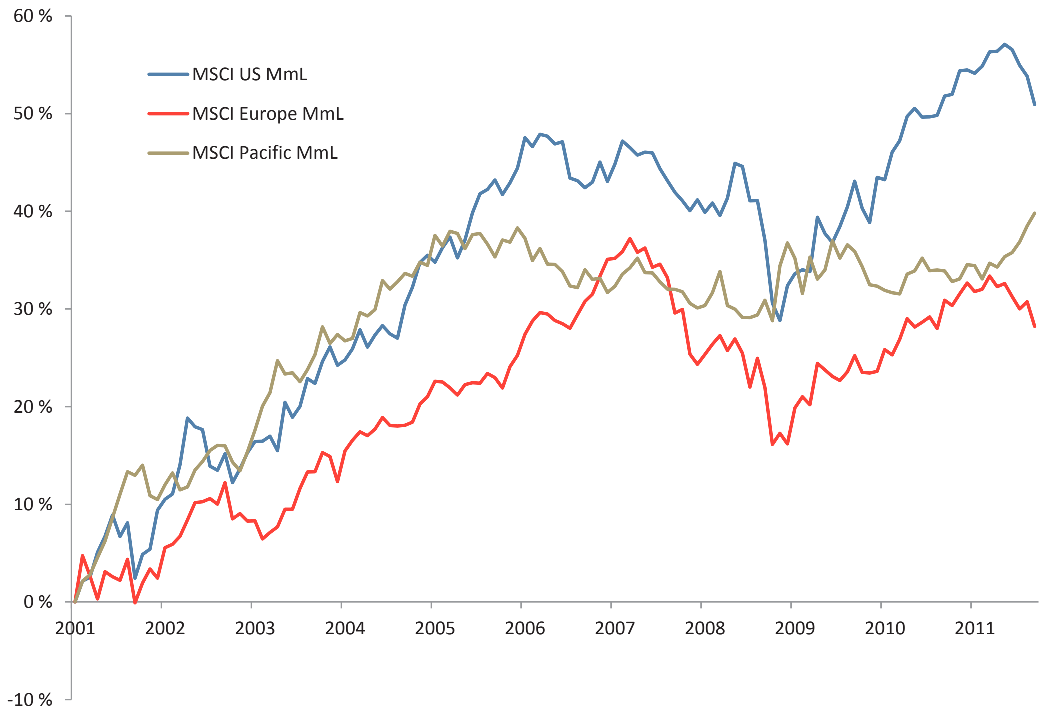

Source: MSCI, NBIM calculations

Chart 13: Comparing small-cap and mid-cap factor returns, 2001-2011

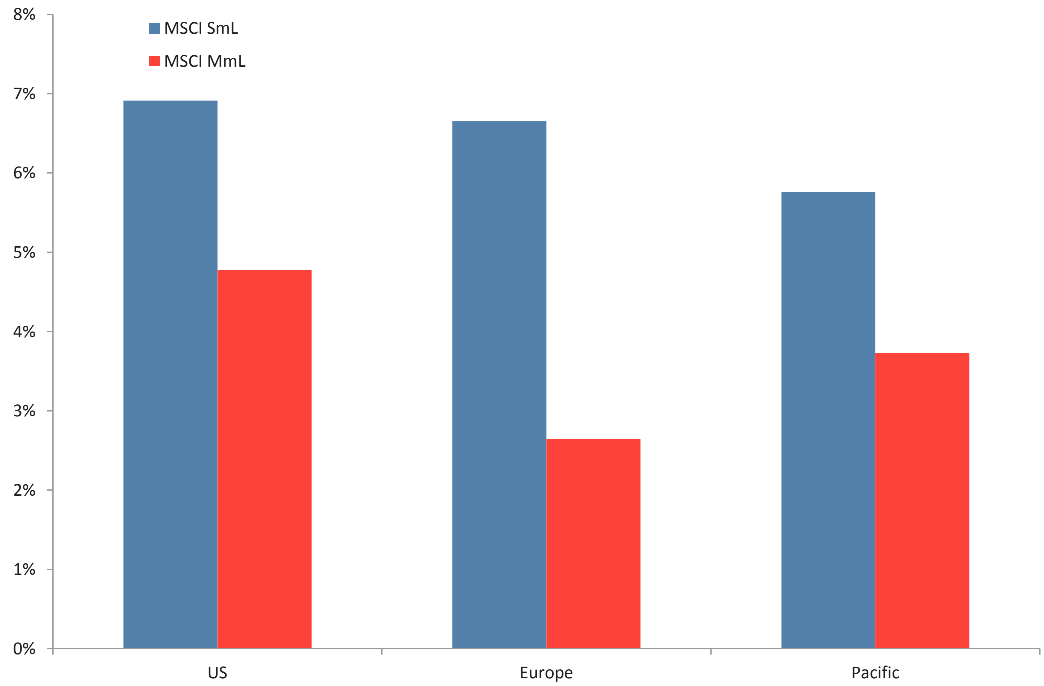

Source: MSCI, NBIM calculations

Chart 14: Comparing Sharpe ratios for small-cap and mid-cap factor returns, 2001-2011

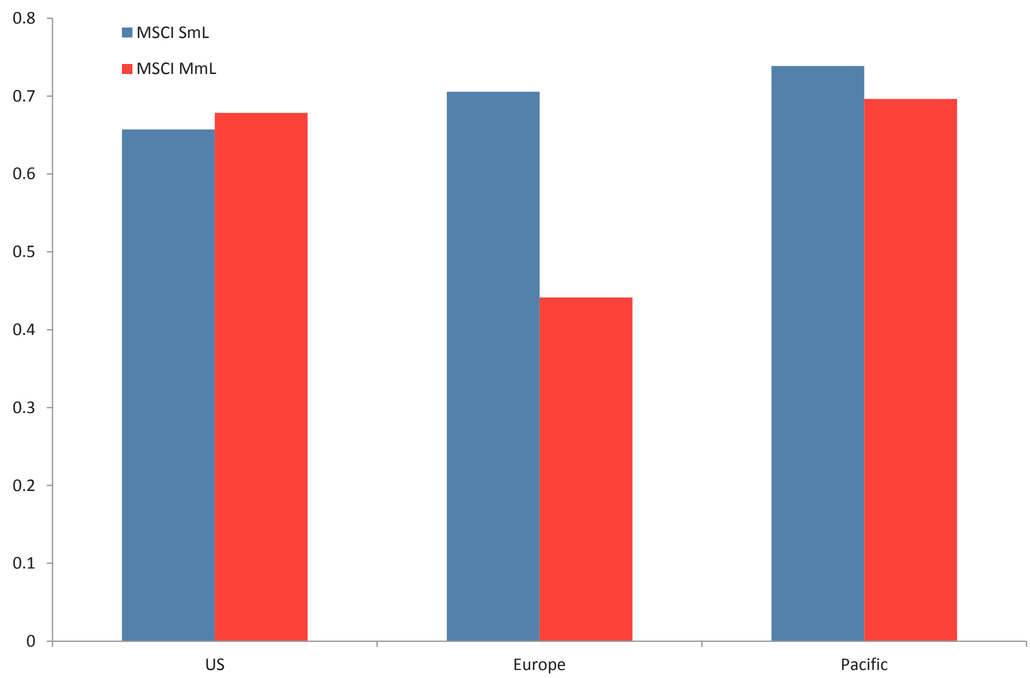

Source: MSCI, NBIM calculations

Imposing sector neutrality
We use our own size factor in order to analyse the impact of imposing sector neutrality. As explained above, we calculate the size premium using the FTSE World large/mid-cap universe. Hence, this specification aims at capturing the mid-cap premium rather than the small-cap premium. As for the value premium, imposing sector neutrality on the size strategy increases returns and reduces both volatility and drawdowns. Moreover, the beta tilt is more than halved.

