An EDHEC-Risk Institute Publication

Factor Investing in Liability-Driven and Goal-Based Investment Solutions

March 2020

with the support of

Amundi Asset Management | ETF Indexing Smart Beta

EDHEC-RISK Institute

This research has benefited from the support of Amundi in the context of the "ETF and Passive Investment Strategies" Research Chair.

Printed in France, March 2020.

Abstract

A new approach known as factor investing has recently emerged in investment practice, which recommends that allocation decisions be expressed in terms of risk factors, as opposed to standard asset class decompositions. While the relevance of factor investing is now widely accepted amongst sophisticated institutional investors, an ambiguity remains with respect to the exact role that risk factors are expected to play in an asset-liability management investment process. The main objective of this paper is precisely to contribute to the acceptance of factor investing by providing useful pedagogical clarification with respect to the benefits of factor investing within the liability-driven investing paradigm. To this end, we draw an important distinction between the benefits of factor investing in the performance-seeking portfolio and its benefits in the liability-hedging portfolio. We also argue that adopting a factor investing lens offers new useful insights with respect to the improvement of the interaction between performance-seeking and liability-hedging portfolios. Overall, our paper can be regarded as a first step towards the introduction of a comprehensive investment framework blending liability-driven and factor investing, widely recognized as the two most significant advances in institutional money management over the last two decades.

ABOUT THE AUTHORS

Lionel Martellini is Professor of Finance at EDHEC Business School and Director of EDHEC-Risk Institute. He has graduate degrees in economics, statistics, and mathematics, as well as a PhD in finance from the University of California at Berkeley. Lionel is a member of the editorial board of the *Journal of Portfolio Management* and the *Journal of Alternative Investments*. An expert in quantitative asset management and derivatives valuation, his work has been widely published in academic and practitioner journals and he has co-authored textbooks on alternative investment strategies and fixed-income securities.

Vincent Milhau is a research director at EDHEC-Risk Institute. He holds master's degrees in statistics and economics (ENSAE) and financial mathematics (Université Paris VII), as well as a PhD in Finance from the University of Nice-Sophia Antipolis. His research areas include static and dynamic portfolio optimisation, factor investing and goal-based investing. He has co-authored a number of articles published in several international finance journals.

Executive Summary

Executive Summary

A new approach has recently emerged in investment practice known as factor investing, which recommends that allocation decisions be expressed in terms of risk factors, as opposed to standard asset class decompositions. While the relevance of factor investing is now widely accepted amongst sophisticated institutional investors, a number of questions remain with respect to the exact role that risk factors are expected to play in an asset-liability management investment process. The main objective of this paper is to contribute to the widespread acceptance of factor investing by providing some clarification about the benefits of factor investing within the liability-driven investing paradigm.

Modern portfolio theory provides conceptual justification for the use of factor investing in liability-driven investing strategies. It also provides some guidance for the choice of factors and suggests that adopting a factor investing lens can offer useful new insights on the proper design of both the performance-seeking and liability-hedging portfolios.

The concept of policy portfolio has long been a cornerstone of institutional money management, where it refers to a portfolio intended to strike a balance between performance and risk relative to a benchmark, the benchmark being the value of liabilities for investors facing commitments. However, modern portfolio theory, pioneered by the work of Harry Markowitz, William F. Sharpe and Robert C. Merton (all awarded the Nobel Prize in economics), shows that the optimal trade-off between risk and return is in principle obtained by combining a "performance-seeking" building block and a "minimum risk" portfolio, which, in asset-liability management, is the liability-hedging portfolio. In the liability-driven investing framework, the relative allocation to these two building blocks depends on outstanding dollar and risk budgets, and also on its periodic revision in reaction to changes in the opportunity set.

In parallel, the recent emergence of factor investing is also connected with advances in financial economics, and more specifically with academic research on asset pricing, notably including the work of Eugene F. Fama (another Nobel Prize winning economist) and Kenneth R. French. Factor models have long been used for the analysis of a portfolio risk and performance, but starting with an influential article published in 1993 by Fama and French,[1] factors have gained a new status as explanatory variables for stylized facts that were previously regarded as puzzles. Fama and French proposed to interpret the size and value effects in equities – broadly speaking, small stocks tend to outperform large ones and stocks with high book-to-market ratios tend to outperform those with low ratios – in terms of factor exposures: if the market capitalization and the book-to-market ratio proxy for exposures to undiversifiable risk factors, then stocks with a small capitalization or a high ratio are more exposed to these risks, which justifies a premium. Although it does not identify the underlying factors, this explanation fits the class of theoretical asset pricing models previously developed by Robert Merton with the Intertemporal Capital Asset Pricing Model (ICAPM) and by Stephen A. Ross with the Arbitrage Pricing Theory (APT). Beyond such risk-based interpretations of patterns related to size, the book-to-market ratio or

1 - Fama, E. F. and K. R. French. 1993. Common Risk Factors in the Returns on Stocks and Bonds. *Journal of Financial Economics* 33(1): 3–56.

other characteristics, there is another category of economic rationales according to which these patterns reveal some form of market inefficiency or incomplete rationality of market participants.

In investment practice, the interest in factor investing has been driven by several forces. First, there has been increasing recognition that traditional cap-weighted indices have a rather bad risk-return profile, which can be improved by investing in stocks endowed with certain observable characteristics, such as low size or high book value relative to capitalization. Second, investors have become increasingly concerned over active management fees and the broad lack of robustness in generating positive and persistent alpha. As a result, they have started to search for added-value investment vehicles with a performance that can be justified by solid economic arguments and does not require complex and expensive processes to select and allocate to securities. Third, the 2008 financial crisis has led to renewed interest in sound risk management practices, so investors are increasingly inclined to ask what risks they face for the returns they earn. These changes have attracted attention to systematic factor investing, now regarded as an approach that blurs the line between passive investing, which involves replicating a cap-weighted index, and active investing, which involves proprietary selection and/or timing skills.

While the relevance of factor investing is now widely established, the discussion around the choice of these factors is ongoing, and a number of questions remain with respect to the exact role these factors should play. In academia, research has produced many (perhaps too many) candidate "pricing factors", defined as factors that explain differences in expected returns between assets, but not all of them appear to be statistically and economically significant. In investment practice, the notion of factor is more polysemic, and a case can be made that different applications call for different definitions. This paper illustrates the flexibility of the factor investing paradigm by explaining how factors can be used at each stage of the liability-driven investing (LDI) process.

In the performance-seeking portfolio, the objective is to harvest risk premia.

In the performance-seeking portfolio the objective is to harvest risk premia across and within asset classes in the most efficient way possible, which is achieved by diversifying away unrewarded risk exposures. Equities is the asset class in which the concept of factor investing is the most mature, and a few factors, including size, value, momentum, volatility, investment and profitability, have been shown to lead to a robust and economically justified risk premium.

According to fund separation theorems, the performance-seeking portfolio (PSP) should be the one that maximizes the Sharpe ratio, regardless of the existence or nature of the investors' liabilities. Unfortunately, this prescription is difficult to implement because the maximum Sharpe ratio portfolio depends on the expected returns of constituents, which are very hard to estimate, and its out-of-sample performance is severely plagued by estimation errors.

As an alternative to statistical analysis, economic models can offer some help to find the composition of this portfolio. The Capital Asset Pricing Model (CAPM), which was introduced by William F. Sharpe and John V. Lintner in 1964 and 1965, identifies it with the "market portfolio", which consists of all assets weighted by their market capitalization and has been traditionally proxied as a broad cap-weighted index of stocks. But the model hinges on rather unrealistic assumptions, including the fact that all investors have identical expectations, which makes it doubtful that a cap-weighted index is the sought after efficient portfolio, and it is well known that alternative construction methods, as simple as weighting constituents equally, produce higher Sharpe ratios. Thus, in spite of advantages like low turnover and high liquidity, cap-weighted indices are now regarded as unsatisfactory proxies for efficient benchmarks.

One possible improvement over cap-weighted indices would be to address their lack of diversification, which results from their excessive concentration in a few large stocks, by changing the weighting scheme. This approach has been taken in so-called "smart-weighted" equity indices such as equally-weighted indices and minimum variance indices, among others. A second, non-exclusive option is to revise the model to relax some of its controversial assumptions, as in the IACPM and APT, respectively introduced in 1973 and 1976. A common feature of these models is that they predict that expected returns depend not only on how securities co-move with a single factor, namely the market portfolio as in the original CAPM, but also on their co-movements with multiple factors.
This prediction is an appealing property of these models since it meets the empirical finding that expected returns tend to be associated with multiple, non-redundant attributes. In equity markets, four decades of empirical research have led to a long enumeration of more than 300 candidate characteristics, but only a handful of them are statistically robust and economically plausible. Low size and value were the first to be reported, followed by momentum and low volatility, with low investment and high profitability being the latest to join the short list of characteristics with a robust statistical track record and a risk-based or behavioral justification.

By selecting securities with the appropriate characteristics, one can construct portfolios with expected returns above the market without using active management. It is important to note, however, that this profitability usually requires a long investment horizon to be observed, and that in the short run, significant risks of drawdown and underperformance with respect to the cap-weighted index subsist. Examples of such portfolios are given by "investable factor indices", which provide exposure to selected "factors", understood here as long-only strategies with expected returns above the market or as long-short strategies with positive expected returns. This is apparent from Exhibit 1, where the factor indices outperform the broad cap-weighted index by 26 to 187 basis points per year between 1972 and 2016, but display maximum relative drawdowns from 17.91% to 38.39%. The second step in the construction of alternative equity indices would involve improving the diversification of specific risks by deviating from the cap-weighted scheme.

Finally, the investable factors thus constructed can be combined within multi-factor portfolios, the objective of which is to take advantage of multiple factor premia and the relative decorrelation of the factors. One possible objective of such multi-factor allocations is to reduce volatility, as shown in Exhibit 1, where the minimum variance portfolio has an annual volatility of 14.59%, less than that of the least volatile constituent (14.71%). Factor combination can also be used to improve the risk-return characteristics of the broad index while managing the relative risk with respect to this benchmark. Thus, a simple equally-weighted portfolio of the six constituents outperforms the broad index by 112 basis points per year, with a tracking error of 1.83% and a maximum relative drawdown of 12.08%, both of which are less than those of the individual constituents.

In conclusion, factors are not only useful for designing models to estimate risk and expected return parameters, but they now can also be regarded as building blocks for the construction of a well-rewarded performance-seeking portfolio.

Factor exposure matching techniques are traditionally employed for the construction of liability-hedging portfolios in institutional money management. Another relevant practical application in individual money management would be the construction of bond portfolios delivering replacement income for a fixed period of time, typically the first 20 years of retirement. Such "retirement bonds" would be very helpful for individuals investing money for retirement.

Exhibit 1: Examples of portfolios invested in a single or multiple factor(s) with data from September 1972 to December 2016

	Cap-weighted	Mid-Cap	Value	High-Mom	Low-Vol	High Prof.	Low Inv.	Equally-weighted	Min vol	Min TE
Ann. ret. (%)	10.09	11.96	11.37	10.98	10.37	10.35	11.50	11.21	10.99	10.61
Volatility (%)	16.51	17.85	16.43	17.07	14.71	16.75	15.37	15.82	14.59	16.23
Sharpe ratio	0.34	0.39	0.39	0.35	0.37	0.32	0.43	0.40	0.41	0.35
Max. drawdown (%)	53.78	58.28	61.58	50.44	49.09	52.78	52.57	52.14	43.58	52.84
Tracking error (%)	–	5.75	5.25	4.05	4.89	3.40	3.71	1.83	4.48	0.98
Info. ratio	–	0.33	0.24	0.22	0.06	0.08	0.38	0.61	0.20	0.53
Max. rel. drawdown (%)	–	28.57	32.60	18.71	38.39	17.91	25.64	12.08	41.69	3.08
ENC (%)	–	16.67	16.67	16.67	16.67	16.67	16.67	100.00	24.49	61.14
ENCB in vol. (%)	–	16.67	16.67	16.67	16.67	16.67	16.67	99.14	24.49	59.06
ENCB in TE (%)	–	16.67	16.67	16.67	16.67	16.67	16.67	42.41	21.60	61.75

Notes: The broad cap-weighted index and the six long-only equity indices are taken from the Scientific Beta database. The six indices are tilted respectively towards mid-cap stocks, stocks with high book-to-market value, past year winners, low volatility stocks, high profitability stocks and low investment stocks, and they are weighted by capitalization. The last three columns represent portfolios invested in the six factor indices, respectively an equally-weighted portfolio, a minimum volatility and a minimum tracking error with respect to the broad index. The tracking error, the information ratio and the maximum relative drawdown are with respect to the broad index. The last three rows are diversification metrics, and they display the effective number of constituents as a percentage of the nominal number of constituents, which is 6, and the effective number of correlated bets in volatility or in tracking error.

Executive Summary

The second building block of an LDI strategy is a liability-hedging portfolio (LHP), whose role is to replicate the performance of liabilities. Because the nature and horizon of the liabilities are specific to each investor, the LHP has to be customized to meet the needs of each investor. From a factor investing standpoint, the relevant factors are also not necessarily the same as for the PSP. Indeed, the objective here is to track changes in the value of liabilities so the relevant factors here are those that explain the time series of liability returns. Liabilities often consist of a set of payments to be made at predetermined dates, so their present value can be obtained as the sum of the discounted of future cash flows, and the main risk factor is the level of interest rates, the exposure of liabilities to this factor being measured by their duration. If payments were indexed on realized inflation, the discount rates would be real rates and the risk factor would be the real rate level. In these examples, and more generally when a set of risk factors with large explanatory power is available, a standard replication method is to align the exposures of assets with those of liabilities, thus generalizing the duration hedging strategy to account for the presence of multiple factors. One possible such extension involves hedging not only against unexpected changes in the level of interest rates but also in slope and convexity of the yield curve.

These methods could be profitable in individual money management.

While these methods are commonly employed in institutional money management, notably by pension funds, they could be equally profitable in individual money management, where the counterpart of LDI is the goal-based investing (GBI) paradigm and the equivalent of the LHP is known as a goal-hedging portfolio (GHP). For instance, individuals who save money to generate replacement income in retirement would clearly benefit from having access to an asset or portfolio that pays fixed cash flows at regular intervals in the decumulation phase. However, none of the currently available retirement products or financial securities satisfactorily addresses this need. Standard coupon-paying bonds pay cash flows that are not deferred in the future, and balanced funds and target date funds offer no predictability in terms of the replacement income that they will produce. Deferred annuities would be the ideal risk-free asset, but they suffer from a number of shortcomings including their perceived costliness, lack of transparency and reversibility, and the absence of wealth transfer to heirs. It can be argued that if annuities are useful to hedge against the risk of unexpectedly long life, they are not necessary to generate replacement income for a fixed period of time, e.g. for the life expectancy of an individual at retirement. For this purpose, the risk-free asset would be a forward-starting bond with progressive redemption of principal in such a way that the periodic cash flows are constant. Economists Robert Merton and Arun Muralidhar have called for the creation of such "SeLFIES" (for Standard of Living indexed, Forward-starting, Income-only Securities), and a recent paper co-authored by members of academics from the EDHEC-Risk Institute and the

Princeton University describes similar "retirement bonds".[2]

In the absence of these bonds in sovereign debt auction programs, one can replicate them with existing fixed-income securities, which provides further scope for the application of factor exposure matching techniques. Exhibit 2 shows summary statistics for two replication strategies, which respectively match the modified duration or the exposure to the level factor calculated from the four-factor model developed by Nelson, Siegel and Svensson. The constituents are "constant-maturity bonds", which are monthly roll-overs of bonds with a constant maturity chosen at a value of 1 year, 2 years, ... until 30 years. At each rebalancing date, i.e. every quarter in this table, the portfolio is invested in the two constituents with the closest durations or exposures to those of the retirement bond. By testing different accumulation periods, we can see that the GHPs constructed by matching a factor exposure are consistently closer to the bond than a strategy that simply rolls over long-term bonds or short-term money market instruments. Over 11 years of accumulation, the cumulative return of a GHP deviates from that of its benchmark by 0.90% to 5.69%. With long-term bonds, the deviation is from 3.57% to 16.13%, and a cash account, which is often regarded as safe because it never loses money at any horizon, appears highly risky when it comes to securing a certain amount of replacement income: in the best scenario, a cash investor recovers only 68.20% of the initial purchasing power of their savings in terms of replacement income.