Table 12: Size premia in FTSE World large/mid-cap universe, 1994-2011

	nbimSize Region neutral	nbimSize Region & Sector neutral
Performance		
Ave Returns, Gross	3.9 %	4.3 %
Ave Returns, Net	3.4 %	3.9 %
Risk		
Volatility	10.0 %	7.2 %
Shortfall 5%	-6.2 %	-4.5 %
Max Drawdown, 6m	-24.8 %	-13.5 %
Sharpe		
Sharpe, Gross	0.39	0.60
Sharpe, Net	0.35	0.54
Beta	0.11	0.05
Resilience		
- Ave Ret when Mkt>0	0.5 %	0.4 %
- Ave Ret when Mkt<0	0.1 %	0.4 %
- Ave Ret when Mkt<VaR5%	-1.8 %	-0.9 %

Chart 15: Imposing sector neutrality. Cumulated returns, 1994-2011

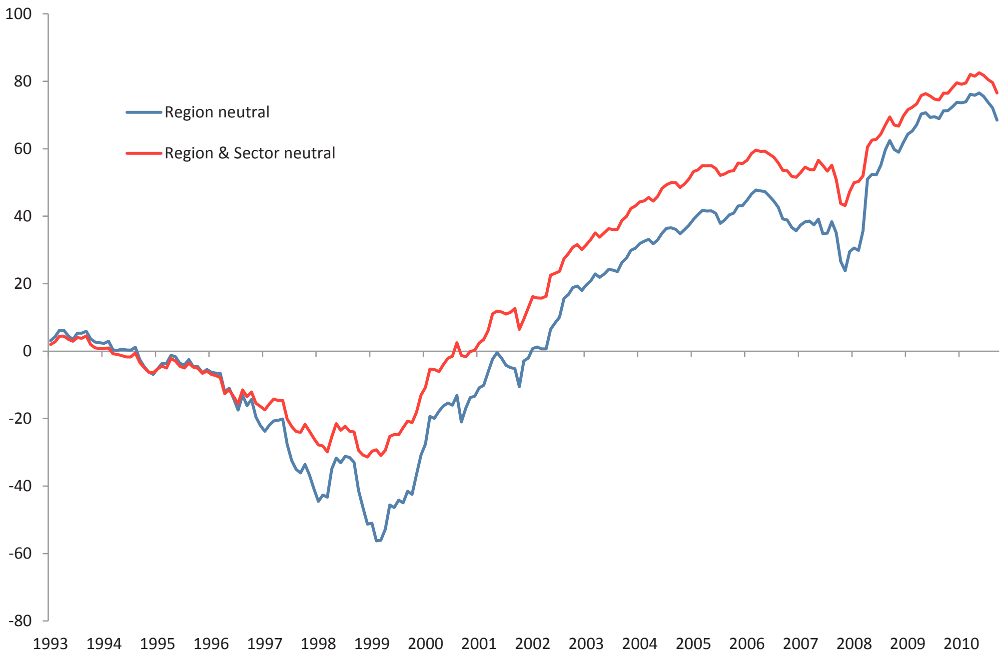

Source: FTSE, NBIM calculations

Rebalancing frequency

The size factor is also a slow-moving factor in the sense that it requires low turnover. As a result, the factor returns are resilient to the rebalancing frequency, see Table 9.

Table 13: Small-cap premia with different rebalancing frequencies, 2001-2011

	Turnover	Ret	Ret Net	Vol	Sharpe	Sharpe Net
Monthly						
- Region neutral	139 %	3.9 %	3.4 %	10.0 %	0.39	0.35
- Region & Sector neutral		4.3 %	3.9 %	7.2 %	0.60	0.54
Quarterly						
- Region neutral	92 %	4.0 %	3.7 %	11.4 %	0.35	0.33
- Region & Sector neutral		4.1 %	3.8 %	7.5 %	0.54	0.50
Semiannually						
- Region neutral	71 %	3.5 %	3.3 %	12.5 %	0.28	0.27
- Region & Sector neutral		3.8 %	3.6 %	8.5 %	0.45	0.42
Annually						
- Region neutral	59 %	4.2 %	4.0 %	15.6 %	0.27	0.26
- Region & Sector neutral		3.8 %	3.7 %	11.5 %	0.33	0.32

Chart 16: Comparing Sharpe ratios net of costs for different rebalancing frequencies, 2001-2011

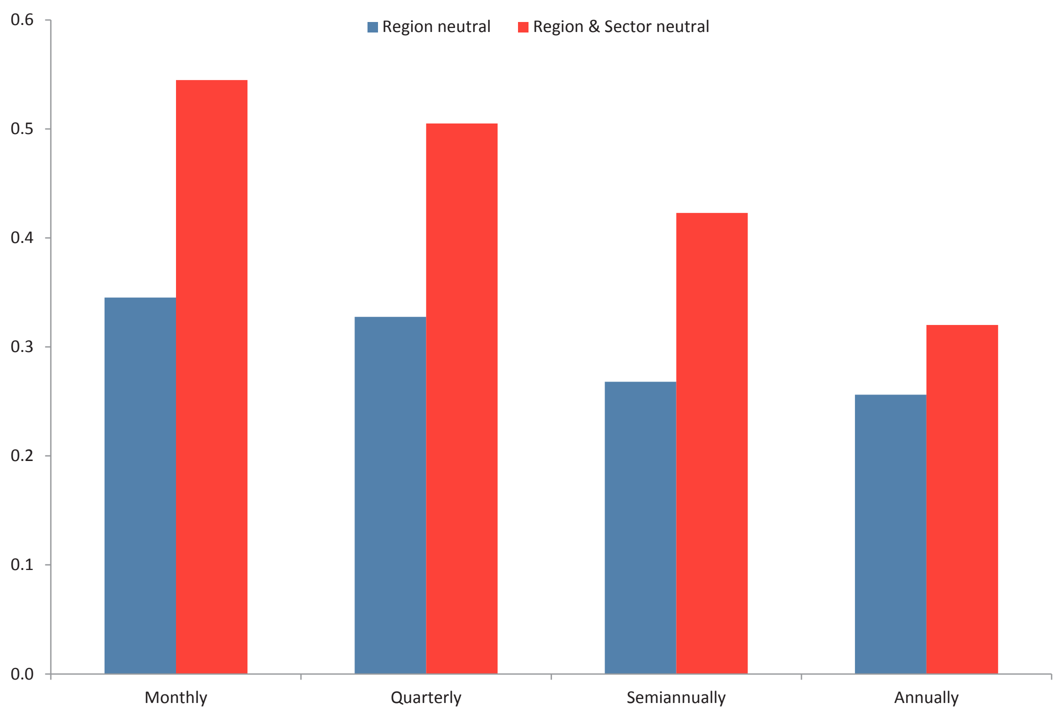

Source: FTSE, NBIM calculations

Momentum

Comparing methodologies

Momentum strategies are characterised by a high frequency of rebalancing and high turnover, features considered undesirable from the perspective of passive, low-turnover, low-cost indices. As a result, the number of publicly available momentum indices is low.

The Fama-French momentum factor has an impressive performance over the last century, as shown in Chart 1 above. The main objection to the momentum factor, however, is that it is hard to capture it once investability and trading costs are taken into account. Compared to the Fama-French factor, the investment process indices include measures to limit turnover in addition to taking into account various risk considerations. This reduces the magnitude of the momentum premium, but also reduces risk significantly. As a result, the Sharpe ratios for the investment process indices hold up very well compared to the Fama-French momentum factor, see Table 14. The returns reported in Table 14 are gross returns. In order to take transaction costs into account, we have to resort to our internally calculated momentum factor, which will be described below.

Focusing on the last decade, the momentum strategy has been less profitable than in previous decades. The Fama-French momentum factor actually has a negative return in the period 2003-2011, whereas the return in the period 1975-2011 was close to 8 percent on average per year. This is mostly due to the drawdown period in 2009 when the market recovered from its drawdown in 2008. The weak momentum performance in 2009 shows up in the investment process indices, as can be seen from Chart 16. We will revert to the reasons for this underperformance below.

Table 14: Momentum premia in different regions, 2003-2011

	Ret	Vol	Sharpe	Beta	Shortfall	Resilience
Theoretical Benchmark						
Fama-French, US, 1926-2011	8.4 %	16.6 %	0.50	-0.30	-12.1 %	3.9 %
Fama-French, US, 1975-2011	7.9 %	15.7 %	0.50	-0.13	-11.4 %	2.1 %
Fama-French, US	-2.5 %	18.2 %	-0.14	-0.30	-12.1 %	3.9 %
Investment Process Indexes						
FTSE All-World ActiveBeta Momentum Index	0.9 %	5.6 %	0.16	-0.07	-3.4 %	0.3 %
MSCI USA Barra Momentum Index	0.7 %	4.2 %	0.18	-0.06	-3.3 %	0.9 %
MSCI Europe Barra Momentum Index	3.0 %	4.8 %	0.62	-0.02	-3.9 %	-0.4 %
MSCI Europe Momentum Tilt Index	2.7 %	4.0 %	0.66	-0.02	-2.2 %	0.5 %

Chart 17: Momentum factors. Cumulated returns, 2003-2011

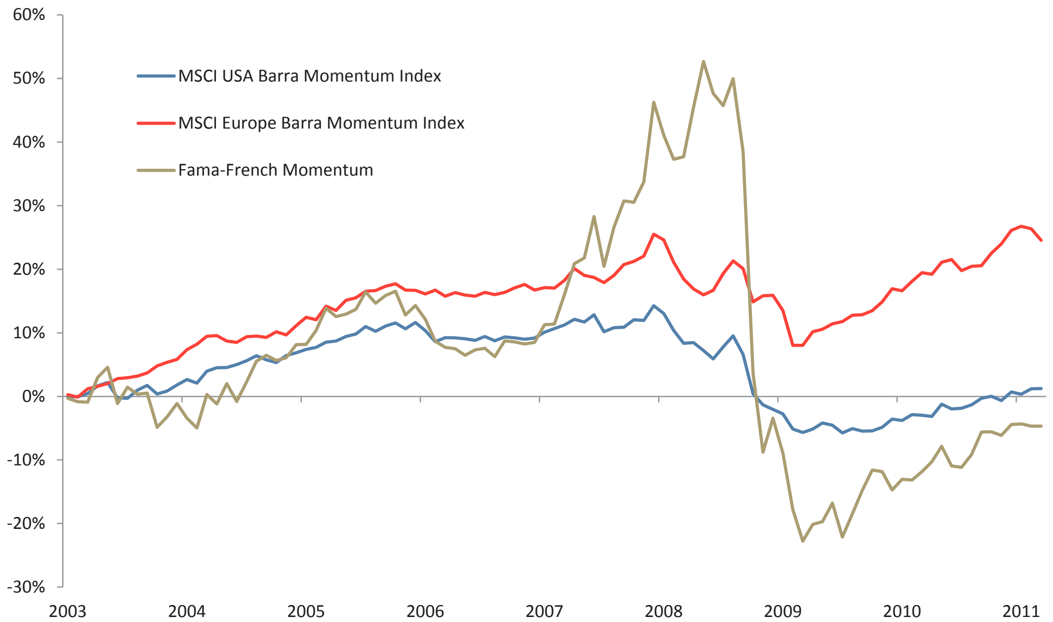

Source: MSCI, Kenneth French's website, NBIM calculations

Rebalancing frequency
As mentioned, the crucial point regarding the momentum factor is whether it is robust to transaction costs. Based on our internal momentum factor, we have calculated the momentum premium net of transaction costs under various assumptions regarding rebalancing frequencies, see Table 15.

As expected, turnover is massively higher than for the value and small-cap premia. The gross momentum premium with monthly rebalancing is estimated at 5.5 percent per year over our sample period 1994-2011. However, after taking transaction costs into account, the average momentum return drops to 4 percent per year, giving a Sharpe ratio of 0.23. Less frequent rebalancing obviously reduces turnover and transaction costs, but the gross return also becomes smaller.

Table 15: Momentum premia with different rebalancing frequencies, 1994-2011

	Turnover	Ret	Ret Net	Vol	Sharpe	Sharpe Net
Monthly						
- Region neutral	493 %	5.5 %	4.0 %	17.4 %	0.32	0.23
- Region & Sector neutral		2.7 %	1.3 %	11.8 %	0.23	0.11
Quarterly						
- Region neutral	363 %	3.2 %	2.2 %	17.3 %	0.19	0.12
- Region & Sector neutral		2.3 %	1.2 %	12.2 %	0.19	0.10
Semiannually						
- Region neutral	299 %	2.8 %	1.9 %	13.7 %	0.20	0.14
- Region & Sector neutral		2.0 %	1.1 %	9.9 %	0.20	0.11
Annually						
- Region neutral	241 %	0.0 %	-0.6 %	18.6 %	0.00	-0.03
- Region & Sector neutral		0.5 %	-0.2 %	13.5 %	0.04	-0.01

Chart 18: Sharpe ratios net of costs for different momentum strategies, 1994-2011

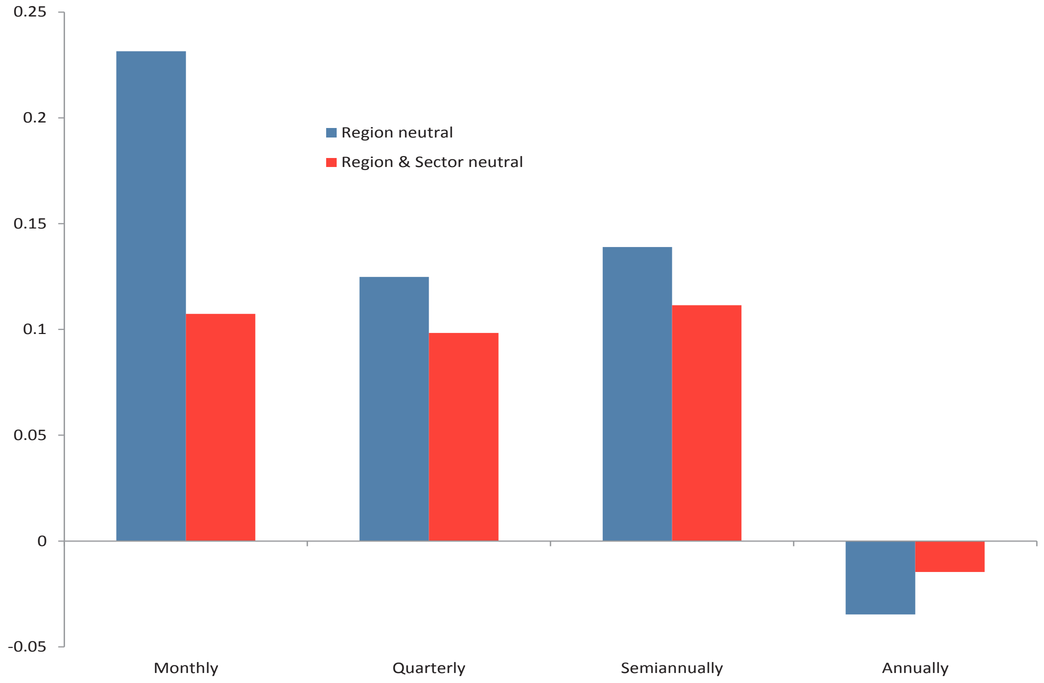

Source: FTSE, NBIM calculations

Imposing sector neutrality
According to academic research, the momentum effect is to a great extent related to sector effects (Moskowitz and Grinblatt 1999). As a result, we expect there to be large differences between a sector-neutral specification and a non-sector-neutral specification. Table 16 and Chart 17 confirm this. The momentum returns are much higher if we accept sector tilts in the momentum portfolios. The volatility and drawdowns will also be higher, but the Sharpe ratio is still higher for the non-sector-neutral specification.