The replication exercise can also be performed in decumulation, where the relevant reporting metric is the maximum amount that can be withdrawn every year from the investment portfolio without exhausting savings before the end of the 20-year decumulation period, and without running a final surplus. The results presented in the paper again show that the factor exposure replication strategies lead to withdrawal rates that are much closer to that of the retirement bond target compared to the use of a roll-over of bonds or a money market account.

At the allocation stage in liability-driven investing, the degree of overlap between assets and liabilities can be measured with a suitable multi-class factor model. The model can also be used to construct equity portfolios with better liability-hedging properties than a standard broad cap-weighted index. With these more "liability friendly" portfolios, investors can in principle allocate more to the performance-seeking portfolio for the same risk budget and thus enjoy higher returns.

The last stage of the LDI process involves choosing an allocation to the PSP and LHP. Following fund separation theorems, this allocation depends on risk budgets, typically expressed either in terms of a target tracking error or maximum relative drawdown relative to the liabilities. The risk of the LDI strategy depends on the investment policy and the risk of each building block, as well as their correlation. For instance, its tracking error with respect to the liabilities for a given allocation increases with the tracking error of the LHP, but also with that of the PSP, which is not controlled at the portfolio construction stage since fund

2 - Martellini, L., V. Milhau and J. Mulvey. 2019. 'Flexicure' Retirement Solutions: A Part of the Answer to the Pension Crisis? *Journal of Portfolio Management*: 45(5) 136-151.

Exhibit 2: Simulation of level-matching and duration-matching portfolios in accumulation.

Retirement year	Beginning of accumulation	Strategy	Ann. return (%)	Gross relative return (%)	Ann. volatility (%)	Tracking error (%)	Max. relative drawdown (%)
2000	Jan. 1989	Ret. bond	10.67	-	10.73	-	-
		GHP Lev.	10.76	100.90 1	1.50	2.06	3.61
		GHP Dur.	10.78	101.15	11.51	2.08	3.68
		15-year bond	10.98	103.19	12.95	3.90	11.37
		Cash	5.28	57.74	0.21	10.73	52.94
2005	Jan. 1994	Ret. bond	8.69	-	11.51	-	-
		GHP Lev.	8.99	103.15	12.32	1.72	2.39
		GHP Dur.	9.00	103.26	12.40	1.77	2.49
		15-year bond	9.56	109.17	13.58	3.76	8.87
		Cash	3.93	61.12	0.18	11.51	53.25
2010	Jan. 1999	Ret. bond	6.55	-	11.75	-	-
		GHP Lev.	6.97	104.44	12.49	1.45	2.99
		GHP Dur.	6.98	104.55	12.59	1.51	3.10
		15-year bond	6.89	103.57	14.22	4.20	11.56
		Cash	2.91	68.20	0.17	11.75	51.31
2015	Feb. 2004	Ret. bond	7.80	-	12.19	-	-
		GHP Lev.	8.03	102.29	12.84	1.78	6.54
		GHP Dur.	8.12	103.27	13.00	1.84	6.55
		15-year bond	9.06	113.47	14.20	3.53	8.49
		Cash	1.42	51.37	0.14	12.19	53.24
2020	Feb. 2009	Ret. bond	5.27	-	13.36	-	-
		GHP Lev.	5.74	104.71	13.78	1.08	1.39
		GHP Dur.	5.83 1	05.69	13.90	1.14	1.46
		15-year bond	6.80 1	16.13	13.13	2.53	5.57
		Cash	0.46	61.69	0.05	13.36	52.05
2025	Jan. 2014	Ret. bond	8.56	-	11.35	-	-
		GHP Lev.	8.81	101.25	11.57	0.68	1.84
		GHP Dur.	8.87	101.55	11.65	0.71	1.86
		15-year bond	7.69	95.72	10.02	1.79	6.94
		Cash	0.79	66.87	0.07	11.35	36.78

Retirement takes place on the first day of the year indicated in the first column. The returns of retirement bonds and the various strategies are simulated from the beginning of accumulation, which is the first day of the month in the second column, until the retirement date or 1 June 2019, whichever comes first. Simulations are based on the US zero-coupon rates published on the website of the Federal Reserve. The beginning of accumulation is chosen so as to ensure that the maturity of the last replacement income cash flow does not exceed 30 years, since this is the longest maturity of US Treasury bonds.

separation principles recommend that the PSP be designed with no hedging concern in mind. This property has interesting practical consequences because it implies that by decreasing the tracking error of the PSP with respect to the liabilities, an investor can allocate more to this portfolio while

Exhibit 3: Liability-driven investing strategies with matched relative maximum drawdown; data from September1971 to December 2016.

Performance-seeking portfolio	Equ.-broad	Min distance		Min systematic tracking error		Max ENB	
Target TE w.r.t. Equ.-broad (%)	–	None	2	None	2	None	2
Cap-weighted indices PSP							
Annual return (%)	10.15	11.84	11.17	10.47	10.86	10.47	10.69
Tracking error (%)	19.16	20.59	18.76	17.32	18.19	17.32	18.92
Max relative drawdown (%)	77.38	67.90	68.34	68.24	70.51	68.24	73.63
LDI Strategy							
Allocation to PSP (%)	40.00	50.87	51.32	50.40	47.72	50.40	44.09
Annual return (%)	9.72	10.92	10.43	9.94	10.13	9.94	10.02
Volatility (%)	9.46	10.43	9.86	9.50	9.53	9.50	9.53
Cumulative relative return (%)	64.80	169.99	121.27	80.97	95.58	80.97	86.58
Gain in funding ratio w.r.t. reference (%)	–	63.83	34.27	9.81	18.68	9.81	13.22
Tracking error (%)	7.51	10.22	9.47	8.62	8.55	8.62	8.20
Information ratio	0.16	0.24	0.20	0.17	0.19	0.17	0.18
Max relative drawdown (%)	40.79	40.79	40.79	40.79	40.79	40.79	40.79
Minimum variance indices PSP							
Annual return (%)	10.15	12.97	12.90	12.65	12.90	12.65	12.90
Tracking error (%)	19.16	18.02	17.94	16.12	17.94	16.12	17.94
Max relative drawdown	77.38	63.44	57.57	55.94	57.57	55.94	57.57
LDI Strategy							
Allocation to PSP (%)	40.00	55.96	64.47	66.75	64.47	66.75	64.47
Annual return (%)	9.72	11.55	11.81	11.61	11.81	11.61	11.81
Volatility (%)	9.46	9.72	10.54	9.83	10.54	9.83	10.54
Cumulative relative return (%)	64.80	249.75	288.60	258.45	288.60	258.45	288.60
Gain in funding ratio w.r.t. reference (%)	–	112.23	135.80	117.51	135.80	117.51	135.80
Tracking error (%)	7.51	9.91	11.43	10.67	11.43	10.67	11.43
Information ratio	0.16	0.31	0.29	0.29	0.29	0.29	0.29
Max relative drawdown (%)	40.79	40.79	40.79	40.79	40.79	40.79	40.79

Minimum distance, minimum systematic tracking error and maximum ENB portfolios are invested in the six equity factor indices, and some of them are subject to a 2% tracking error constraint per year with respect to the broad index. The allocation to the PSP is calculated so as to match the maximum relative drawdown of the reference strategy over the sample period, the reference strategy being the one invested in the broad index and the LHP. Liabilities are represented by a 10-year constant-maturity bond. The gain in funding ratio with respect to the reference strategy is calculated as $[1 + r_2]/[1 + r_1] – 1$, where r_1 and r2 are the respective cumulative returns of the reference strategy and the alternative one.

staying within the limits of a given risk budget. This leads to an increase in upside potential, unless the performance of the more "liability-friendly" PSP that replaces the original PSP is too inferior.

To construct an equity performance-seeking portfolio with better liability-hedging properties than a broad capweighted index – which is often the default option –, it is useful to start by measuring the overlap between the PSP and liabilities. This can be done by measuring their respective exposures to a set of risk factors, provided the risk factors have been shown to explain a large fraction of the common time variation of assets and liabilities. We propose to use an eight-factor asset-liability management model with the equity market factor, the long-short size, value and momentum equity factors from Ken French's library, the "betting-against-beta" equity factor, the level of interest rates, the term spread and the credit spread. (The betting-against-beta factor is the excess return of a portfolio of low beta stocks over a portfolio of high beta stocks.) The model captures between 90% and 100% of the variance of equity factor indices, and 96.6% of that of liabilities. Liabilities are mostly exposed to term structure factors, as expected, but equity indices are also exposed to these factors, sometimes significantly in the statistical sense, and not all of them have the same exposure, which suggests that they have heterogeneous hedging abilities. For instance, we find that the low volatility index has the most negative exposure to the level factor, thus making it the most "bond-like".

With the factor model at hand, a variety of new PSPs can be constructed using different weighting methods, including minimizing the distance between the exposures of the equity portfolio and those of liabilities, minimizing the systematic tracking error with respect to the liability portfolio (defined as the part of the tracking error that arises from factor exposures), or by maximizing the "effective number of bets" (ENB), i.e. by maximizing diversification across the eight factors. Variants of these weighting schemes are additionally obtained by constraining the tracking error with respect to the broad index to be less than a cap, say 2% per year. As is clear from Exhibit 3, the relative risk of an equity portfolio is reduced by replacing the broad index with a PSP constructed from equity factor indices, especially when the constituents are minimum variance as opposed to cap-weighted portfolios. For instance, the maximum relative drawdowns of alternative PSPs range from 55.94% to 73.63%, versus 77.38% for the broad index, and with the provision of the unconstrained minimum distance portfolio of cap-weighted indices, these PSPs also have lower tracking errors by 24 to 204 basis points per year. In addition, they outperform the broad index in this sample period because the expected long-term outperformance of long-only factors over the equity market materialized in this sample period.

Thus, starting from a reference strategy invested in the broad index and the perfect LHP with respective weights of 40% and 60%, one can allocate more than 40% to each alternative PSP while keeping the maximum relative drawdown of the strategy unchanged. As a result, each strategy using an alternative PSP outperforms the reference strategy by an amount that depends both on the increase in allocation and in the gain in average return within the PSP. For instance, the strategy in which the PSP is the portfolio that minimizes the systematic tracking error earns 9.94% per year, versus 9.72% for the one that uses the broad index. The annual gain may seem to be modest, but after 45 years, it translates into a gain of 9.81% in funding ratio. With minimum variance versions of these factor portfolios, the gain in funding ratio over the sample period rises to 117.51%, thanks to the higher annual return of the PSP.

In conclusion, while factor investing and liability-driven investing relate to two separate strands of the academic literature, there is a case for combining these approaches. Each of the three steps of a liability-driven investing process, namely the construction of a well-rewarded performance-seeking portfolio, the construction of a safe liability-hedging portfolio and an efficient allocation to these building blocks, can be better addressed by taking a factor perspective. Our paper can be regarded as a first step towards the introduction of a comprehensive investment framework blending liability-driven investing and factor investing.

Executive Summary

1. Introduction

1. Introduction

Two major new investment paradigms, respectively known as "liability-driven investing" and "factor investing", have had a profound impact on the strategic asset allocation decisions made by institutional investors in general, and on pension fund investment practices in particular.

Historically, asset allocation practices were firmly grounded on one overarching foundational concept, the policy portfolio, which had in particular dominated pension fund investment practices for decades. This policy portfolio, a theoretical reference portfolio allocated to different asset classes according to a mix deemed most appropriate for the pension fund, was constructed either from a simple asset-only perspective, or from a more sophisticated surplus optimization perspective. In the latter case, the goal was to efficiently allocate to various risky asset classes so as to achieve the best possible compromise between the risk relative to the pension fund liabilities and the excess return that the investor could hope to obtain through the exposure to rewarded risk factors. Over the past 15 years or so, this old paradigm has gradually been recognized as obsolete, and its death was announced, or rather predicted, by Peter Bernstein in the March 2003 edition of his "Economics and Portfolio Strategy" newsletter. Two main reasons explain this change.

First, academic research has shown that the main concern of a typical pension fund, which is to generate performance without excessively deviating from the evolution of its liabilities, is best addressed by investing in two reference portfolios with well-defined objectives. A performance-seeking portfolio (PSP) aims to collect the premia on risky assets in the most efficient way, and a liability-hedging portfolio (LHP) is in charge of replicating liabilities. This functional separation is in line with the fund separation theorem developed by Sharpe (1964), which implies that any mean-variance efficient portfolio is an overlay of the maximum Sharpe ratio portfolio and the risk-free asset. In the presence of liabilities, the risk-free asset changes in nature and becomes the LHP, or the portfolio that maximizes the absolute correlation with liabilities if perfect replication is not attainable, and the theoretically optimal PSP is the same as in an asset-only context: hence, a new fund separation theorem is obtained which advocates separation of assets between a PSP and an LHP (see Martellini and Milhau (2012)).

> Cost-efficient insurance against shortfall risk is better implemented through a dynamic management of risk budgets.

The second reason for the decreasing popularity of the policy portfolio as a cornerstone of asset-liability management (ALM) is the recognition that it does not provide a strategic allocation solution consistent with the state of the art of research in asset allocation. Indeed, the presence of a substantial amount of predictability in time-varying risk and return parameters for financial assets invalidates the relevance of any portfolio that would be held constant by investors for a sustained period of time, with no revisions of weights as a function of changes in market conditions. Moreover, cost-efficient insurance

1. Introduction

against shortfall risk, whether in absolute terms or relative to liabilities, cannot be implemented with static strategies, for such strategies must be calibrated against the worst case and thus incur prohibitive opportunity costs in the more likely scenarios where extremely adverse market conditions do not materialize (see Deguest, Martellini and Milhau (2014)). A better approach uses dynamic management of risk budgets, as in constant proportion portfolio insurance, and this reduces opportunity costs by allowing for higher exposure to risky assets when minimum funding levels are not at risk. The bottom line is that one should think in terms of dynamic allocation strategies as opposed to searching for a static portfolio. Taken together, these driving forces have led to the progressive emergence of a modern approach to ALM for pension funds, known as "dynamic liability-driven investing". The main principle is that assets should be split into a risky but efficient PSP, and a safe LHP, and that the split should evolve as a function of changes in market conditions and also as a function of prudential risk budgets defined by the regulators, who have increased their focus on the respect of minimum funding ratio levels.

In parallel, a new approach known as factor investing has emerged in investment practice, which recommends that allocation decisions be expressed in terms of "factors", as opposed to standard asset class decompositions. But it turns out that there is no single definition of what constitutes a "factor", and the multiple definitions that have been proposed focus either on performance or on risk. In the academic literature, the term "factor" was first used to refer to a "common risk factor" that affects the returns on all securities in a given universe, e.g. the market factor for stocks, the common shocks impacting all firms in an industry, or the level of interest rates for bonds. Sharpe's (1963) one-factor model for portfolio optimization, which was introduced before the Capital Asset Pricing Model, uses the word in that sense, and the notion of risk factor is central in the literature on the Arbitrage Pricing Theory (APT), which began with Ross (1976). The APT pioneered the class of multi-factor models, which have become the norm for risk decomposition, in particular with Barra risk models.