Table 16: Imposing sector neutrality, 1994-2011

	nbimMomentum Region neutral	nbimMomentum Region & Sector neutral
Performance		
Ave Returns, Gross	5.5 %	2.7 %
Ave Returns, Net	4.0 %	1.3 %
Risk		
Volatility	17.4 %	11.8 %
Shortfall 5%	-12.3 %	-8.7 %
Max Drawdown, 6m	-47.7 %	-38.1 %
Sharpe		
Sharpe, Gross	0.32	0.23
Sharpe, Net	0.23	0.11
Beta	-0.33	-0.19
Resilience		
- Ave Ret when Mkt>0	-0.2 %	-0.1 %
- Ave Ret when Mkt<0	1.4 %	0.8 %
- Ave Ret when Mkt<VaR5%	4.7 %	2.0 %

Chart 19: Cumulated momentum returns, 1994-2011

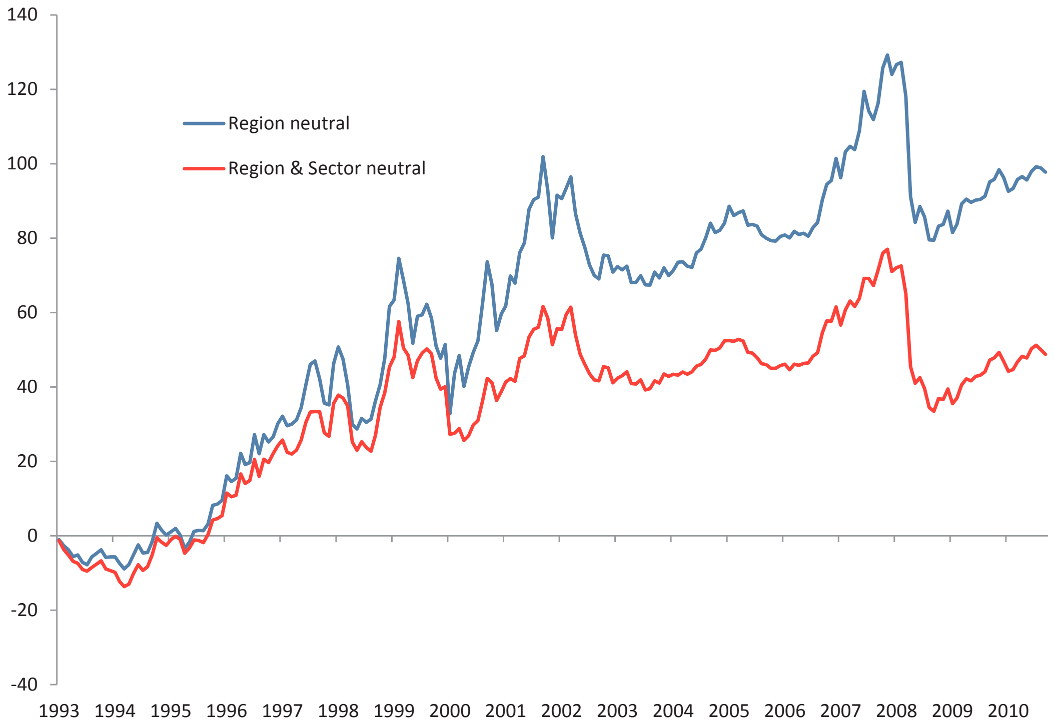

Source: FTSE, NBIM calculations

Volatility

There are a number of non-market-capitalisation-based indices that have significant tilts towards low-volatility stocks where the performance partly reflects the volatility premium. These indices are addressed in a separate note and are not discussed here. Instead, we focus on our internally-developed volatility strategy based on the methodology in Frazzini and Pedersen (2010). We compare this strategy to the performance and risk of the low-volatility factor index calculated by MSCI, which belongs to the same family of indices as the factor indices discussed above. An important difference between our volatility strategy and the MSCI strategy is that our strategy imposes ex-ante beta neutrality on the factor-mimicking portfolio, in addition to region and sector neutrality, but no restrictions with respect to other sources of systematic risk. MSCI minimises exposures to a broad set of systematic risk factors, including market beta, but allows some leeway. As a result, the MSCI indices have a bigger negative beta tilt than our internal volatility strategy, as can be seen from Table 17. As a result of the optimisation process built into the MSCI portfolio construction process, both volatility and drawdowns are smaller for the MSCI indices. However, tighter risk management also leads to lower returns.

Our results indicate that the beta-neutral strategy succeeds in capturing a volatility premium of more than 3 percent per year on average net of transaction costs. This is lower than the returns reported by Frazzini and Pedersen (2010). This may be due partly to the strong performance of the volatility factor prior to 1994.

Chart 19 plots the cumulated returns on our beta-neutral volatility strategy along with the cumulated returns on the FTSE World index and a simple volatility strategy. The simple volatility strategy is constructed by overweighting the least volatile stocks and underweighting the most volatile stocks on a cash-neutral basis. As a result, this simple volatility strategy has a significant negative beta tilt, which reduces the performance significantly relative to the beta-neutral strategy. The chart also shows that the beta-neutral strategy has a big drawdown in 2009. However, as illustrated in Chart 20, this is not related to a beta tilt in the strategy, as the realised beta is close to 0 in 2008 and 2009 when the market drops and recovers. We will revert to this issue below.

Table 17: Risk and return of volatility premium, 1994-2011

	nbimVolatility Monthly	nbimVolatility Quarterly	Frazzini & Pedersen (2010)1)	MSCI Europe Barra Low Volatility Index2)	MSCI Europe USA Low Volatility Index2)
Performance					
Ave Returns, Gross	4.4 %	3.9 %	8.6 %	1.3 %	1.5 %
Ave Returns, Net	3.3 %	3.2 %			
Risk					
Volatility	7.3 %	8.1 %	10.9 %	4.9 %	6.6 %
VaR 5%	-3.2 %	-2.0 %		-1.9 %	-2.2 %
CVaR 5%	-4.9 %	-3.0 %		-2.2 %	-3.9 %
Max Drawdown, 6m	-23.8 %	-19.7 %		-6.6 %	-13.6 %
Sharpe					
Sharpe, Gross	0.60	0.48	0.79	0.27	0.23
Sharpe, Net	0.45	0.39			
Beta (realized)	-0.01	-0.03	0.02	-0.12	-0.25
Resilience					
- Ave Ret when Mkt>0	0.32 %	0.33 %		-0.3 %	-0.6 %
- Ave Ret when Mkt<0	0.43 %	0.31 %		0.7 %	1.5 %
- Ave Ret when Mkt<VaR5%	-0.82 %	0.51 %		2.1 %	2.6 %

1) BAB-factor for global equities 1984-2009, monthly rebalancing
2) Period 2003-2011

Chart 20: Cumulated volatility returns, 1994-2011

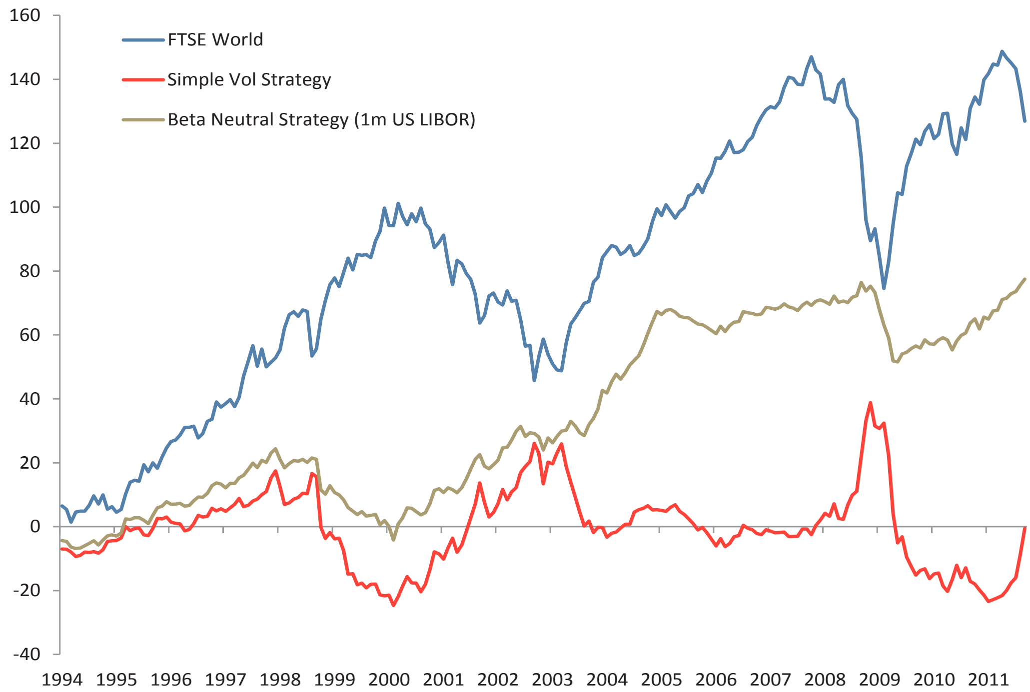

Source: FTSE, NBIM calculations

Chart 21: Cumulated volatility returns and realised beta, 1994-2011

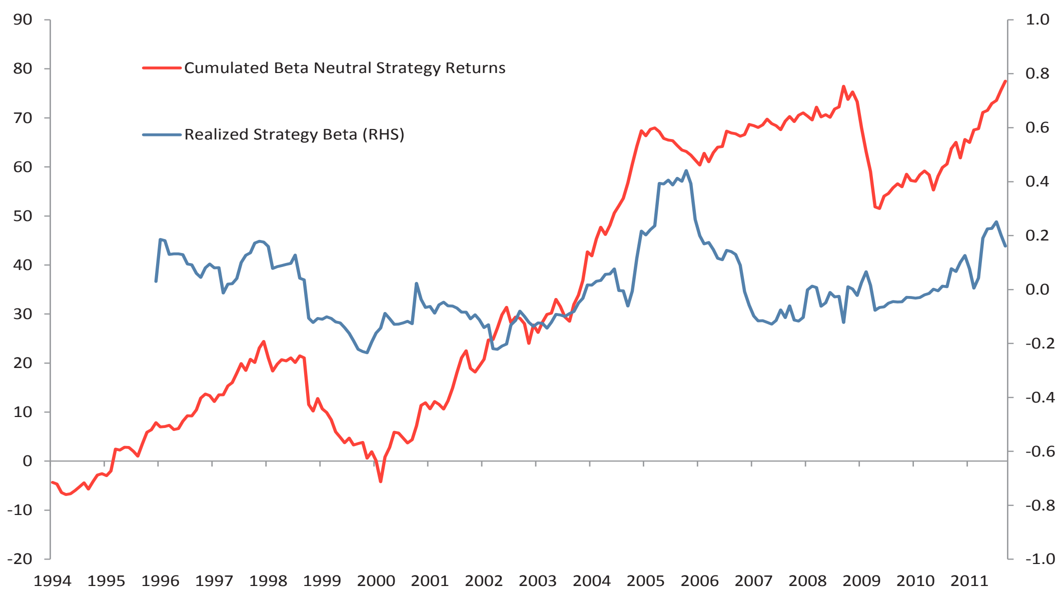

Source: FTSE, NBIM calculations

A portfolio of risk premia

There are several reasons why the systematic risk premia in the equity market should be captured within a unified framework for all systematic risk premia. As pointed out above, all risk premia have positive expected returns over a long horizon, and the correlations between the different risk premia are generally low. A combined capture of several risk premia will therefore give a higher Sharpe ratio and lower drawdowns. In addition, each risk premium has indirect exposures to other sources of systematic risk, such as liquidity risk. A combined capture of several risk premia enables better management of both direct and indirect risk exposures as well as turnover. We will address these points in more detail below.

Liquidity premia in factor returns
A large empirical literature documents the impact of illiquidity on returns. Illiquid assets provide, on average, higher returns than liquid assets, and this excess return is often called the liquidity premium. The liquidity premium arises for several reasons. Investors may require a premium in order to cover higher transaction costs related to illiquid stocks. However, academic studies show that the liquidity premium may not just be a compensation for the characteristic of being illiquid, but also for an asset's sensitivity to fluctuations in aggregate liquidity conditions, called liquidity risk. The two features, which come from viewing illiquidity as both a cost and a risk, are closely related and hard to disentangle.

Liquidity premia are hard to capture for several reasons. First, liquidity is hard to define and measure. Academic research focuses on a number of different aspects of liquidity, and each aspect can be modelled using different specifications (see survey in Ilmanen 2011). Second, even given a specification of liquidity, it is hard to isolate the illiquidity effect from other sources of systematic risk.

Amihud and Mendelson (1986) show that liquidity is priced as a characteristic in the sense that less liquid stocks have higher expected gross returns. The authors argue that investors price assets based on their expected net (after-cost) returns. Heterogeneous trading horizons and limited capital lead to clientele effects where short-horizon investors focus on the most liquid assets while long-horizon investors focus on less liquid assets. Long-horizon investors can earn rents because patient capital is in short supply. Empirical evidence from the US supports the authors' predictions, as stocks with higher bid-ask spreads earn higher average returns.

Liquidity conditions vary over time and liquidity shocks exhibit commonality across assets, suggesting that liquidity may be a priced risk factor. This means that expected returns may be related to stocks' sensitivity to shifts in aggregate liquidity – and not just to the liquidity of the stock per se. Pastor and Stambaugh (2003) focus on liquidity as a risk factor, where they construct a measure of market-wide liquidity based on temporary price fluctuations induced by order flow. The authors show that stocks with high sensitivity to this liquidity factor outperform stocks with low liquidity betas after adjusting for other sources of systematic risk.

Academic research points out that several of the more well-known systematic risk factors, such as value, size and volatility, are linked to liquidity. In particular, Amihud and Mendelson (1986, 1989) find that the size effect disappears after controlling for bid-ask spread. Brennan, Chordia and Subrahmanyam (1998, 2005) and Chordia, Subrahmanyam and Anshuman (2001) find that the size-return relation becomes either insignificant or positive after controlling for trading volume. These results suggest that the size effect may be fully explained by liquidity. Other studies find evidence of a liquidity premium also after controlling for size and other sources of systematic risk (Pastor and Stambaugh 2003; Ibbotson, Chen and Hu 2011). A difficulty in testing the liquidity-based explanation against the risk-based explanation lies in the fact that stock liquidity is very hard to measure and is usually inextricably correlated with firm size. In fact, the market capitalisation of equity, a standard measure of firm size, is also frequently used as a measure of stock liquidity (Kluger and Stephan 1997). As a result, it is very difficult to disentangle the genuine size effect and the liquidity effect.