While the APT established a clear relationship between the expected return on a security and its factor loadings, the connection between security characteristics known to have an impact on expected returns and the notion of factor was not made immediately thereafter. In the academic world, Merton (1973) developed the Intertemporal Capital Asset Pricing Model, which predicts that expected returns depend on a security's covariances with changes in investment opportunities, e.g. with interest rates, in addition to its covariance with the market, and empirical research identified a number of patterns in stock returns that could not be explained with Sharpe's CAPM and were thus referred to as "anomalies" or "puzzles". Examples of such characteristics include the market capitalization of a firm (Banz 1981), its book-to-market ratio (Stattman 1980) and its price-earnings ratio (Basu 1983). In investment practice, some of these attributes were used in the context of "styles" or "themes" as part of a long-standing tradition: for instance, the idea of buying stocks with low price relative to fundamentals, which describes strategies focusing

1. Introduction

on book-to-market or dividend-price, among many other possible ratios, dates back to at least Graham and Dodd (1934).

It seems that Fama and French (1993) were the first to re-interpret the size and book-to-market effects in average returns in terms of "factors", by arguing that they are not the result of market inefficiencies, but that they can be perfectly explained by rational asset pricing models if one accepts the idea that size and book-to-market proxy for exposures to undiversifiable risk factors like those of Ross' APT, or to state variables that investors want to hedge against, as in Merton's ICAPM. The risk-based interpretation is not universally accepted, and some authors (see e.g. Daniel and Titman (1997)), without questioning the existence of the effects in the first place, argue that they reveal some form of irrationality or inability to process information on the part of investors or the existence of market frictions. Thus, risk-based explanations and behavioral or friction-based explanations co-exist for many risk premia identified in stock returns. A number of papers subsequently studied the cross-sectional determinants of expected returns and documented effects that are not captured by Fama and French's three-factor model (see the momentum effect in Jegadeesh and Titman (1993) and the volatility effect in Ang et al. (2006)), leading to more than 300 "factors" in stocks, according to the literature survey conducted by Harvey, Liu and Zhu (2016). In view of this huge number, a key question for researchers is how many of these candidate factors are not just data mining products, are robust out of sample, and are thus really useful for estimating expected returns.

On the practical side, factor investing has gradually emerged as a paradigm that blurs the traditional distinction between active and passive investing. Passive investing has long been taken as synonymous of investing in broad capweighted indices, while active management allowed investors to incorporate views and convictions through security picking and/or proprietary models. The situation has changed with the recognition that cap-weighted indices do not have an efficient risk-return profile, owing to their high concentration in a few very large stocks, and that the simple approach of weighting constituents equally improves both volatility and average returns (see Table 1 in this paper for an illustration). There has also been increasing concern over fees and transparency in active investing: in the absence of detailed information on the investment process, it is difficult to figure out how successful it is likely to be in the future, and thus if fees are justified. If positive long-term performance and/or risk reduction can be achieved by applying systematic rules for security selection and weighting, investors may not be willing to pay active fees. The aim then becomes to identify the sources of performance that have not only proved robust in the past but are also likely to persist in the future, and to find selection and weighting procedures that capture this performance at the best possible costs. While factor models have a long-standing tradition in risk and performance analysis, investors have only more recently started to consider investing directly in factors, so as to have direct control over the sources of risk and return in their portfolios. This move was favored by the financial crisis that began in 2008, which fostered the interest in risk management practices. In parallel, "smart-weighted" products have emerged and grown in the equity class, whose aim is to improve the risk-return profile of cap-weighted indices by better diversifying individual risks across securities. The smart-weighting approach and the factor investing perspective can be combined to create "smart beta" products, which target both exposures to selected factors and efficient diversification of individual risks.

Based on a survey of the academic literature, Martellini and Milhau (2018a) identify four categories of factors. The first is that of asset pricing theory, as exposed by Cochrane (2005): a pricing factor helps explain the cross section of expected returns in the sense that securities with different exposures to this factor but identical exposures to other pricing factors have different expected returns. The second category is broader and corresponds to a definition of factors that is currently popular in investment practice: a factor is a strategy with a positive long-term expected return, but possibly a risk of loss in the short run, whether this premium is justified by risk exposures as in an asset pricing model, by market frictions or by investors' incorrect assessment of prices, provided these mistakes cannot be completely arbitraged away. The third definition of a factor, which is also commonly adopted in practice, is as a common risk factor: it is in this sense that factors were first addressed in the academic literature. The last definition is inspired by Merton's ICAPM: a factor can be a state variable that describes investment opportunities at a given point in time. Among such variables are those that characterize risk premia, i.e. those that have some forecasting power for future returns. Standard examples include the dividend yield for stocks and forward rates for bonds. The appendix A of this paper presents the four definitions in greater detail. While they refer to distinct sets of financial or macroeconomic variables, these definitions are not necessarily inconsistent, as some variables may fall into multiple categories. Indeed, Merton's ICAPM predicts that state variables are pricing factors, while Ross' APT emphasizes the pricing role of risk factors. The identification of pricing factors can be exploited to construct zero-dollar strategies with positive expected returns if one can take long positions in the securities with the highest expected returns, and short positions in the others.

The main objective of this paper is to explain how factor investing can be used at each stage of a liability-driven investing process, i.e. in the construction of the performance-seeking portfolio, that of the liability-hedging portfolio and the allocation to these building blocks. Interestingly, the factor perspective is flexible enough to be used at the portfolio construction stage with very different purposes. Consider for instance the PSP construction, where the objective is to efficiently harvest the risk premia of risky assets. Theoretically, this should be done by maximizing the Sharpe ratio – the expected return in excess over the risk-free rate, divided by the volatility – and, as shown by Martellini and Milhau (2015), the resulting portfolio contains zero unrewarded risk with respect to pricing factors, so the portfolio risk is entirely explained by the factor exposures. In practice, this approach is difficult to implement because it heavily relies on expected return estimates, which are hard to obtain, but the objective of efficient collection of risk premia can be attained with alternative diversification methods involving less estimation risk (see Martellini, Milhau and Tarelli (2014a)). Now consider the construction of the LHP. The objective here is very different since it is to track the returns on the value of liabilities as accurately as possible. The first-order question is not to identify rewarded factors, but to have a list of risk factors that have good explanatory power for the returns of liabilities and those of the hedging securities, and to equate the exposures of the asset and liability sides of the balance sheet to these factors. The duration hedging technique, in which one equates the exposures to a change in the level of interest rates, is a traditional example. The notion of state variable is also useful in the portfolio construction process, since one may want to incorporate signals so as to overweight or underweight certain securities when market conditions change. This approach is theoretically motivated by Merton's model for multi-period portfolio selection, in which the composition of each building block of the optimal portfolio strategy depends on the time-varying risk premia, volatilities and correlations. In practice, time-varying volatilities can be estimated with dedicated models such as ARCH and GARCH. Estimating time-varying expected returns can be much more challenging unless one can rely on good signals. A recent paper by Maeso, Martellini and Rebonato (2019d) provides an example with conditional carry strategies in the bond market: depending on the expected term premium estimated by the method of Cieslak and Povala (2015), the investor decides to target a longer or shorter duration than a benchmark.

Finally, the factor perspective can also be usefully adopted at the last stage of the liability-driven investing process, which is the allocation to the building blocks. In the last section of this paper, we introduce a multi-factor model for the analysis of strategies invested in both a PSP and an LHP, and we apply it to the decomposition of the risk (absolute or relative to liabilities) of these strategies. We then revisit the problem studied by Coqueret, Martellini and Milhau (2017), which is the construction of equity portfolios with better liability-hedging properties. The motivation for this part of the analysis is that by decreasing the tracking error of the PSP with respect to liabilities, a fund manager can allocate a bigger fraction of assets to this building block without increasing the tracking error of the strategy. If the PSP with reduced relative risk has at least as high an expected return as the original PSP, typically benchmarked against a broad cap-weighted equity index, this larger allocation results in higher average returns at the level of the multi-block strategy. A performance gain can still be achieved even if the alternative PSP underperforms the original one, but this underperformance is not too large compared to the increase in allocation. Coqueret, Martellini and Milhau (2017) show that selecting stocks with high dividend yield and/or low volatility is an effective way of aligning the PSP with liabilities exposed to nominal interest rate risk, and that further risk reduction is achieved by deviating from the cap-weighting scheme, e.g. with equal weighting or variance minimization. In this paper, we study the same

Improving the liability-hedging properties of the performance-seeking portfolio allows investors to enjoy larger potential returns.

1. Introduction

question from the factor perspective. We find that minimizing the distance between the factor exposures of assets and liabilities gives mixed results in terms of replication, but minimizing the systematic tracking error, which is the part of the error explained by the factors, leads to a reduction in the ex-post tracking error of equity portfolios.

These results have potentially important implications for all institutional investors facing liability constraints, and for defined-benefit (DB) pension plans in particular. It is important, however, to recognize that the transition from an asset-based to a factor-based liability-driven investing framework raises profound governance issues. For example, if the performance of managers of the equity segment within the PSP is purely measured in terms of their ability to outperform a given commercial benchmark, they will have little if any incentive to improve the alignment of the factor exposure in their portfolio with respect to the factor exposures within the LHP portfolio, if such an alignment may hurt the performance from an asset-only perspective, even if it is likely to add value from an ALM perspective. In addition to having strong potential relevance for defined-benefit pension plans, the results presented in this paper may also be of interest to defined-contribution (DC) pension plans. Indeed, even if, indeed precisely because, the shift from a DB to a DC structure implies that the risk has been transferred from the sponsors of the pension plan to the beneficiaries of the plan, it is of critical importance for DC pension plan managers to adopt a liability-driven perspective in the portfolio management process so as to account for the presence of relevant goals for beneficiaries even if they are not to be treated as formal liabilities. In this context again, the factor investing paradigm appears ideally suited to optimize the adequacy of outcomes with respect to the needs of individuals expressed as target levels of replacement income.

The rest of the paper is organized as follows. In Section 2, we examine the benefits of factor investing for efficient harvesting of risk premia, with particular emphasis on equity markets, where this approach is most mature. In Section 3, we analyze how the proper matching of factor exposures between the liability and asset sides allows for accurate replication of investors' liabilities, with an application to retirement investing decisions by defined-benefit or defined-contribution pension plans. In Section 4, we study equity portfolios constructed to have better hedging properties than a broad cap-weighted index. We present our conclusions in Section 5.

1. Introduction

2. Efficient Factor Risk Premia Harvesting: Improving the Performance Benefits of the Performance-Seeking Portfolio with Factor Investing

2. Efficient Factor Risk Premia Harvesting: Improving the Performance Benefits of the Performance-Seeking Portfolio with Factor Investing

Any liability-driven investing strategy must include a performance-seeking portfolio (PSP), which should in theory target the highest Sharpe ratio, i.e. the highest expected return per unit of risk taken. But when it comes to implementation, the MSR portfolio suffers from a major shortcoming, namely its strong dependence on expected returns, which are very hard to estimate accurately using purely statistical methods, as Merton (1980) has shown. A more promising road is to estimate expected returns by using more than past return series and by using the information that we usually have about securities. This approach has been notably adopted in the equity class, where a large body of empirical research has identified a number of characteristics that tend to be associated with higher average returns. It is for this reason that we focus on this asset class in this section, and also because it tends to form the bulk of investors' performance-seeking portfolios. We first review the statistical and economic evidence for six characteristics that are known to have an impact on the expected returns of stocks, namely size, value, momentum, volatility, investment and profitability, and we then describe the construction principles for investable forms of these factors, known as factor indices. Finally, we show that the benefits of a factor approach to the construction of equity portfolios over investing in standard cap-weighted indices can be further improved by combining several factors through allocation techniques.

2.1 Factors in the Equity Class

Equities are the asset class in which the concept of factor investing is the most mature, both in research and in investment practice. It is in this class that "anomalies" in returns, understood as patterns that cannot be explained with the CAPM, were first uncovered through the size, value and momentum effects. This class has also been the primary test universe for the multi-factor models developed since Fama and French (1993). These research advances can perhaps be explained by the availability of good long-term data, at least on the US market, and the fact that the equity market as a whole is less affected by liquidity issues than other alternative classes. On the investment side, equities form the bulk of institutional investors' performance-seeking portfolios, and in the benchmarking industry, interest in factor indices has grown steadily since the 2008 financial crisis and the recognition that a significant proportion of the returns of active strategies could in fact be captured by taking a passive approach.

2.1.1 The Market Factor

CAPM: Market as the Single Pricing Factor

In the CAPM, the "market factor" is the return on the market portfolio, which is the portfolio containing all capital assets weighted by their market capitalization. It is not equity-only, but it has been traditionally identified with the returns on a broad equity portfolio since the first empirical tests of the CAPM. Weighting stocks by capitalization in the index seems to be the option that is most consistent with the definition of the theoretical market portfolio, and has thus been a common approach in empirical tests of the CAPM at least since Jensen (1969). However, equally-weighted indices have also often been used, including in highly influential studies like those by Black, Jensen and Scholes (1972) and Fama and MacBeth (1973), who consider the stocks listed on the New York Stock Exchange, and by DeBondt and Thaler (1985), who include all securities listed on the Center for Research on Security Prices database.

In the CAPM, differences in exposures to the market factor explain the differences in expected returns across assets, as stated in Equation (A.2) in the Appendix, so that a testable implication of the model is the nullity of the alpha in the following regression:

$$\tilde{r}_i = \alpha_i + \beta_i \tilde{r}_{mkt} + \varepsilon_i, \tag{2.1}$$

where $\tilde{r}_i$ is the excess return on security i and $\tilde{r}_{mkt}$ is the excess market return.

The empirical performance of this basic version of the CAPM is poor. Test assets in the left-hand side of the regression (2.1) are in general taken to be portfolios, rather than individual stocks, in order to address the issue of cross-sectional dependence between idiosyncratic returns (Black, Jensen and Scholes 1972), and to diversify estimation errors in individual betas (Fama and MacBeth 1973). One result obtained in the early studies is that, in contrast with the model prediction, alphas are negative for high-beta and positive for low-beta portfolios: Table 2 in Black, Jensen and Scholes (1972) displays a monotonic pattern of alphas with respect to betas. This leads, in most samples, to an estimated factor premium that is lower than the historical average return on the market portfolio, and an estimated risk-free rate that is greater than the average rate on Treasury bills (see Table 4 in Fama and MacBeth (1973)).

Beyond CAPM

At this point, several approaches are possible to improve our understanding of the relationship between expected return and risk. The first is to improve the measurement of market returns, and also possibly of the risk-free rate, following the observations of Roll (1969), and the criticism by Roll (1977), who argues that the validity of the CAPM cannot be tested until the market portfolio is known. To this end, Mayers (1972) proposes to introduce human capital, which is an important non-marketable asset for most agents, and this idea is further explored by Fama and Schwert (1977), who also complete the equity portfolio with US Government bonds[3], but conclude that changes in aggregate income are not sufficiently correlated with returns on marketable assets to improve the explanatory power of the market portfolio.

Another route is to deviate from the assumptions of the CAPM. Black (1972) relaxes the assumptions that investors can lend and borrow at the risk-free rate, and finds that a modified version of the central CAPM prediction holds: expected excess returns are still linearly related to betas, but they must be taken with respect to the expected return on a "zero-beta" portfolio, not the risk-free rate. Thus, Black's version of the CAPM involves

3 - In Fama and Schwert's study, the market value of bonds is proxied as the face value of these bonds (p. 98 of their paper).

estimating another parameter in addition to the market risk premium, namely the expected return on the zero-beta portfolio. Gibbons (1982) and Shanken (1985) introduce dedicated econometric procedures to perform this estimation. It should be noted that in this model, the alpha of a security, defined as its expected return minus its beta times the expected market return, is no longer zero. Gibbons (1982) still finds that the restrictions imposed by the model are rejected by the data.[4]

Jagannathan and Wang (1996) relax the assumption that investors live for a single period, and therefore consider a version of the model in which Equation (A.2) holds conditionally at the beginning of each period. In the unconditional version of the model, which serves as the basis for empirical work, expected returns not only depend on the market beta, but also on the beta with respect to the conditional market premium. In the conditional CAPM, the (unconditional) expected return of an asset increases in the covariance of the conditional beta with the conditional market premium, so that if an asset has higher market beta when the market premium is high, it is expected to earn a higher average return than an asset whose beta tends to be lower when the premium is high. The intuition is that if the asset is less exposed to the market when the expected market return is high – a characteristic of "bad" economic times –, it will be attractive to mitigate the adverse impact of negative market shocks on wealth, and therefore deserves a lower premium than an asset whose exposure covaries positively with the market premium.