Furthermore, Asness, Moskowitz and Pedersen (2009) document liquidity effects in value and momentum returns. The authors investigate value and momentum strategies for stock selection within a broad universe of stocks as well as for country selection, government bond selection, currency selection and commodity selection. In addition to documenting the existence and co-movement of value and momentum premia in all asset classes, they seek to link the premia to underlying risks such as macroeconomic indicators and indicators of liquidity risk. They find a weak relationship between the risk premia and the macroeconomic indicators. To explore the role played by liquidity risk, the authors regress the value and momentum premia on a wide range of different funding liquidity risk indicators, such as the US Treasury-Eurodollar (TED) spread and a global average of TED spreads as well as other funding liquidity measures used in the literature by Pastor and Stambaugh (2003), Sadka (2006), Acharya and Pedersen (2005), Adrian and Shin (2007) and Krishnamurthy and Vissing-Jørgensen (2008). In addition, the authors compute an illiquidity index that takes a weighted average of all these measures. The use of the TED spread as a measure of funding liquidity is motivated by Brunnermeier and Pedersen (2008), who show that funding liquidity is a natural driver of common market liquidity risk across asset classes and markets. Also, Moskowitz and Pedersen (2008) show empirically that funding liquidity measures based on TED spreads are linked to the relative returns of liquid versus illiquid securities globally. For both levels and changes in liquidity indicators, the authors find a consistent pattern among value and momentum strategies. Specifically, value loads positively on liquidity risk, whereas momentum loads either negatively or zero on liquidity risk, depending on the measure. Put another way, value strategies do worse when liquidity is poor and worsening, and momentum seems to do better during these times. One reason for this may be that value companies typically have higher leverage. This is consistent with what we find in our internally-calculated value strategy, see Chart 21.

This means that the illiquidity exposure of value strategies explains part of the value premium. Moreover, the different loading on liquidity risk for value and momentum premia may help explain part of why value and momentum strategies are negatively correlated. However, the negative illiquidity exposure in the momentum premium only deepens the puzzle presented by its high returns. Momentum strategies seem to do better when the market is illiquid, presumably a characteristic investors would pay for in terms of lower expected returns.

Frazzini and Pedersen (2010) investigate the illiquidity exposure in their BAB factor along the same lines, using the level of and change in the TED spread as indicators of funding liquidity risk. They find a significant relationship between liquidity risk and the return on the BAB factor. Chart 22, which plots the cumulated return on our internally-calculated volatility factor together with the TED spread, indicates a similar relationship. The drawdown in this factor in 2009 is not directly related to the market beta, as explained above, but more to the change in the liquidity risk proxied by the TED spread.

Chart 22: Exposure to leverage in value factor, 1994-2011

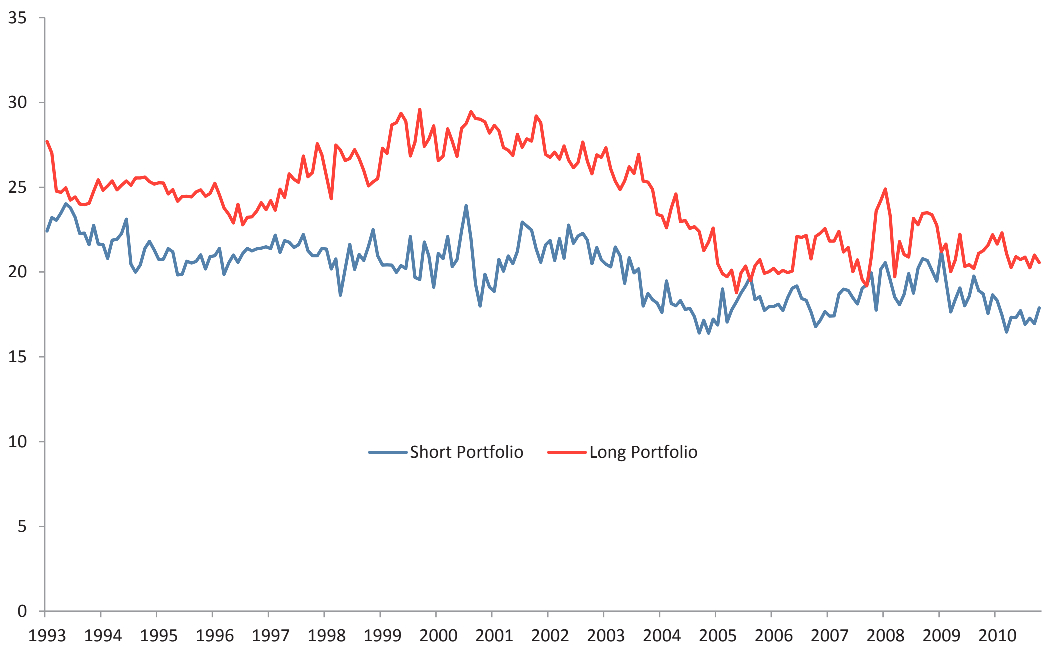

Source: FTSE, NBIM calculations

Chart 23: Funding liquidity risk and cumulated returns on volatility factor, 1994-2011

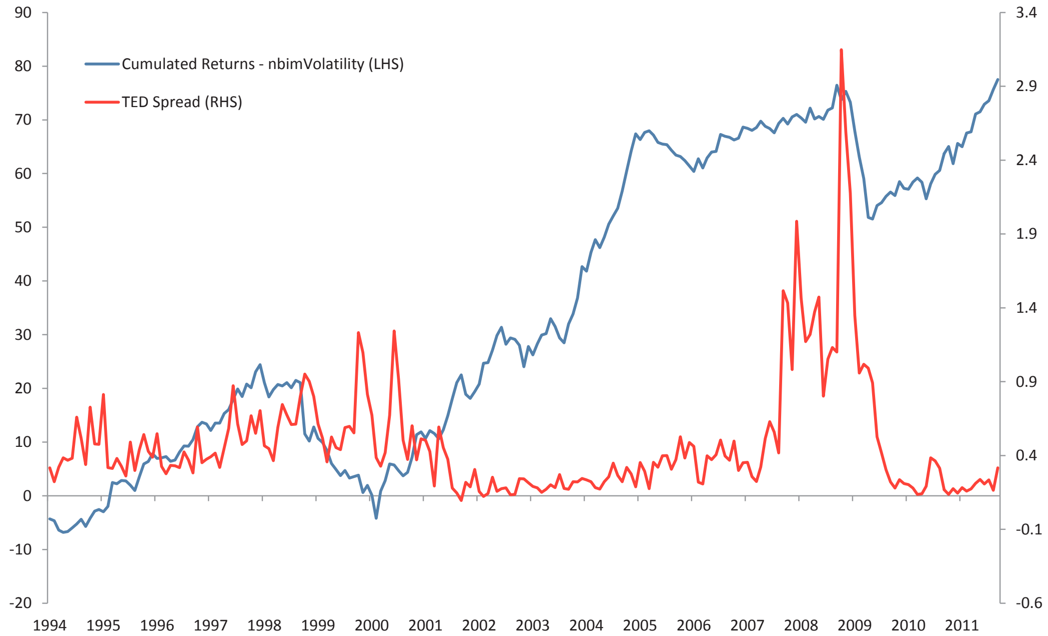

Source: FTSE, NBIM calculations

All in all, there are good reasons to believe that the size, value and volatility premia are all partly compensation for illiquidity risk. Due to the challenges related to specifying a robust and transparent methodology for capturing illiquidity premia directly and in isolation, we would argue that a better way would be to capture illiquidity premia indirectly through the other sources of systematic risk.

Other indirect factor exposures

In addition to the indirect exposures to liquidity risk, the four factors we investigate also have indirect exposures both to each other and to other sources of systematic return. For example, our internally-developed size factor has, at times, significant exposures to value and volatility, as illustrated in Charts 23 and 24. Furthermore, these indirect exposures vary considerably over time. For example, the size factor had a big negative value tilt during the tech bubble in 1999-2000, which has since reversed. The same reasoning can be applied to all the other risk premia that we have addressed.

Chart 24: Value exposure (P/E) in size premium, 1994-2011

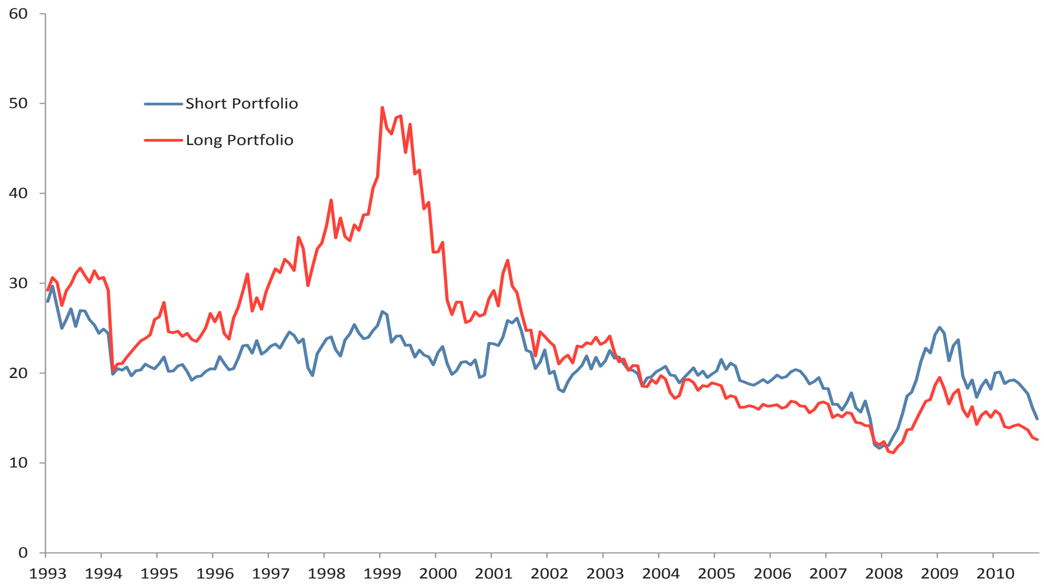

Source: FTSE, NBIM calculations

Chart 25: Volatility exposure in size premium, 1994-2011

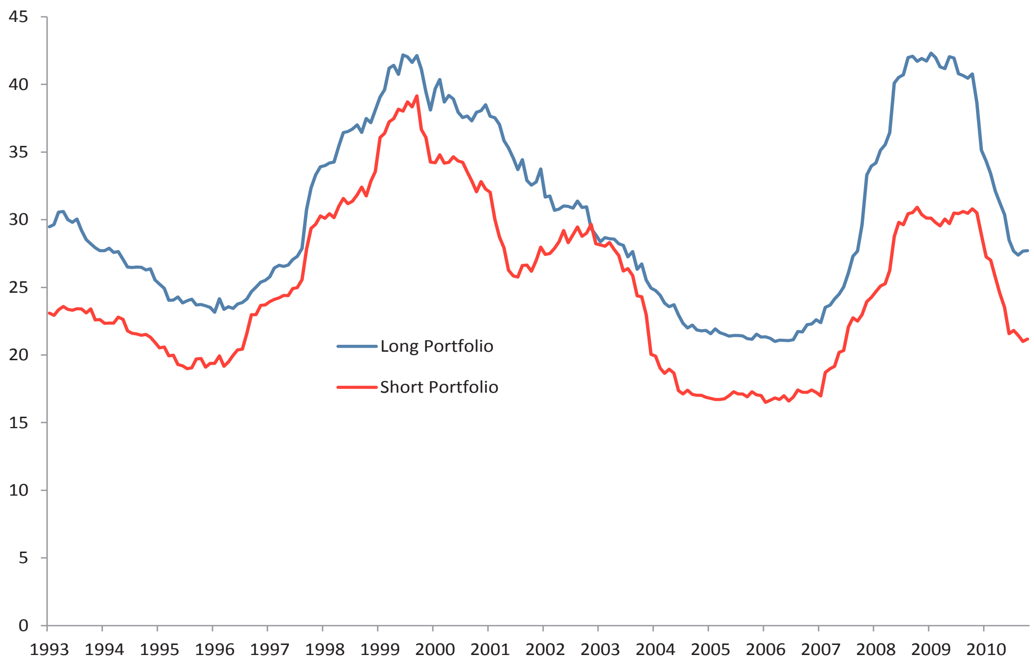

Source: FTSE, NBIM calculations

A diversified approach to systematic risk premia

Above, we documented the low correlations between the Fama-French factors. We compare this with the correlation matrix for our internally-developed risk factors for the period 1994-2011 in Table 18. As explained above, all our internal factors are calculated based on a global universe of stocks, but where we impose region neutrality on all factors. All factors are also sector-neutral, except for the momentum factor, as there are reasons to believe that the momentum factor is related to sector effects. Furthermore, our factors are calculated for the large-cap/mid-cap segment, excluding small caps and emerging markets in order to increase investment capacity in the factors. In spite of the methodological differences, the correlations in Table 18 are broadly consistent with the correlation matrix for the Fama-French factors: the correlations between our internal factors are generally low, and the momentum factor is negatively correlated with both the market and the other risk factors. As a result, we expect the diversification benefits from a combined capture of the various risk premia to be substantial. This is confirmed in Table 18.

Although the correlations between the different factors are low on average over our sample period, the correlations will be time-varying. If the correlations increase towards 1 in times of distress, the diversification benefits will obviously disappear. However, as documented in Table 19, this is not the case. Although each of the individual factors has substantial drawdowns, these drawdowns do not coincide. As shown in Table 19, each of the risk factors has big drawdowns, but the drawdown of the combined factor is substantially less than for any of the factors on a standalone basis. Also, the shortfall of the combined factor is lower than for any of the other factors.

Table 18: Correlations between systematic risk factors, 1994-2011

	FTSE World	nbimValue	nbimSize	nbimMom	nbimVol
FTSE World	1.00				
nbimValue	0.10	1.00			
nbimSize	0.12	0.38	1.00		
nbimMom	-0.30	-0.63	-0.28	1.00	
nbimVol	-0.02	0.00	-0.16	0.24	1.00

Table 19: Risk and return of systematic risk factors, 1994-2011

	nbimValue	nbimSize	nbimMom	nbimVol	nbimCombo
Performance					
Ave Returns, Gross	7.3 %	4.3 %	5.5 %	4.4 %	5.4 %
Ave Returns, Net	6.7 %	3.9 %	4.0 %	3.3 %	4.5 %
Risk					
Volatility	8.1 %	7.2 %	17.4 %	7.3 %	4.4 %
Shortfall 5%	-4.5 %	-4.5 %	-12.3 %	-3.2 %	-2.4 %
Max Drawdown, 6m	-17.6 %	-13.5 %	-47.7 %	-23.8 %	-9.3 %
Sharpe					
Sharpe, Gross	0.90	0.60	0.32	0.60	1.21
Sharpe, Net	0.82	0.54	0.23	0.45	1.01
Beta	0.05	0.05	-0.33	-0.01	-0.06
Resilience					
- Ave Ret when Mkt>0	0.63 %	0.36 %	-0.16 %	0.32 %	0.29 %
- Ave Ret when Mkt<0	0.57 %	0.36 %	1.40 %	0.43 %	0.34 %
- Ave Ret when Mkt<VaR5%	-0.43 %	-0.94 %	4.68 %	-0.82 %	0.62 %