The Intertemporal CAPM of Merton (1973) is also a multi-period version of the CAPM, but it does not presuppose that Equation (A.2) holds, even conditionally, and it instead derives a conditional multi-factor version of this equation by assuming that investors optimize their welfare over several periods. The factors are the state variables that determine the time-varying investment opportunities, as summarized in volatilities, risk premia, correlations and interest rates.[5]

The last option is to try and improve the CAPM prediction summarized in Equation (A.2) by introducing other asset pricing factors in the right-hand side. This approach has been largely employed in empirical research since the work of Fama and French (1996), who show how certain patterns in average returns that the CAPM left unexplained could be captured by introducing "size" and "value" factors. One advantage of this method is its flexibility, since Fama and French's procedure for constructing factors as the returns on long-short portfolios can be applied to virtually any stock characteristic that is proved empirically to have an impact on expected returns: the first two examples of such sorting criteria are the market capitalization and the book-to-market ratio of stocks. The danger with this approach is the risk of factor proliferation if new factors are added on the fly every time a new "anomaly" in returns is discovered that is not explained by existing models. So caution is needed when adding empirical asset pricing factors. As pointed out by Fama (1991) and Campbell (1996), one possible safeguard is to require that factors explain not only the cross section of expected returns, but

4 - The likelihood ratio tests in his Table 1 have p-values that lead to model rejection at the 1% confidence level in five out of ten sub-periods between 1926 and 1975.
5 - Nielsen and Vassalou (2006) show that the set of relevant state variables in fact only comprises those variables that impact the short-term interest rate and the Sharpe ratio of the maximum Sharpe ratio portfolio. In other words, a state variable acts as a pricing factor only to the extent that it has an impact on one or both of these metrics.

also have some ability to predict returns, which is a restriction imposed by Merton's ICAPM.

Market as a Common Risk Factor

Even though the market factor proxied as the return on a broad equity portfolio has limited power to explain the cross-section of expected stock returns, it still clearly represents a common risk factor. This means that, by itself, it explains a substantial fraction of the time-series variance of returns, as can be seen by calculating the R^2 in a linear regression of an equity portfolio's returns on the market factor. For instance, in Table 4 of Fama and French (1993), the R^2 of the 25 portfolios sorted on size and book-to-market are all greater than 60%.

In fact, the market factor was proposed as a risk factor by William Sharpe himself even before he introduced the CAPM. As explained in Appendix A.3, the "simplified model for portfolio analysis" of Sharpe (1963) assumes that security returns are generated by a factor model with a single factor, and one possible choice for this factor is a stock market index. The model is then applied to the estimation of covariances and expected returns for the calculation of efficient portfolios in the sense of Markowitz (1952). The article does not offer direct evidence that the factor captures a large common fraction of variance, but shows that the efficient portfolios derived under the assumption of a single factor structure are close to those obtained by running an unconstrained analysis, in which this restriction is not imposed.

The Low Beta Anomaly

It was reported early on that the CAPM empirically underestimates the returns on low beta securities and overestimates those on high beta securities, so that alphas tend to be positive for the former and negative for the latter. This is still the case in Black's CAPM without a risk-free asset (see Black, Jensen and Scholes (1972) and Gibbons (1982)).

More recently, Frazzini and Pedersen (2014) conducted a systematic study of the relationship between alphas and betas in several asset classes, including US and international equities, Treasury bonds, corporate bonds, commodities and foreign exchange rates. The tendency for alpha to decrease as the beta increases is observed in all these classes.[6] It is important to note, however, that this finding applies to alphas, but not necessarily to average excess returns, which sometimes display the reverse pattern or no pattern at all. For instance, their results show that the average returns of portfolios of US stocks formed on beta tend to increase in beta (Table 3), and that the pattern is even clearer for Treasury and corporate bonds (Tables 6 and 7): leaving distressed bonds aside, average returns increase as one moves from investment-grade to high-yield bonds. These findings motivate the introduction of the "betting against beta" (BAB) factor, defined as the excess return on a zero-beta portfolio, that goes long the low beta securities and shorts the high beta ones, leveraging the long leg and de-leveraging the short one so that they both have a beta of 1. Frazzini and Pedersen (2014) report that the BAB factor earns significantly positive average returns in all asset classes.

6 - See their Figure 1.

This result is an "anomaly" for the standard CAPM, because this model predicts that the expected excess returns of both the leveraged long leg and the de-leveraged short leg should equal the market premium, so the BAB factor should earn a zero average return. The explanation proposed by Frazzini and Pedersen (2014) relies on the same general principles as Black's (1972) CAPM with restrictions on borrowing. Investors face such restrictions, so that those who have preferences for high volatility and high expected returns must pick assets with high betas rather than leveraging a portfolio of low beta securities. Thus, they bid up high beta securities, leading them to earn less per unit of systematic risk than the low beta ones do. In support for this theory, Frazzini and Pedersen (2014) find that the returns on the BAB factor are negatively related to the Treasury-Eurodollar spread: to the extent that a high TED spread means that funding is more difficult to obtain, this confirms the model prediction.[7] They also find that investors who are less likely to be leverage-constrained, like private equity funds and LBO funds, tend to hold low beta securities, while individuals and mutual funds, who are more likely to face leverage constraints, hold high beta securities.[8]

2.1.2 Size

Empirical Evidence

The size effect was one of the first reported patterns in stock returns. Banz (1981) shows that average stock returns over the period from 1936 to 1975 are significantly negatively related to firm size, even after controlling for the market beta,[9] and whether the market factor is a stock-only or a stock-bond portfolio. Moreover, a model linear in the market exposure and the firm size still leaves positive unexplained returns for small firms and negative returns for large ones, so the actual relationship between expected returns and size is not linear. A similar negative relationship between average return and firm size is simultaneously reported by Reinganum (1981), who finds that this effect "largely subsumes" the decreasing relationship between average returns and the price-earnings ratio.

As noted by Fama and French (1992), it is important to control for market beta in some way when measuring the returns on size-sorted portfolios, for betas are highly negatively correlated with size – meaning that small firms tend to have higher betas than large ones –, so it would be difficult otherwise to separate variation in returns due to size from variation due to beta. They do this by sorting stocks on beta, then on size, and they go on to show that returns tend to decrease in size within a given beta decile, confirming that the market exposure does not account for the entire return spread between small and large stocks.

International evidence for the size effect is given by Heston, Rouwenhorst and Wessels (1995) for twelve European countries, and by Liew and Vassalou (2000) for Australia, Canada, France, Germany, Italy, Japan, the Netherlands, the UK and the US.[10] The exception is Switzerland, where a negative size premium is observed.

7 - See the negative loadings on the change in TED spread in their Table 9.
8 - See their Table 11.
9 - See his Table 1. Specifically, he groups stocks every year in 25 portfolios sorted on market beta (estimated over the past five years) and size (measured at the end of the five years), and estimates the beta of each portfolio over the five years after the portfolio formation period. Next, excess portfolio returns are cross-sectionally regressed on the betas and the size (the weighted average of constituents' sizes) to produce time series of size and beta premia. Finally, the premia are regressed in time series against the excess market returns, and the alphas of the regressions are taken as estimates of the constant premia.
10 - See their Table 3, where they calculate the average returns of each country's SMB factor.

2. Efficient Factor Risk Premia Harvesting: Improving the Performance Benefits of the Performance-Seeking Portfolio with Factor Investing

Figure 1: Annual returns on portfolios formed on size from Ken French's library (in %)
(a) July 1926 - January 1980.

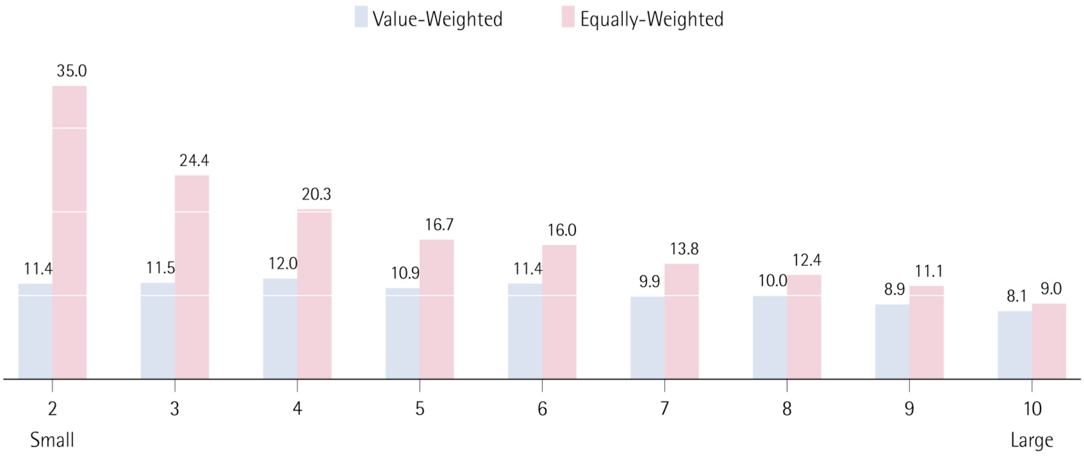

(b) January 1980 - January 2000.

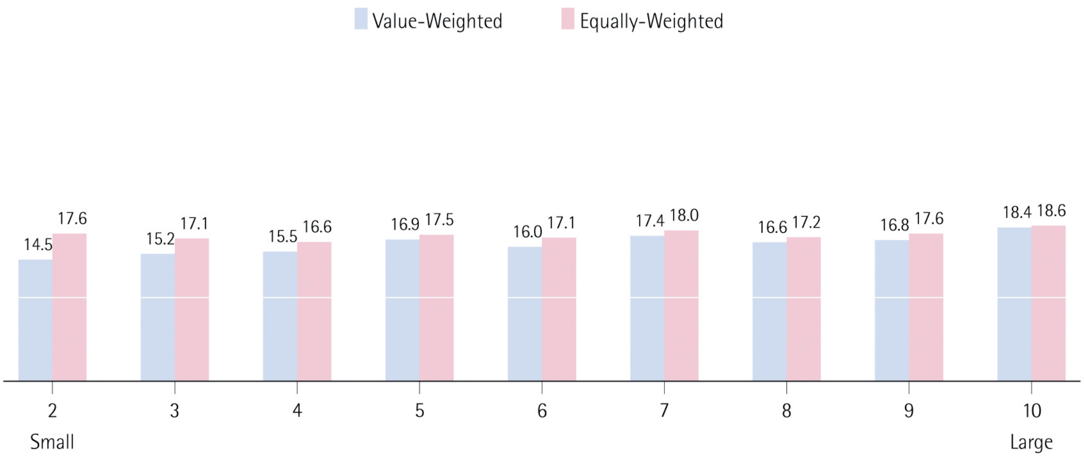

Questions on Robustness

There has been some controversy about the significance of the size effect, even in the original sample, and about its persistence after its discovery. Banz (1981) himself finds that the relationship between average returns and size, while statistically significant over the period from 1936 to 1975 and negative in all sub-periods, is not significant in all of these sub-periods. Doubt about its persistence was reinforced after the size premium apparently disappeared in the 1980s and the 1990s. For instance, Horowitz, Loughran and Savin (2000) find that small stocks underperformed large ones after 1982, and Gompers and Metrick (2001) come to the same conclusion for the period from 1980 to 1996. Figure 1 illustrates the difference between the pre-1980 sample, which roughly corresponds to Banz's original study, and the post-1980 sample, in which the relationship between size and the average return is at best flat, or even slightly increasing, thus inverted with respect to the pre-1980 period. Raw data for this figure consists of the annualized returns on portfolios formed by grouping stocks into size deciles and was obtained from Prof. French's library.[11] The base universe includes all NYSE, AMEX and NASDAQ stocks, hence it includes very small and illiquid stocks, so the returns on the portfolios of the smallest deciles, which can exceed 20% per year, should not be taken at face value. With this caveat in mind, a size pattern clearly emerges in the first sample, especially when portfolios are equally weighted, but in the subsequent twenty years, the relationship between performance and size appears to be much flatter or even perhaps slightly increasing. Incidentally, it can be noted that equally-weighted portfolios outperform their value-weighted counterparts, a well-known effect referred to as the rebalancing premium. The natural bias of equally-weighted portfolios towards small stocks inflated the returns of the portfolios formed with the smallest stocks in the pre-1980 sample, which is the period in which these stocks did outperform.

In a comprehensive study of the size effect, Van Dijk (2011) finds that the value-weighted excess returns of the smallest 20% of US stocks over the largest 20% shows pronounced cyclicality and is sometimes negative, but examination of the data for the period from 1927 to 2010 does not point to a vanishing size effect.[12] For Gompers and Metrick (2001), the underperformance of small stocks from 1980 to 1996 could be due to the growth in stock holdings by large institutional investors, who tend to favor large stocks.

In a recent contribution, Alquist, Israel and Moskowitz (2018) dispute the actual significance of the size premium in Banz's original sample, arguing that his estimates are plagued by the delisting bias documented by Shumway (1997): delisted stocks, which have large negative returns and are likely to be more represented among small than large stocks, are eliminated from the empirical analysis. Statistical concerns about studies of the size effect are not new, and were first raised right after Banz's article was published by Roll (1981), who points that because stocks of smaller firms are less frequently traded, portfolios made up of these stocks have more positively auto-correlated returns, which results in downward bias of beta estimates and thus leads to beta-adjusted returns being overstated. But Reinganum (1982) notes that Roll did not test his adjustment of betas for autocorrelation directly with firm size data (he instead studies the return spread between an equally-weighted portfolio and a cap-weighted portfolio), and concludes that the bias in betas does not completely explain the abnormal returns on small firms.

Alquist, Israel and Moskowitz (2018) also revive the debate on the seasonality of the size effect, which was first raised by Keim (1983): the latter article shows that from 1963 to 1979, the excess returns of the smallest 10% of stocks over the largest 10% is larger in January than in any other month, and in all other months the return spread

11 - Data is available at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.
12 - See his Figure 1.

often turns out to be negative.[13] The 2018 article by Alquist et al. adds to this evidence by showing that investing in the small-minus-big factor of Fama and French (1993) in January and in cash in other months produced positive cumulative returns from 1927 to 2017, while investing in the long-short factor in non-January months generated close to zero or negative cumulative returns.[14]

All in all, while some studies have questioned the existence or the significance of a size effect by raising a number of legitimate concerns about more positive studies, no definitive evidence against it has been provided to date. The size effect may be partly explained by market exposures, as the CAPM predicts, but as shown by Fama and French (1992), size is a more powerful attribute than beta when it comes to explaining the cross-section of stock returns.[15] It may have become weaker since its publication, but it is still present, even in the results of Alquist, Israel and Moskowitz (2018), who report a higher Sharpe ratio for the small-minus-big factor after its discovery.[16] Small stocks may have delivered very low returns in the 1980s and the 1990s, but as noted by Hou and Van Dijk (2019), small firms experienced negative profitability shocks and large firms positive shocks in this period, and in any case, these negative realized returns do not mean that expected returns are negative too.

Economic Rationale

Banz (1981) conjectures that small firms could earn higher returns than large firms because their stocks are less "desirable" to investors, and this lack of appeal could be attributed to the smaller amount of information that they have about small firms. A formal version of this argument is given by Merton's (1987) CAPM with limited information, in which the market alpha of a stock can depend on its volatility, its size and the percentage of the population of investors who know about it. He shows that his model can generate alphas that are decreasing in size – more specifically, in the weight of the stock in the market portfolio.[17]

> Small stocks could be stocks that investors often neglect, or stocks more sensitive to adverse economic conditions.