Chart 26: Cumulated returns for systematic risk factors, 1994-2011

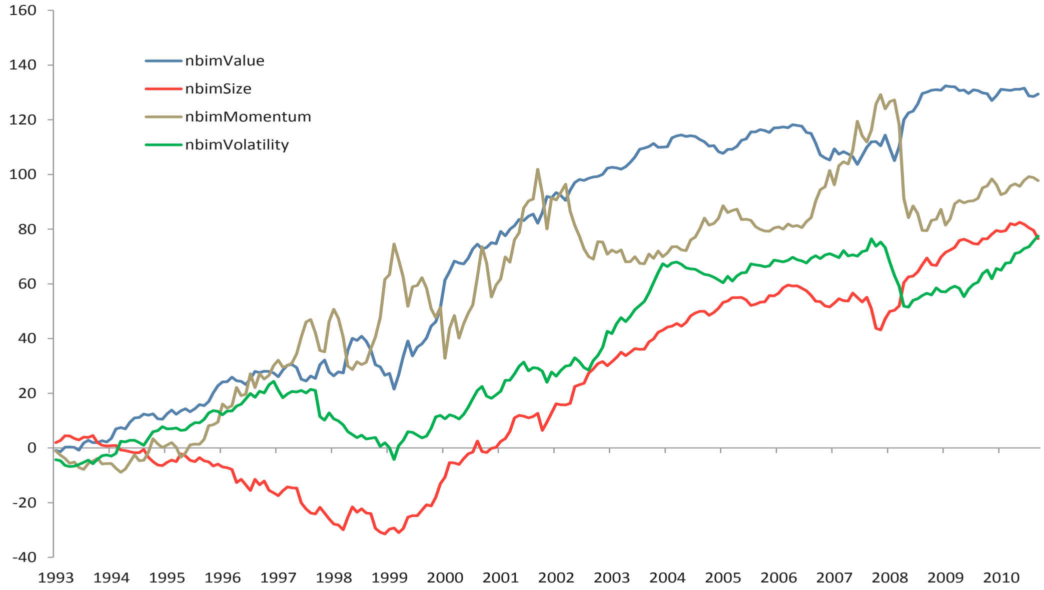

Source: FTSE, NBIM calculations

Chart 27: Cumulated returns for NBIM's combined factor, 1994-2011

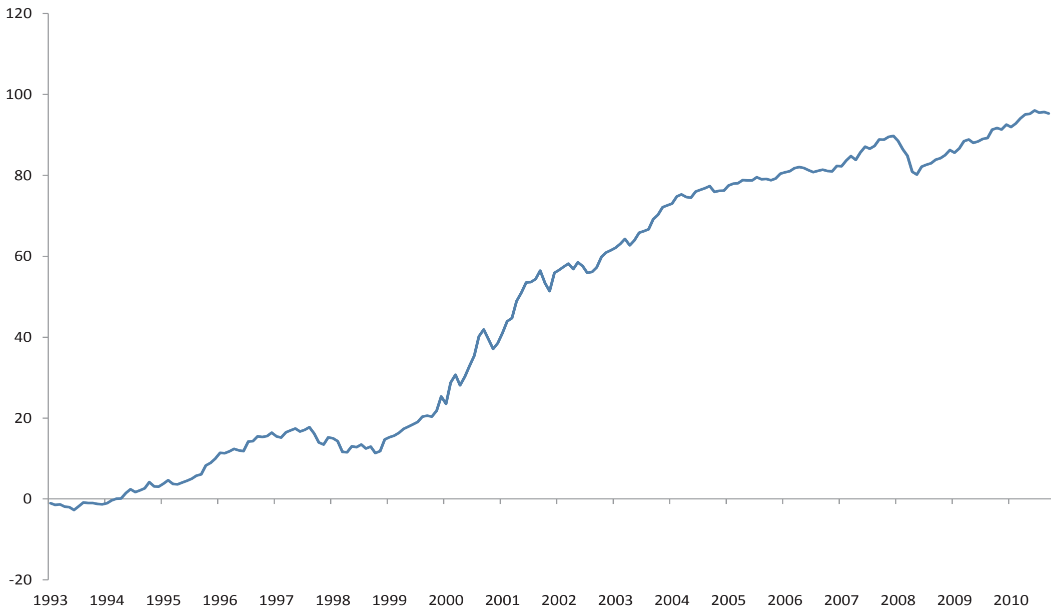

Source: FTSE, NBIM calculations

References

Acharya, V. and L.H. Pedersen (2005): "Asset pricing with liquidity risk", *Journal of Financial Economics*, 77, 375-410.

Adrian, T. and H.S. Shin (2007): "Liquidity and Leverage", working paper, FRB New York and Princeton University.

Amihud, Y. (2002): "Illiquidity and Stock Returns: Cross-section and time-series effects", *Journal of Financial Markets*, 5, 1, 31-56.

Amihud, Y. and H. Mendelson (1986): "Asset pricing and the bid-ask spread", *Journal of Financial Economics*, 17, 2, 223-249.

Amihud, Y. and H. Mendelson (1989): "The Effects of Beta, Bid-Ask Spread, Residual Risk and Size on Stock Returns", *Journal of Finance*, 44, 479-486.

Asness, C., T.J. Moskowitz and L.H. Pedersen (2009): "Value and Momentum Everywhere", working paper, AQR Capital Management.

Brennan, M.J., T. Chordia and A. Subrahmanyam (1998): "Alternative factor specifications, security characteristics, and the cross-section of expected stock returns", *Journal of Financial Economics*, 49, 345-373.

Brunnermeier, M. and L.H. Pedersen (2008): "Market Liquidity and Funding Liquidity", *Review of Financial Studies*, forthcoming.

Carhart, M. (1997): "On Persistence in Mutual Fund Performance", *Journal of Finance*, 52, 57-82.

Chan, L.K., N. Jegadeesh and J. Lakonishok (1996): "Momentum strategies", *Journal of Finance*, 51, 1681-1713.

Chordia, T., A. Subrahmanyam and V. Anshuman (2001): "Trading Activity and Expected Returns", *Journal of Financial Economics*, 59, 3-32.

Cremers, M., A. Petajisto and E. Zitzewitz (2010): ""Should Benchmark Indices Have Alpha? Revisiting Performance Evaluation", working paper.

Fama, E. and K. French (1993): "Common risk factors in the returns on stocks and bonds", *Journal of Financial Economics*, 33, 3-56.

Franzoni, F. (2006): "Where is beta going?", working paper.

Houge, T. and T. Loughran (2006): "Do investors capture the value premium?", *Financial Management*, 35, 2, 5-19.

Huij, J. and M. Verbeek (2007): "On the Use of Multifactor Models to Evaluate Mutual Fund Performance", working paper.

Ibbotson, R.G., Z. Chen and W.Y. Hu (2011): "Liquidity as an investment style", working paper.

Ilmanen, A. (2011): "Expected Returns. An Investor's Guide to Harvesting Market Rewards", Wiley Finance.

Jegadeesh, N. and S. Titman (1993): "Returns to buying winners and selling losers: Implications for stock market efficiency", *Journal of Finance*, 48, 1, 65-911.

Kluger, B.D. and J. Stephan (1997): “Alternative liquidity measures and stock returns”, *Review of Quantitative Finance and Accounting*, 8, 1, 19-36.

Lakonishok, J., A. Shleifer and R.W. Vishny (1994): “Contrarian investment, extrapolation and risk”, *Journal of Finance*, 49, 5, 1541-1578.

Melas, D., R. Suranarayanan and S. Cavaglia (2010): “Efficient Replication of Factor Returns”, *Journal of Portfolio Management*, 36, 2, 39-51.

Moskowitz, T.J. and M. Grinblatt (1999): “Do industries explain momentum?”, *Journal of Finance*, 54, 4, 1249-1290.

Moskowitz, T.J. and L.H. Pedersen (2008): “Measuring Global Market and Funding Liquidity”, working paper, University of Chicago.

Pastor, L. and R.F. Stambaugh (2003): “Liquidity Risk and Expected Stock Returns”, *Journal of Political Economy*, 111, 3, 185-222.

Sadka, R. (2006): “Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk”, *Journal of Financial Economics*, 80, 309-349.

Sharpe, W.F. (1988): “Determining the Fund’s Effective Asset Mix”, *Investment Management Review*, 2, 6, 59-69.

Sharpe, W.F. (1992): “Asset Allocation: Management Style and Performance Measurement”, *Journal of Portfolio Management*, 18, 7-19.

Norges Bank Investment Management (NBIM)
Bankplassen 2
Postboks 1179 Sentrum
N-0107 Oslo

Tel.: +47 24 07 30 00
www.nbim.no