Another category of economic justifications for the size effect holds that size proxies for the exposure to some undiversifiable risk factor, so that investors require a premium to hold small stocks. Chan and Chen (1991) report that dividend cuts – which can be regarded as an indicator of a difficult financial situation – are more numerous in small than large firms, and that small firms tend to be more leveraged.[18] As a result, these firms are more likely to be "marginal firms", with financial and operational problems that make them more sensitive to adverse economic shocks. Fama and French (1996) conjecture that small firms are more likely to be financially distressed, and the study by Vassalou and Xing (2004) supports this view, showing that firms' distance to default – a

13 - See his Table 2.
14 - See their Exhibits 9, 10 and 11.
15 - See e.g. their Table I, where the average return of portfolios formed on beta ranges from 1.14% to 1.36% per month across deciles, while that of portfolios formed on size ranges from 0.89% to 1.52%, and their Table AIII, where the relationship between average returns and betas is positive but statistically insignificant, while the relationship with the logarithm of the market cap is negative and significant.
16 - See their Exhibit 6.
17 - See Equations (31.a), (31.b) and (31.c) of his article, and the analysis pp. 496-7.
18 - See their Table III.

measure inspired by the corporate bond pricing model of Merton (1974) – tends to increase with size.[19] They also find that when the aggregate distance to default is added to Fama and French's three factors to construct a pricing model, this variable captures some of the role of the small-minus-big (SMB) factor, which suggests that the two variables share some informational content.

2.1.3 Value

Broadly speaking, the "value effect" refers to the fact that, on average, stocks with a "low" price relative to some accounting metrics outperform those with a "high" price. In fact, the literature has reported several of these effects, depending on how the "cheapness" of a stock is measured.

Empirical Evidence

The most established measure in academic studies is the book-to-market value of equities (BE/ME), since the work of Fama and French (1992) and Fama and French (1993), which led to the introduction of the three-factor model named after the authors. It appeared well before, in studies by Stattman (1980) and Rosenberg, Reid and Lanstein (1985). The latter article shows that a zero-dollar strategy that went long stocks with a high BE/ME and shorted those with a low BE/ME produced average monthly returns that were significantly positive from 1973 to 1984.[20] Fama and French (1992) provide concurring evidence by showing that the returns on decile portfolios sorted on BE/ME increase almost monotonically with this ratio,[21] and by exhibiting positive premia associated with BE/ME in Fama-MacBeth regressions.

These regressions aim to estimate the coefficients Λ_k in models of the form described by Equation (A.1), and, in this case, are performed by replacing the betas in the right-hand side by observable characteristics, namely BE/ME or firm size. The sensitivity of expected returns to BE/ME turns out to be even greater in magnitude than the (negative) sensitivity to size.[22]

Fama and French (1993) introduce the "high-minus-low" (HML) factor, defined as the excess return of a portfolio of high BE/ME stocks over one made up of low BE/ME ones, with weights proportional to market cap, and Liew and Vassalou (2000) provide international evidence for a value premium by showing that the HML factor has positive and significant returns (at the 5% level, with t-statistics greater than 2) in Australia, Canada, France, Germany, Italy, Japan, Switzerland, the UK and the US.[23] Fama and French (2012) also measure the returns of the HML factor in North America, Europe, Japan and Asia Pacific, and find that the average return is positive in all four zones.[24] Lakonishok, Shleifer and Vishny (1994) provide concurring evidence, showing that when stocks are sorted in deciles on BE/ME and grouped in portfolios, the returns in each of the five years after the formation date increase almost monotonically with BE/ME.[25]

Other ratios than BE/ME mixing an accounting variable and a measure of the market value of a

19 - See Panel B of their Table VI.
20 - The weighting rules, described in p. 10 of their article, are rather complex, and correspond to a minimum variance portfolio subject to a number of neutrality constraints, with respect to 55 industry groups and several macro and microeconomic variables (including firm size) suspected to have an impact on returns.
21 - See their Table IV.
22 - See their Table III. To be exact, the characteristics against which returns are regressed are the logarithms of BE/ME and the market cap.
23 - See their Table 3.
24 - See their Table 1.
25 - See Panel A in their Table I.

firm have been proposed to explain differences in average returns. The price-earnings ratio (P/E) is one of them. Basu (1977) shows that sorting stocks on their (P/E) ratio and grouping them in portfolios produces average returns that decrease in the P/E, but the market betas of the portfolios do not line up with these average returns.[26] So stocks with higher earnings relative to price tend to earn higher subsequent returns. Lakonishok, Shleifer and Vishny (1994) consider the cash flow-to-price ratio (C/P) and growth in sales (GS) over the past five years, in addition to BE/ME and P/E. Across the four variables, GS has a special status since it is the only one that is a pure accounting variable, which does not involve any market value. These alternative metrics are interesting to consider in place of, or in conjunction with, BE/ME: a sort on C/P produces a wider return spread than a sort on BE/ME,[27] and GS maintains the ability to discriminate between stocks within BE/ME groups, with the low GS portfolio outperforming the high GS one.

Fama and French (1993) show that the P/E effect is explained by their three-factor model: portfolios sorted on P/E have roughly the same market betas, but the betas with respect to the SMB and HML factors are decreasing in P/E, so the average return is decreasing too.[28] Similarly, Fama and French (1996) find that the three-factor model explains the return spreads between portfolios sorted on C/P or past five-year sales growth: the beta with respect to HML increases with C/P and decreases with sales growth, so the model predicts higher returns for the higher C/P portfolios and for the lower sales growth portfolios. The vast majority of the residual alphas are statistically insignificant.[29]

Economic Rationale

Like for other anomalies, explanations can fit into rational asset pricing theory or invoke irrational behavior among investors. Fama and French (1993) suggest an interpretation of the value effect in terms of distress: low BE/ME stocks would be "stocks with persistently high earnings on book equity that result in high stock prices relative to book equity" (p. 50), while high BE/ME stocks would be stocks with low earnings and low prices. In support of this claim, Fama and French (1995) report that stocks with high BE/ME on a given date tend to have a lower ratio of earnings to book equity from four years before the measurement date to five years after it.[30] Vassalou and Xing (2004) provide some evidence supporting the idea of a connection between the value effect and distress risk, by showing that book-to-market ratios align well with distances to default at the market level and that portfolios closer to default deliver higher average returns.[31] However, the alignment between ratios and the distance to default is not observed in all sub-universes, in particular in groups of stocks with lower default risk,[32] so BE/ME is not just a proxy for the default likelihood.

The risk-based interpretation is disputed by Lakonishok, Shleifer and Vishny (1994), who note that value strategies outperform "glamour" strategies even during NBER recessions and bear market years.[33] In their analysis, value portfolios

26 - See his Table 1.
27 - Compare Panels A and B in their Table I.
28 - See their Table 11.
29 - See their Table III.
30 - See their Figure 1.
31 - See their Table III
32 - See Panels B and C of their Table V.
33 - See their Figure 2.

are invested in stocks with high C/P and low GS, and glamour portfolios in stocks with low C/P and high GS. More generally, they find no evidence for an association between the relative return of value to glamour strategies and the general direction of the equity market (bullish or bearish).[34] They conclude that the underperformance of growth stocks is more likely due to investors' failure to correctly anticipate mean reversion in earnings after a period where earnings grew. Instead, investors unduly extrapolate the high past earnings of growth stocks to conclude that these stocks are good investment opportunities, and therefore neglect value stocks.

The behavioral model of Barberis and Huang (2001) combines features of consumption-based asset pricing models with behavioral biases, namely "narrow framing" and "loss aversion", which cause investors to worry about gains and losses on individual stocks in addition to their consumption level. The authors show that the equilibrium prices generated by this model exhibit a dividend-price pattern: stocks with a higher dividend-price ratio have higher average returns and higher alphas with respect to the market factor.[35]

Among risk-based theories, a sub-class of models emphasizes the role of time-varying risk. The idea is that value stocks might not be fundamentally riskier than growth firms in all market conditions, but they might be in "bad" economic times, when investors are least willing to endure losses, and their higher average return would be a reward for living with this undesirable characteristic. But "risk" and "bad times" must first be defined. The aforementioned study by Lakonishok, Shleifer and Vishny (1994) considers the ex-post market return or NBER definition as indicators of the latter, but the exact definition in asset pricing theory is times of high marginal utility, and Petkova and Zhang (2005) argue that ex-post returns are not an accurate proxy for marginal utility (p. 198). They propose instead to identify bad economic times as those with a high ex-ante market premium. Combining the idea of a time-varying market premium with the CAPM's definition of risk as the covariance with the market portfolio, several papers have studied conditional versions of the CAPM to see whether the value premium can be explained by these models. In the conditional CAPM, both the beta and the conditional market premium can vary over time, and the expected return of an asset not only depends on its expected beta and the unconditional market premium, but also on the covariance between its time-varying beta and the conditional market premium. If the beta-premium covariance is positive, the expected return is higher than for an otherwise identical asset for which the covariance would be zero. Lettau and Ludvigson (2001b) show that a conditional CAPM is indeed able to explain the returns on portfolios sorted on size and book-to-market about as well as the three-factor model.[36] They show that except within the segment of largest stocks, value stocks have higher consumption betas than growth stocks in "bad" times identified by high equity market premium, i.e. by higher risk aversion on the part of the representative investor.[37] Petkova and Zhang (2005) also find evidence in favor of

34 - See their Table VII.
35 - See their Table IV, which is based on simulated data.
36 - Compare subfigures b and d in their Figure 1.
37 - See their Table 5.

an explanation based on time-varying risk, by reporting that the conditional beta of value stocks covaries positively with the market premium, while that of growth stocks covaries negatively.[38] However, they find that this covariance is too small in magnitude to fully explain the empirical value premium.

A difficulty inherent in conditional models is how the information set is modeled.

A difficulty inherent in conditional models is that the conditional market premium and the conditional betas are unobservable, so they must be estimated by specifying a conditioning information set. As a matter of fact, not all studies follow the same modeling approach, and they sometimes reach different conclusions regarding the ability of the conditional CAPM to account for the value premium. Lettau and Ludvigson (2001b) use a weighted average of aggregate consumption, aggregate financial wealth and aggregate income, but Lewellen and Nagel (2006) use another approach that does not require the specification of state variables, by assuming that betas are stable within a month or quarter: under this assumption, conditional alphas can be estimated by running ordinary linear regressions of stock returns on the market returns over a month or quarter. The latter authors find a significantly positive alpha for a value-minus-growth strategy, which contradicts the conditional CAPM.[39] Petkova and Zhang (2005) return to an approach with state variables to model the conditional market premium,[40] and they reach a less negative conclusion, which is that the conditional CAPM reduces, but does not eliminate, the alphas of value-minus-growth strategies.[41]
Zhang (2005) also invokes time-varying risk, but his model is developed on the producer side rather than on the investor side, so the origin of risk lies in the production technology. All firms have assets in place and growth options, but value firms have more of the former and growth firms have more of the latter. In general, capital stocks are more costly to reduce in bad times than they are to expand in good times, so value firms face bigger difficulties in bad times when they want to adjust their production than growth firms do in good times when they exercise their growth options. This implies a bigger dispersion of risk between value and growth firms in bad times. As Zhang notes, the presence of time variation in the market price of risk reinforces this effect, as the premium is higher in bad times, thereby depressing the discounted values of future projects, and encouraging value firms to disinvest even more.

2.1.4 Momentum

Empirical Evidence

The seminal study of the momentum effect in stock returns is that of Jegadeesh and Titman (1993). Specifically, they consider "relative strength portfolios", defined as long-short portfolios in which stocks with the highest past returns are purchased and stocks with the lowest past returns are sold short. Both the ranking

38 - See their Table 2. There are a few exceptions depending on the sample and the way conditional betas are measured, with a negative covariance for value stocks and a positive one for growth stocks.
39 - See their Table 3.
40 - The state variables are the dividend yield, the default spread, the term spread and the T-bill rate.
41 - Compare their Tables 1 and 3.

period, used to calculate past returns and sort stocks, and the holding period, over which the long-short portfolio is held, range from 1 to 4 quarter(s), and the full sample spans the period from January 1965 to December 1989. For all combinations of period lengths, the relative strength portfolio is shown to have positive performance, most often statistically significant. Carhart (1997) provides similar evidence for mutual funds invested in equities, rather than individual stocks.[42]

The horizon is important. When stocks are ranked over longer periods, a reversal is observed, instead of a continuation. The long-term reversal effect was documented before momentum by DeBondt and Thaler (1985): when past returns are measured over periods from one to five years, portfolios formed with past losers tend to outperform the market, defined as a broad equity portfolio with equal weights, while portfolios of past winners tend to underperform.[43] The authors also show that the excess returns of the loser portfolio are greater than the excess returns of the market over the winner portfolio, so the profitability of a long-short strategy that exploits reversal is more due to the short side. At short horizons, Jegadeesh (1990) documents a reversal in stock returns, with a negative serial correlation between returns in consecutive months. Conrad and Kaul (1998) make a comprehensive study of the performance of strategies buying winners and selling losers, letting the ranking period range from one week to three years and taking an equally long holding period.[44]

The momentum effect has proved to be remarkably robust over time and across geographic zones. First of all, it survived its discovery, as shown by Jegadeesh and Titman (2001), who document that portfolios of stocks sorted on their past 6-month return yield returns that are almost monotonically decreasing from past winners to past losers, both in the pre-1989 period, which corresponds to the sample of Jegadeesh and Titman (1993), and in the subsequent period from 1990 to 1998. Their results show no evidence of a weakening momentum effect, even for large stocks, which are in principle cheaper to trade and thus more subject to hunting by investors seeking to take advantage of an anomaly.[45] Momentum has also withstood very long-term tests. Israel and Moskowitz (2013) find that a long-short strategy that selects stocks based on their past twelve-month returns, goes long the top 10% of past winners and shorts the bottom 10% of past losers, and holds the portfolio for one year earns positive and significant returns in different subperiods between 1926 and 2011. Evidence is stronger than for size and value, whose premium, although it remains positive, fails to be significant in some subperiods.[46] Geczy and Samonov (2016) conduct a test on the US stock market from 1801 to 2012, and also find evidence for momentum. International evidence in twelve European countries is provided by Rouwenhorst (1998) for the period from 1980 to 1995: a strategy that goes long the top 10% of winners and shorts the bottom 10% of losers, earns positive average returns.[47]

The outperformance of past winners is not captured by the three-factor Fama-French model,

42 - See the monthly excess returns on portfolios of funds sorted on past twelve-month returns in his Table III.
43 - See their Table I.
44 - See their Table 1.
45 - See their Table I.
46 - See their Table 1.
47 - See their Table III.

as Fama and French (1996) show: when regressed on the three factors, portfolios of stocks sorted on past twelve-month returns still have statistically significant alphas that increase in the past return rank, as the average returns do. What is worse, the model does not reduce the anomaly: market betas do not exhibit a well-defined pattern, while the betas with respect to the SMB and HML factors tend to be higher, not lower, for past losers than for past winners. So the model seems to predict lower returns for past winners than for past losers, the exact opposite of what is empirically observed. The situation is different for the long-term reversal effect, when portfolios are formed on the basis of past five-year returns: alphas are insignificant, and the SMB and HML betas are lower for portfolios invested in past five-year winners, thus predicting lower returns for these portfolios than for those made of past losers. As a result, the model produces a range of alphas that is broader than the range of average returns.[48] In other words, the model exacerbates, rather than reduces, the anomaly.

The inability of the three-factor model to explain the outperformance of past short-term winners is the motivation to introduce a fourth factor, which Carhart (1997) defines as the return of an equally-weighted portfolio of the 30% of stocks with the highest past one-year returns, minus the return of a portfolio of the 30% of stocks with the lowest returns. He shows that the four-factor model better explains the differences in the average returns of portfolios of mutual funds sorted on past performance than the three-factor model does. Nevertheless, the model does not account for the underperformance of the worst-return funds, for which statistically significant and negative alphas remain.[49] Prof. Ken French publishes a momentum factor in his online data library, but he weights stocks by their market value and combines the momentum sort with a sort on size to avoid any implicit size bias in the long and short legs of the factor.

Economic Rationale

Behavioral explanations have often been proposed to explain the profitability of momentum strategies. For instance, Jegadeesh and Titman (1993) invoke "delayed price reactions to firm-specific information" (p. 67), suggesting that news about firms is only gradually reflected in prices, and winners continue to be winners and losers losers during the information absorption process. This explanation is consistent with the view that positive autocorrelation in returns arises because of underreaction, and negative autocorrelation due to overreaction. But the challenge for any economic theory of the momentum effect is that it must reconcile it with reversal at long (three years and longer) or very short (from one week to one month) horizons. Daniel, Hirshleifer and Subrahmanyam (1998) propose such a model, based on overconfidence of investors in their ability to estimate expected returns, and on biased self-attribution: investors tend to attribute past success to their skills and past failure to external causes. In their model, momentum arises because investor confidence augments when investment choices prove to be successful: this way, investors overestimate their ability to pick winners, so they continue to

48 - In their Table VI, the average excess returns on decile portfolios range from 0% per month for past losers to 1.31% for past winners, while in Table VII, the three-factor alphas range from −1. 15% to 0.59% for the same portfolios.
49 - See his Table III.

buy these stocks, which raises their prices until expectations eventually adjust.

Hong and Stein (1999) also develop a model in which short-term underreaction and long-term overreaction coexist but arise from the slow diffusion of information among two different types of investors, as opposed to being the consequence of cognitive biases: "newswatchers" process information about future fundamentals but not about past prices, while "momentum traders" trade based on an analysis of past prices. In this model, positive news about a firm is partially absorbed by newswatchers, raising the price of its stocks, but not to the point where it reaches its new long-run value, and momentum traders subsequently use this price signal to buy more stocks, until the price exceeds the long-run value. At this time, exploitation by momentum traders of the initial underreaction eventually results in overreaction. Hong, Lim and Stein (2000) find empirical support for this theory, showing that momentum strategies are more profitable in small stocks and in stocks with low analyst coverage, two categories of securities for which information is likely to spread more slowly.[50]

While limited rationality and market frictions are popular explanations of the momentum premium, other studies have proposed theories that are compatible with market efficiency. Conrad and Kaul (1998) provide a mathematical decomposition of the expected returns on strategies buying winners and selling losers and show that there are two sources of profits, namely the auto-correlation in stock returns and the cross sectional dispersion in expected returns.

So, predictability can imply positive average profits, but returns do not need to be predictable for the strategy to be profitable. As a matter of fact, Conrad and Kaul show that the cross sectional variance in expected returns accounts for a significant fraction of profits.[51]

In other papers, momentum profits arise as the reward for bearing undiversifiable risk, especially risk of a macroeconomic nature. Chordia and Shivakumar (2002) note that momentum strategies tend to pay off in expansionary phases and deliver negative returns in recessions,[52] which is suggestive of a link between momentum returns and the marginal utility of consumption. They also show that after adjusting returns over the ranking period for exposures to a set of macroeconomic variables, momentum strategies yield insignificant returns, which they interpret as evidence for the role of time-varying expected returns in the profitability of these strategies.[53] Bansal, Dittmar and Lundblad (2005) show that the return spread between the past winner and past loser portfolios is well explained by the betas with respect to aggregate consumption of growth rates in security cash flows.[54] Liu and Zhang (2008) consider the growth rate in industrial production, a variable taken from the macroeconomic asset pricing model of Chen, Roll and Ross (1986), and find that past winners are indeed more exposed than past losers,[55] although differences in exposures tend to vanish seven months after the portfolio formation and reverse twelve months

50 - See their Table III for momentum profits by size: the excess return of past winners over past losers (P3 – P1) is decreasing in size, when the smallest 20% of stocks are excluded. See their Table IV for profits by analyst coverage, and Table V for profits by size and coverage.
51 - See their Table 2.
52 - See their Table II.
53 - See their Table III. The macroeconomic predictors of stock returns are the dividend yield, the three-month T-bill rate, the term spread and the default spread.
54 - See the high (greater than 60%) R2 of cross sectional regressions of average returns on betas in their Table IV, and the good alignment of fitted expected returns with realized average returns in their Figure 4.
55 - See Panel C in their Table 1.

after formation.[56] They calculate that the Chen, Roll and Ross model explains 60% to 90% of the profitability of the winners-minus-losers portfolio.[57]

2.1.5 Low Volatility

The CAPM predicts a positive relationship between the expected return and the volatility of efficient portfolios. Indeed, the expected excess return of a portfolio, whether or not it belongs to the efficient frontier, is an increasing function of its market beta, and if it is efficient, it has no idiosyncratic risk, so its volatility is proportional to its beta. As a result, the expected excess return of an efficient portfolio is an increasing linear function of its volatility. This theoretical prediction is in line with common sense, which suggests that risk-averse investors demand compensation in the form of higher expected return for holding more risky assets. However, empirical evidence is much more mixed, and the famous study by Ang et al. (2006) even points out that the relationship between risk and return within stocks may in fact be inverted. This "low volatility" effect has encouraged a lot of empirical work and is still the subject of debate.

Empirical Evidence

Some of the early studies of the CAPM specifically focused on the relationship between expected return and volatility, as opposed to that between expected return and beta, and found empirical support for the CAPM's prediction that expected return will increase in volatility. Sharpe (1965) analyzes a sample of 34 mutual funds, assumed to follow efficient diversification strategies,[58] finds a significant positive relationship between their average returns and their volatility[59] and concludes that "the data do lend considerable support to the theory tested" (p. 422). Soldofsky and Miller (1969) conduct a study at the asset class level, with government bonds, corporate bonds, preferred stocks and common stocks, and also find a positive and significant slope when regressing average returns on volatilities, though the coefficient appears to decrease over time.[60] However, examination of their Chart 1 suggests that the positive relationship is largely driven by the presence of two groups of securities, namely bonds and stocks. Indeed, bonds are less volatile than stocks and earn lower average returns, which lends support for the CAPM's prediction. But within stocks, the risk-return relationship appears to be much flatter or even slightly negative. Confirming this observation, Haugen and Baker (1996) estimate the expected returns on individual stocks by relating them to 41 characteristics through a model similar to the Barra risk model, and they report higher volatilities among securities with the lowest expected returns than among those with the highest.[61]

The seminal study on "low volatility puzzles" is by Ang et al. (2006), who report that the 20% of stocks with the lowest total volatilities or idiosyncratic volatilities with respect to the Fama-French model underperform the 20% of stocks with the highest risk, and that this return spread is explained neither by the CAPM nor by the Fama-French model since the same pattern is observed in alphas. These models do not even reduce the anomaly, and instead tend to reinforce it, since there is a bigger spread in alphas than

56 - See their Table 2.
57 - See their Table 6.
58 - "Lacking any satisfactory test for reasonable diversification, all 34 funds were assumed to have chosen efficient combinations of securities." (p.417).
59 - See his Figure 1.
60 - See the estimated slopes in their Table 2.
61 - See their Table 2, which shows that realized returns are positively related to expected returns, and their Table 3 for expected returns. However, volatility is not monotonically decreasing across expected return deciles.

in average returns.[62] The authors also show that while stocks with large idiosyncratic volatility could be more exposed to aggregate volatility risk, and exposure to aggregate volatility is negatively priced in the cross section,[63] controlling for volatility exposure leaves the return spread between low and high risk stocks almost unchanged.[64] International evidence is reported by Blitz and Van Vliet (2007) for the US, Europe and Japan, and by Ang et al. (2009) for Canada, France, Germany, Italy, Japan and the UK.[65] In Blitz and van Vliet's study, the effect appears to be quantitatively comparable to the value and the momentum patterns, when looking at the spread of alphas from the Fama-French model.

Questions on Robustness

Perhaps because it contradicts the intuition that higher average returns come as a reward for higher risk taking, the volatility effect has been particularly controversial in the literature, and a number of studies have produced discordant findings. Bali and Cakici (2008) point out that the "IV puzzle" – the short name for the "idiosyncratic volatility puzzle" – seems to lack robustness with respect to a number of measurement details. Specifically, the significant negative relationship between idiosyncratic risk and average returns is obtained for value-weighted portfolios when volatility is calculated from daily returns in the previous month, as in Ang et al.'s (2006) article, but it becomes insignificant when volatility is based on past monthly returns,[66] or when volatility is daily but portfolios are equally-weighted or weighted by the inverse of idiosyncratic volatility.[67] The authors also show that the anomaly is only significant in small stocks and insignificant in large ones,[68] and they argue that the results of Ang et al. (2006) are driven by small and illiquid stocks. Similarly, Martellini (2008) finds a positive rather than a negative relationship between the total volatility and the average return when volatility is calculated from past monthly returns and stocks are grouped in equally-weighted quintile portfolios.[69]

Other studies have discussed the choice of the volatility measure. Fu (2009) argues that idiosyncratic risk varies substantially over time, so the sample volatility of idiosyncratic returns over a given past period is not a good estimate of forward-looking idiosyncratic volatility, which is the measure that appears in Merton's (1987) model predicting that specific risk has an effect on expected returns. So he proposes an exponential generalized autoregressive conditional heteroskedasticity (EGARCH) model to estimate it and finds a positive relationship between expected returns and expected future idiosyncratic volatility. However, Guo, Kassa and Ferguson (2014) argue that this result is driven by a look-ahead bias that is present in the estimation procedure for the EGARCH model because the volatility in month t estimated in month $t-1$ is based on parameters estimated over a sample that includes month t. When removing this bias, they find that the relationship between expected returns and expected future volatility is still positive, but statistically insignificant.[70]

62 - See their Table VI.
63 - They estimate the price of volatility risk to lie between–0. 08% and–0. 07%, depending on what other factors are included in the pricing model (see their Table V).
64 - Compare the difference in the alphas of the low and high risk groups in Tables VI and IX.
65 - See Exhibit 3 in Blitz and Van Vliet (2007) and Table 2 in Ang et al. (2009).
66 - See their Tables 3 and 4.
67 - See their Tables 1 and 8.
68 - See their Table 6.
69 - See his Exhibit 1.
70 - Compare the t-statistics in their Table 5, which uses the out-of-sample estimates, to those in their Table 4, which is based on the in-sample estimates.

The lack of robustness with respect to measurement details poses a challenge to potential theories of the volatility effect, as they have to reconcile a number of apparently contradictory findings. Huang et al. (2010) show that the presence of short-term reversals in returns can explain the difference in returns between equally-weighted and value-weighted portfolios. Indeed, equally-weighted portfolios re-allocate gains from past winners to past losers, so they take advantage of return reversals, while value-weighted portfolios overweight past winners, which are likely to be the next period's losers. These authors show that in the high idiosyncratic volatility group, value weighting substantially underperforms equal weighting, such that the value-weighted portfolio of the most risky stocks underperforms the value-weighted portfolio of the least risky stocks, although equally-weighted returns are roughly independent of volatility.[71] They argue that this can be attributed to the large proportion of winners and losers contained in the high idiosyncratic volatility universe.

Economic Rationale

In frictionless asset pricing models, investors are assumed to be rational and to diversify away any idiosyncratic risk, so they neither require nor pay an extra premium for holding securities with higher specific risk. In the presence of frictions, specific risk can have an impact on the expected return, but the first models to incorporate such features predict a positive relationship between the two. For instance, in Merton's (1987) CAPM with limited information, investors hold a security only if they know about it, so not all of them hold the market portfolio. Merton shows that for such a stock, the market alpha increases with the specific volatility, defined here with respect to a common production factor. This reinforces the low volatility puzzle.

One possible explanation for the volatility effect is that investors do not have the type of concave preferences typically assumed in asset pricing models and that, for instance, they value lottery-like payoffs, which provide them with opportunities for large gains. Barberis and Huang (2008) develop a pricing model based on the cumulative prospect theory, in which agents have preferences for positive skewness, and they show that this can lead to high prices for securities with a small probability of a "jackpot". High volatility stocks could be examples of such securities, as noted by Blitz and Van Vliet (2007) and Baker, Bradley and Wurgler (2011). Mitton and Vorkink (2007) explore the asset pricing implications of the existence of "Lotto investors", with a positive preference for skewness, alongside investors with traditional mean-variance preferences, and they show that the former agents can choose portfolios that are inefficient in the mean-variance sense, but they are compensated by higher skewness. Several aspects of institutional investment may also contribute to explain the anomaly. Following the above lottery argument, Baker and Haugen (2012) suggest that the demand from money managers for volatile stocks may be explained by compensation rules: the typical compensation package has an option-like profile, with the probability of a bonus being higher if the manager chooses a high volatility portfolio. The authors

71 - See their Table 3.

also show that more volatile stocks receive higher coverage by analysts and media, and so are more likely to be recommended by analysts to CIOs. Finally, most managers are concerned about their relative performance with respect to a benchmark, and Blitz and Van Vliet (2007) note that investing in high volatility or high beta stocks appears to be the most straightforward way to generate outperformance. Baker, Bradley and Wurgler (2011) develop a similar argument, and conclude that the practice of benchmarking tends to exacerbate the demand for the most risky stocks.

2.1.6 Investment and Profitability

These factors are more recent to academic research and investment practice than size, value and momentum.

Theoretical Foundation

The dividend-discount model of Gordon and Shapiro (1956) is often invoked as an ex-ante justification to argue that expected returns should be linked to investment and profitability. Fama and French (2006) note that the price of a firm's stock, P, can be written as a function of its future dividends, $\mathbb{E}[D]$, and the expected return, r, as

$$P_t = \sum_{h=1}^{\infty} \frac{\mathbb{E}_t[D_{t+h}]}{[1+r]^h}. \tag{2.2}$$

They then use the accounting equation that says that dividend per share equals earnings per share, Y, minus the change in book equity per share, and they derive the following equation (their Equation (3)):

$$\frac{ME_t}{BE_t} = \frac{1}{BE_t} \sum_{h=1}^{\infty} \frac{\mathbb{E}_t[Y_{t+h} - [BE_{t+h} - BE_{t+h-1}]]}{[1+r]^h}. \tag{2.3}$$

As explained by Fama and French, this equation implies that holding BE/ME and the growth in book equity fixed, higher future earnings relative to book equity, i.e. higher expected future profitability, results in a higher discount rate and thus in higher expected returns: this is the profitability effect. Similarly, other things being equal, a higher growth in book equity that would follow from higher investment decreases the discount rate: this is the investment effect. These arguments are re-used by Novy-Marx (2013).

Asness, Frazzini and Pedersen (2019) introduce the related notion of "quality": a "quality stock" is defined as one with desirable characteristics, including profitability. They consider a special case of Equation (2.2) in which expected future dividends grow at the constant rate g, so that after a few algebraic manipulations, Equation (2.2) becomes

$$P_t = \frac{\mathbb{E}_t[D_{t+1}]}{r-g},$$

which they rewrite as

$$\frac{ME_t}{BE_t} = \frac{1}{r-g} \times \frac{\text{dividends}}{\text{profits}} \times \frac{\text{profits}}{BE_t}. \tag{2.4}$$

This equation shows that the market value increases with the expected growth rate in dividends, with the profitability (measured as the profit-to-book ratio) and with the payout rate (the dividends-to-profits ratio), and decreases with the expected rate of return, hence with risk to the extent that riskier firms command a higher expected return.

2. Efficient Factor Risk Premia Harvesting: Improving the Performance Benefits of the Performance-Seeking Portfolio with Factor Investing

Empirical Evidence

Unlike for the size and the value attributes, which are conventionally identified in the literature with the market capitalization and book-to-market ratio, there seems to be no consensual measure of profitability and investment, and the metrics vary across papers. Titman, Wei and Xie (2004) are concerned with the question whether increases in investment have an impact on subsequent returns, and they measure "abnormal capital investment" by comparing capital expenditure in the current year to the average capital expenditure over the past three years. They find that portfolios of low investment stocks outperform those made up of high investment stocks.[72] In Fama and French (2006), investment is measured as growth in total assets, which "[the authors] judge gives a better picture of investment" than the change in book equity, although the latter is the metrics that appear in Equation (2.3). Because the valuation equation (2.3) involves expected asset growth, not past growth, they also test a forward-looking measure, which is the fitted value from a regression to predict asset growth. Returns are significantly negatively related to past asset growth, but are insignificantly positively related to expected growth.[73]

Fama and French (2008) also use asset growth as a proxy for investment, but they conduct a study by size group, and they find that the asset growth effect is highly significant in micro stocks, significant in small stocks, and absent in large stocks,[74] but given the numerical importance of small stocks, it is still detected at the market level.

Differences also arise in the choice of metrics when assessing the profitability of a firm. Fama and French (2006) consider the earnings-to-book ratio, which they find bears a significantly positive relationship with average returns. This significance subsists when the past ratio is replaced with an expected value.[75] Novy-Marx (2013) argues that "[g]ross profitability is the cleanest measure of true economic profitability" (p. 2), and so should be preferred to earnings, and that book assets are a better deflator than book equity because they are independent of leverage. In fact, he shows that gross profits-to-assets have a more significant impact on the returns of industry portfolios than the alternative measures defined as the earnings-to-book and free cash flow-to-book ratios. Like Fama and French (2008), he finds that the profitability effect is present in NYSE stocks, and the excess return of long-short strategies is more significant in small than in big stocks,[76] but unlike Fama and French (2008), he finds that the long-short strategies deliver positive excess returns even in large stocks. A third measure is proposed by Fama and French (2014), who consider operating profitability, calculated as revenue minus costs, expenses and interest payments, all divided by the book value of equity, and investment is the past percentage change in total assets. Like the aforementioned studies, they find that sorts on these characteristics produce greater return spreads among small stocks than large stocks.[77] Finally, Asness, Frazzini and Pedersen (2019) use a composite measure of profitability, obtained by aggregating gross profits over assets, return on equity, return on assets, cash flows over assets,

72 - See their Table 1. However, average returns are not monotonic across deciles of capital investment, since the bottom portfolio underperforms the second decile portfolio.
73 - See their Table 3.
74 - See their Table IV.
75 - See their Table 3.
76 - See his Table 4.
77 - See their Table 1.

gross margin and accruals. They report a significant positive effect of the profitability score on average returns,[78] and they subsequently use this score as a component of a "quality" score that also includes the other dimensions of quality suggested by Equation (2.4), namely growth, safety and payout.

The profitability effect is not explained by the Fama-French market, size and value factors. This can be seen from the results of Novy-Marx (2013), who finds that a strategy that goes long stocks with the highest gross profits-to-assets and shorts those with the lowest ratio delivers positive and significant alphas. These alphas are all the more puzzling for the three-factor model since the long-short strategies have negative betas with respect to the positively rewarded factors, such that the alphas are greater than the average returns themselves.[79]

Motivated by this evidence, Fama and French (2014) introduce two additional factors, defined as the excess returns on portfolios that go long stocks with the highest operating profitability or the lowest investment, and short the mirror groups. The observation that the profitability effect does not have the same intensity across all size groups suggests that it is useful to control for size, and possibly for book-to-market, in the construction of the factors. So the factors are first defined separately for small, big, value, medium and growth stocks, and the market-level factors are obtained by averaging these components. The five-factor model is rejected in the GRS test run by Gibbons, Ross and Shanken (1989), but with lower confidence levels than the three-factor model, which suggests that it better captures patterns in average returns.[80] In a related effort to complete the three-factor model, Hou, Xue and Zhang (2015) develop a "q-factor model" that includes the market factor, a size factor, an investment factor and a return-on-equity factor. The definition of the last two factors follows Fama and French's general approach to defining factors as the excess returns on zero-dollar portfolios. The authors find that their model has similar performance to Carhart's four-factor model in explaining the returns on portfolios sorted on size and book-to-market, but performs better when it comes to explaining the investment effect and the profitability effect, here understood as a return-on-equity effect.

2.2 Construction of Efficient Equity Factor Benchmarks

Shortcomings of Cap-Weighted Indices

Cap-weighted indices are traditionally used as benchmarks by equity investors. They have clear practical advantages. First, they provide a good representation of the equity market by placing greater weights on larger firms, whose movements have more impact on the market. Second, they require minimal rebalancing because changes in weights replicate changes in market capitalization without requiring purchases or sales of stocks. The limit to this rule, and the source of turnover in cap-weighted indices, is the turnover of the underlying universe. Stocks can be deleted from the universe, e.g. because of a merge or

78 - See their Table III.
79 - See his Table 2.
80 - See their Table 5.

failure to meet inclusion criteria, while new stocks can be listed. Third, cap-weighted indices have advantages regarding liquidity because by definition they assign larger weights to larger constituents, which reduces the risk of hitting limits on market capacity. In addition, stocks of large firms tend to be more liquid than those of small firms.[81] Lastly, the construction procedure of cap-weighted indices relies only on observable quantities, namely the market capitalization of each stock, or actually its free-float capitalization. The free-float factor in the calculation of capitalization eliminates stocks that are not available for trading in markets, like those held for strategic purposes by governments or other companies.

Cap-weighted indices are dominated by a few very large stocks.

From a theoretical perspective, weighting stocks by capitalization may seem to be an approach to proxy for the CAPM's market portfolio, in which assets are weighted by their capitalization. But the theoretical justification is elusive. The true market portfolio should encompass all assets in the universe, including non-tradable ones, and even equity investors hold non-equity assets in the form of human capital. Moreover, the market portfolio is optimal under the restrictive assumptions of the CAPM, but many of these are violated in practice: for instance, investors have non-homogenous expectations and horizons and they can rebalance their portfolios.

The economic significance of these departures from the theoretical framework is at first glance unclear, but simple empirical tests show that the risk-return profile of cap-weighted indices is easy to improve, e.g. by weighting stocks equally. As an illustration, the first two rows of Table 1 highlight that an equally-weighted index of the 500 largest US stocks earned an average return 159 basis points above that of the cap-weighted version from 1970 to 2016. Equal weighting is even more parsimonious than cap weighting when it comes to parameter estimation since it requires only the number of securities in the portfolio, excluding any other data, and can therefore be regarded as even simpler than cap weighting.

In fact, a cap-weighted index often fails to match our intuitive notion of a "well-diversified" portfolio because market caps are highly dispersed, thereby giving rise to a few very large weights that result in an index dominated by a few large stocks. This intuition can be laid out in mathematical terms by introducing a quantitative measure of the dispersion of weights, like the Herfindahl index, which is the sum of squared weights, and the "effective number of constituents", defined as the reciprocal of the Herfindahl index. It can be shown that the effective number of constituents (abbreviated as ENC) is maximal for the equally-weighted portfolio and can be close to its minimum of 1 for a concentrated cap-weighted portfolio.

81 - The relationship between firm size and transaction costs has been documented, among others, by Stoll and Whaley (1983), who show that the bid-ask spread and commission rate tend to be larger for smaller stocks (see their Table 5), and Pastor and Stambaugh (2003), who find that their liquidity measure is also positively related to size (see their Table 9). They estimate liquidity by measuring the relationship between a stock's daily excess return over the market and the past day's volume multiplied by the sign of the past day's excess return: the idea is that a stock with low liquidity that outperforms the market on a given day is expected to underperform on the next day by an amount that increases with the current day's volume and the degree of illiquidity. Alquist, Israel and Moskowitz (2018) also document a positive relationship between size and various liquidity measures such as trading costs, the bid-ask spread and price impact (see their Exhibit 20).

2. Efficient Factor Risk Premia Harvesting: Improving the Performance Benefits of the Performance-Seeking Portfolio with Factor Investing

The excessive concentration of cap-weighted indices is not their only shortcoming. Indeed, no attempt is made in their construction procedure to provide exposure to one or more of the rewarded risk factor(s) identified by academic research on equity investing and presented in Section 2.1, so they can only be exposed to these factors "by chance" in specific subperiods, but not in a stable and lasting manner. The case can even be made that some of the factor exposures embedded in cap-weighted indices have the "wrong" sign since they result in lower expected returns. For instance, cap weighting implies a natural bias towards growth stocks – stocks with a large market value relative to book value – while empirical evidence is in favor of a premium for stocks with a low market-to-book ratio (see Section 2.1.3). These shortcomings of cap-weighted indices are also discussed by Amenc et al. (2014b).

Efficient equity factor indices aim to address them. In what follows, we outline the construction procedure advocated by Amenc, Goltz and Lodh (2012), Amenc and Goltz (2013) and Amenc et al. (2014b), in which the stock selection and weighting procedures are separated.

Factor Tilts Through Stock Selection

The first step in the construction of efficient equity indices is the selection of stocks from an underlying universe, so as to achieve exposure to a given factor. All the factors listed in Section 2.1 relate to stock characteristics, so sorting stocks on one of these characteristics produces a sort on a factor exposure: smaller stocks are more exposed to the size factor, stocks with a high book-to-market ratio are more exposed to the value factor and so on. Exposures may not strictly line up with characteristics at the stock level, meaning that the two sorts are not perfectly equivalent, but the approach of grouping stocks in halves or thirds smooths some of the ranking errors, making the group of the x% with the highest/lowest characteristic very close to the group of the x% with the top/bottom exposure.

Sorting is easier when the characteristic is readily observable, which is the case for many factors. The market cap, book-to-market ratio, past twelve-month return, return on equity and total asset growth are observable, though they require accounting data, access to which is costly. The only factor based on a non-observable characteristic is low volatility, because volatility is not a market or accounting attribute, but the volatility of a stock can be estimated from its return history, provided the sampling frequency is sufficiently high (see Appendix A of Merton (1980)).

Efficient Diversification Through Stock Weighting

When stocks have been selected from the universe, the second step is to choose a weighting scheme. The challenge at this stage is to diversify away unrewarded risk by exploiting the risk and return parameters of the stocks, including the correlation structure between them. In theory, the optimal weighting scheme is the maximum Sharpe ratio portfolio, but it involves expected returns, which are difficult to estimate. Sample means provide imprecise estimates, so economic models relating expected return to some measure of risk are more promising. One approach is to postulate a factor model and to express each expected excess return over the risk-free rate as a sum of factor premia weighted by factor exposures, as in Equation (A.1) of the Appendix: in this model, risk has multiple dimensions and is measured by a set of factor betas.

Alternatively, one can relate expected returns to some one-dimensional risk metrics. Martellini (2008) suggests using the total volatility of stocks, assuming that expected excess returns are proportional to this metric, and he shows that the resulting estimated maximum Sharpe ratio portfolios do have higher ex-post Sharpe ratios than the cap-weighted portfolio and even than the equally-weighted one, at least when the covariance matrix is estimated in a structured way:[82] the structured estimator can be the constant correlation estimator used in Elton and Gruber (1973), or the matrix implied by the Fama-French model. But there are reasons to believe that investors also care about higher order moments. Indeed, Scott and Horvath (1980) mathematically show that under mild conditions, investors have a positive preference for skewness and a negative preference for kurtosis. Empirically, Mitton and Vorkink (2007) find that investors who hold ill-diversified portfolios in the mean-variance sense are prone to purchasing stocks with high skewness. Preferences for higher-order moments appear more indirectly through asset pricing implications: extracting these moments from option prices, Conrad, Dittmar and Ghysels (2013) show that stocks with higher skewness tend to have lower subsequent returns, while those with higher kurtosis display higher returns.[83]

Building on these insights, Amenc et al. (2011) assume that expected returns are proportional to downside risk as measured by the semi-deviation of returns,[84] and they show that the "efficient index" thus constructed consistently outperforms the cap-weighted index across regions.[85]

For those who do not want to take a stand on expected returns, a number of alternative diversification schemes bypass the estimation of these parameters. Assuming that all stocks have the same expected return, maximizing the Sharpe ratio is equivalent to minimizing volatility, so the global minimum variance (GMV) portfolio is optimal. A potential concern is that in the presence of a large number of constituents and long-only constraints, the GMV portfolio can exclude many stocks to favor the least risky ones. The risk parity method of Maillard, Roncalli and Teiletche (2010) avoids the concentration issue by imposing the restriction that all constituents must contribute equally to the risk of the portfolio, which guarantees that all weights are non zero. When stocks are uncorrelated, or more generally when they have equal pairwise correlations, this weighting method is equivalent to weighting stocks by the reciprocal of their volatility.

Amenc, Goltz and Lodh (2012) show that a weighting scheme consistent with the diversification objective targeted by an investor is more effective at reaching this objective than a selection of stocks with characteristics that are deemed to serve the objective. Indeed, a suitable diversification methodology better captures the interactions across securities than selection, which

82 - See his Exhibit 3.
83 - See their Table II.
84 - The squared semi-deviation is the expected squared difference between the mean return and returns when this difference is non-negative.
85 - See their Exhibit 11.

focuses on the standalone characteristics of stocks. The importance of using a sound weighting rule is re-affirmed by Amenc et al. (2016), who show that score-weighting approaches that target the highest factor exposure, hence the highest reward, leave a sizable amount of unrewarded risk (see also Amenc et al. (2017)).

Risk-Return Profile of Equity Benchmarks

Table 1 displays the annualized return, volatility and Sharpe ratio of a range of US equity indices calculated by Scientific Beta over a 46.5-year period that starts in June 1970 and ends in December 2016. The broad cap-weighted index has the lowest annualized return (10.59%) and the lowest Sharpe ratio (0.33) across the 28 indices, and in each sub-universe, made up of half the stocks contained in the broad universe, it is also the cap-weighted portfolio that features both the lowest performance and the lowest Sharpe ratio.

Each sub-universe is defined by a stock selection that exposes the portfolio to an equity factor known for being rewarded in the long run. The empirical evidence summarized in Table 1 shows that both deviations from broad Cap-weighted indexation help to improve average returns and the Sharpe ratios, since the cap-weighted index of each universe dominates the broad cap-weighted index with respect to both criteria, and the non-cap-weighted indices bring further improvement.

2.3 Allocation to Efficient Equity Factor Benchmarks

Several reasons can motivate the construction of multi-factor equity portfolios. First, investors may believe in the persistence of multiple factor premia and not want to give up on some by selecting a single factor index. Second, they may want to benefit from performance smoothing over time due to the imperfect correlation between the cycles of various factors. As long as different factors have different "bad times", investors can expect to mitigate the impact of short-term losses on their portfolios. From a theoretical standpoint, investing in multiple factors is justified if more than one factor is needed to explain the cross section of expected returns, and the empirical evidence accumulated so far points in this direction. Indeed, Martellini and Milhau (2015) show that if the cross section is explained by a set of pricing factors and these factors are the excess returns on portfolios (possibly involving short positions), then the maximum Sharpe ratio that could be earned by investing in individual securities can also be collected by constructing the maximum Sharpe ratio portfolio of the pricing factors. This result implies a reduction in the dimension of the optimization problem: as opposed to maximizing the Sharpe ratio over hundreds of individual stocks, one can maximize it over a handful of factors, and there is no loss of optimality provided these factors explain the cross section of returns. But as usual, one has to know which factors to invest in. The factors enumerated in Section 2.1 go a long way towards explaining known facts about the cross section, but the list is likely still open, so the aforementioned

Table 1: Performance and risk indicators for US cap-weighted or smart-weighted equity indices, June 1970 to December 2016

Universe	Weighting	Annualized return (%)	Volatility (%)	Sharpe ratio
All	Cap-weighted	10.59	16.74	0.33
	Equally-weighted	12.18	16.56	0.43
	Minimum variance	12.58	14.23	0.53
	Maximum Sharpe ratio	12.59	15.44	0.49
Mid cap	Cap-weighted	12.50	17.03	0.44
	Equally-weighted	12.96	16.78	0.48
	Minimum variance	13.61	14.05	0.61
	Maximum Sharpe ratio	13.37	15.53	0.54
Value	Cap-weighted	11.60	16.71	0.40
	Equally-weighted	13.34	16.30	0.51
	Minimum variance	13.59	14.17	0.61
	Maximum Sharpe ratio	13.89	15.37	0.58
High momentum	Cap-weighted	11.45	17.13	0.38
	Equally-weighted	12.75	16.96	0.46
	Minimum variance	13.86	14.88	0.60
	Maximum Sharpe ratio	13.26	16.04	0.52
Low volatility	Cap-weighted	10.86	14.99	0.39
	Equally-weighted	12.71	13.69	0.56
	Minimum variance	13.01	12.69	0.63
	Maximum Sharpe ratio	13.06	13.16	0.61
Low investment	Cap-weighted	11.90	15.58	0.44
	Equally-weighted	13.47	15.20	0.56
	Minimum variance	14.08	13.49	0.67
	Maximum Sharpe ratio	13.90	14.50	0.61
High profitability	Cap-weighted	11.03	16.99	0.36
	Equally-weighted	12.69	17.04	0.45
	Minimum variance	13.02	14.56	0.55
	Maximum Sharpe ratio	13.14	15.78	0.52

Notes: Equity indices are constructed from US stocks. The "All" universe contains the 500 large stocks, and the others are made of the 250 stocks with the smallest capitalization, lowest investment (past 2-year asset growth), highest momentum score (return over past 52 weeks minus most recent 4 weeks), highest gross profitability (gross profit divided by total assets), lowest past 2-year volatility or highest book-to-market ratio. Cap-weighted indices are weighted by market capitalization, minimum variance indices are weighted so as to minimize the ex-ante portfolio variance, and maximum Sharpe ratio indices are designed to maximize the ex-ante Sharpe ratio. Non-negativity constraints are applied to rule out short positions, and liquidity and turnover adjustments are made in order to ensure the investability of all indices. Data is from the long-term track records of the Scientific Beta database. The start date is 19 June 1970, and the end date is 31 December 2016. Details on the universe construction and index calculation can be found in Scientific Beta (2018, 2019). In maximum Sharpe ratio indices, expected returns are assumed to be proportional to the semi-standard deviations of stocks.

result does not imply that mixing only these six factors yields as high a Sharpe ratio as what would be achieved by working at the stock level.

A simpler question is whether factor combinations deliver improvement over the broad cap-weighted index and the individual factors. Amenc et al. (2014a) introduce a framework for allocation to multiple factor benchmarks, and multi-factor allocations are also studied by Martellini and Milhau (2015) and Martellini and Milhau (2018a). Table 2 displays performance, risk and diversification metrics for such strategies, invested in six equity factor indices exposed to the six rewarded factors of Section 2.1, namely size, value, momentum, volatility, investment and profitability. A description of the multi-factor strategies follows:

- The equally-weighted portfolio allocates 1/6 to each factor;
- The minimum volatility portfolio, which is borrowed from mean-variance theory, minimizes the ex-ante volatility, subject to long-only constraints;
- The risk parity portfolio in the sense of Maillard, Roncalli and Teiletche (2010) equates the contributions of the six factors to portfolio volatility. It is long only by construction;
- The minimum tracking error portfolio minimizes the ex-ante tracking error under long-only constraints;
- The relative risk parity equates the contributions to the tracking error with respect to the broad cap-weighted index. Like the risk parity portfolio, it is long only by construction;
- The maximum ENB portfolios maximize the "effective number of bets" (ENB) in the portfolio volatility, as explained in what follows. The general idea is to have all risk factors contribute equally to volatility, but the risk factors in question are not specified ex-ante and are treated as implicit. They are estimated either by principal component analysis (PCA), or by minimum linear torsion (MLT). The calculation of ENB is described below;
- The maximum relative ENB portfolios maximize the ENB in tracking error.

Except for the equally-weighted one, all these portfolios require an estimate for a covariance matrix, either that of returns or that of returns in excess of the broad index. This estimate is updated at each rebalancing date – every quarter – and is obtained by taking the sample covariance matrix of returns or excess returns over the past 104 weeks.

Diversifying Across Implicit Risk Factors

The ex-ante portfolio variance can be decomposed across the six factors as a quadratic function of weights with coefficients given by the entries of the covariance matrix. Implicit factors are a linear transformation of the original factors, so they fully explain the variance without any residual risk. But unlike the original factors, which are correlated, they have a zero correlation. As explained in Appendix A.3, there are (at least) two ways of extracting these factors from the data, namely PCA and MLT as in Meucci, Santangelo and Deguest (2015), the main difference being that PCA aims to maximize the explanatory power of the uncorrelated factors with respect to the original factors, while MLT aims to minimize the distance between the original factors and the new ones. Martellini and Milhau (2018b) give an explicit expression for the linear transformation that maps the original factors into the new ones in the MLT, and Appendix C of this paper does the same for PCA factors.

Once a set of factors has been fixed, one can calculate the ex-ante exposures of the portfolio by weighting the constituents' exposures. With six implicit factors, there are six portfolio exposures, β_{p1} through β_{p6}. Because the factors are uncorrelated, the portfolio variance is the sum of factor variances times squared exposures:

$$\sigma_p^2 = \beta_{p1}^2 \sigma_{F1}^2 + \cdots + \beta_{p6}^2 \sigma_{F6}^2.$$

Each term in the sum is a contribution to portfolio variance, and the relative contribution of the k^{th} factor is

$$q_k = \frac{\beta_{pk}^2 \sigma_{Fk}^2}{\sigma_p^2}.$$

These contributions are all non-negative and add up to 100%.

The "effective number of bets" (ENB), defined in Carli, Deguest and Martellini (2014), is an indicator constructed to quantify the extent to which the contributions of the uncorrelated factors are unbalanced. It is calculated by taking the reciprocal of the sum of squared relative contributions. It can be shown that the maximum ENB is the nominal number of portfolio constituents – 6 in this case – and is achieved when all contributions are equal, a situation known as "factor risk parity". More generally, a portfolio with balanced exposures will have a large ENB, i.e. close to 6, and one with concentrated exposures will have an ENB close to 1. In Table 2, ENBs are expressed as percentages of the nominal number of constituents, so they vary from 16.67% to 100%.

The same procedure can be applied to the portfolio tracking error, to analyze the contributions of factors to relative risk. To this end, the PCA and the MLT algorithms are run on the relative covariance matrix, which is the covariance matrix of relative returns with respect to the broad equity index, and this produces a different set of six factors than the factors coming from the analysis of the covariance matrix of returns. The calculation of relative contributions and the ENB is then identical to the absolute case. For clarity, we refer to the ENB in tracking error as the "relative ENB".

Because we restrict all portfolios in Table 2 to have non-negative weights, factor risk parity may not be attainable. Martellini and Milhau (2018b) measure the leverage of factor risk parity portfolios and show that it can be sizable. We therefore consider instead "maximum ENB" and "maximum relative ENB" portfolios, calculated by maximizing the ENB subject to the long-only constraints. As a result, their ENBs and relative ENBs are less than 100%, but their weight profiles are more realistic than those of portfolios that would take short positions on some factors.

Measuring Diversification Across Assets

An easy-to-calculate diversification metric is the "effective number of constituents" (ENC), defined as the reciprocal of the sum of squared weights. It can be shown that for long-only portfolios, it is always comprised between 1 and the number of constituents, so here the ENC divided by 6 is always between 16.67% and 100%. The maximum is attained by the equally-weighted portfolio.

The ENC does not take into account the differences in volatilities across constituents, so the risk of a portfolio may be dominated by some constituents, usually the most risky ones, even if dollar weights are equal. The "effective number of correlated bets" (ENCB) corrects this shortcoming by measuring the dispersion of contributions to volatility. The starting point is the decomposition of variance across constituents and the definition of the contribution for each constituent, as in Maillard, Roncalli and Teiletche (2010). The ENCB is defined like the ENB and ENC by taking the reciprocal of the sum of squared relative contributions. Its maximum of 6 is attained when all assets contribute equally to volatility, i.e. when the portfolio achieves risk parity. The "relative ENC" is defined in the same way, but replacing the contributions to volatility with the contributions to tracking error.

Properties of Multi-Factor Strategies

First, Table 2 confirms some of the results from Table 1 over a slightly different sample period, shorter by two years. All factor indices (considered here in their cap-weighted version) outperform the broad index by 28 to 187 basis points, and although some have higher volatilities, all but the high profitability index have higher Sharpe ratios. Nevertheless, they have substantial tracking errors, ranging from 3.40% to 5.75% per year, and they sometimes experience large losses with respect to the broad index, as appears from the maximum relative drawdowns, which can be as high as 38.39%.

> The minimum volatility portfolio has the lowest volatility, but displays large risk relative to the cap-weighted index.

Interestingly, the mere fact of putting factors together can reduce the volatility risk, though this does not work for any weighting scheme: the equally-weighted portfolio has a tracking error of 1.83% and a relative drawdown of 12.08%, and if an additional effort is made to minimize the tracking error, the scores fall to 0.98% and 3.08% respectively. The minimum volatility portfolio is a counterexample, since it has a bigger relative drawdown than any individual constituent, at 41.69%, and its tracking error is a large 4.48%. But several of the multi-factor portfolios display higher information ratios than the single factor indices. The best two ratios are achieved by the equally-weighted and risk parity portfolios: at 0.61 and 0.63, they largely overwhelm the best information ratio of a single index, which is 0.38, for the low investment factor. In absolute terms, the benefits of combining factors are less obvious: annual returns range from 10.61% to 11.29%, which is similar to the scale from 10.35% to 11.96% for individual constituents. Sharpe ratios also roughly fall within the same

Table 2: Multi-factor equity strategies; data from September 1972 to December 2016.

	Cap-weighted	Mid-Cap	Value	High-Mom	Low-Vol	High Prof.	Low Inv.		
Ann. ret. (%)	10.09	11.96	11.37	10.98	10.37	10.35	11.50		
Volatility (%)	16.51	17.85	16.43	17.07	14.71	16.75	15.37		
Sharpe ratio	0.34	0.39	0.39	0.35	0.37	0.32	0.43		
Max. drawdown (%)	53.78	58.28	61.58	50.44	49.09	52.78	52.57		
Tracking error (%)	–	5.75	5.25	4.05	4.89	3.40	3.71		
Info. ratio	–	0.33	0.24	0.22	0.06	0.08	0.38		
Max. rel. drawdown (%)	–	28.57	32.60	18.71	38.39	17.91	25.64		
ENC (%)	–	16.67	16.67	16.67	16.67	16.67	16.67		
ENCB (%)	–	16.67	16.67	16.67	16.67	16.67	16.67		
Rel. ENCB (%)	–	16.67	16.67	16.67	16.67	16.67	16.67		
ENB PCA (%)	–	18.52	21.89	21.40	27.28	31.95	35.56		
ENB MLT (%)	–	61.71	76.84	73.21	86.43	73.05	89.44		
Rel. ENB PCA (%)	–	30.27	26.88	35.21	49.47	41.33	34.10		
Rel. ENB MLT (%)	–	18.89	26.63	24.48	30.96	38.07	32.02		
	Equally-weighted	Min vol	Risk parity	Min TE	Relative risk parity	Max ENB PCA	Max ENB MLT	Max rel. ENB PCA	Max rel. ENB MLT
Ann. ret. (%)	11.21	10.99	11.24	10.61	10.77	11.01	11.28	11.29	10.63
Volatility (%)	15.82	14.59	15.70	16.23	16.16	15.94	14.80	15.85	16.03
Sharpe ratio	0.40	0.41	0.40	0.35	0.36	0.38	0.43	0.40	0.35
Max. drawdown (%)	52.14	43.58	51.71	52.84	52.77	53.48	46.19	51.11	52.24
Tracking error (%)	1.83	4.48	2.01	0.98	1.08	2.22	3.44	2.08	1.48
Info. ratio	0.61	0.20	0.57	0.53	0.63	0.42	0.35	0.58	0.37
Max. rel. drawdown (%)	12.08	41.69	13.44	3.08	2.69	17.98	26.50	11.38	12.31
ENC (%)	100.00	24.49	99.18	61.14	77.11	59.94	48.34	63.96	76.96
ENCB (%)	99.14	24.49	100.00	59.06	73.83	59.71	49.48	63.47	75.04
Rel. ENCB (%)	42.41	21.60	41.91	61.75	100.00	29.50	32.63	34.42	59.39
ENB PCA (%)	26.27	24.90	26.24	25.85	26.10	89.99	27.13	27.70	25.75
ENB MLT (%)	92.56	93.57	93.74	87.89	88.73	90.31	99.31	90.34	90.15
Rel. ENB PCA (%)	41.60	39.86	41.62	40.17	43.00	39.51	39.95	94.00	42.42
Rel. ENB MLT (%)	60.90	32.08	58.39	51.53	79.04	43.94	41.53	52.23	97.84

Notes: Except for the portfolio fully invested in the broad cap-weighted index, all portfolios are invested in six cap-weighted equity factor indices: mid-cap, value, high momentum, low volatility, high profitability and low investment. All portfolios are rebalanced quarterly and are subjected to long-only constraints. The minimum tracking error portfolio ("Min TE") aims to minimize the tracking error with respect to the broad index, and the relative risk parity equates the contributions of constituents to the ex-ante tracking error. The maximum ENB portfolios minimize the dispersion of factor contributions to volatility (i.e., they maximize the effective number of bets), where the factors are obtained from asset returns by principal component analysis (PCA) or minimum linear torsion (MLT). The maximum relative ENB portfolios minimize the dispersions of contributions to tracking error. The tracking error, the information ratio and the maximum relative drawdown are with respect to the broad cap-weighted index. For the portfolios invested in one or more factor(s), the effective number of constituents (ENC), the effective numbers of correlated bets in volatility (ENCB) and in tracking error (relative ENCB), and the effective numbers of bets are calculated at each rebalancing date based on post-rebalancing weights and the estimated covariance matrices of returns, excess returns over the broad index and factors. They are expressed as percentages of the nominal number of constituents, which is 6. The table displays the averages of these quantities across the rebalancing dates. Raw data (the returns on factor indices) is from the Scientific Beta database.

range, from 0.35 to 0.43, versus 0.32 to 0.43 for factors taken in isolation. Maximum drawdowns are also similar in size. It is therefore in terms of relative metrics with respect to the broad index that multi-factor portfolios post a clear advantage over their constituents.

The diversification indicators reported here (ENC, ENCB and ENB) are in line with the definition of each portfolio. The equally-weighted portfolio has an ENC of 100%, and the risk parity one an ENCB of 100%, and the maximum ENB portfolios dominate the others in terms of their respective metrics. The minimum volatility portfolio is the most concentrated by the ENC and ENCB criteria, because it assigns a large weight to the low volatility index.[86]

86 - The ENC and ENCB of the minimum volatility portfolio are seen to be equal in Table 2. It can be shown that in this portfolio, the contribution to risk of each asset is equal to its dollar weight, and that this property holds both for the long-short minimum variance portfolio and the portfolio subject to non-negativity constraints on weights.

3. Efficient Factor Exposure Matching: Improving the Hedging Benefits of the Liability-Hedging Portfolio with Factor Investing

3. Efficient Factor Exposure Matching: Improving the Hedging Benefits of the Liability-Hedging Portfolio with Factor Investing

As explained in the introduction, the liability-driven investing (LDI) paradigm is intended for investors who face liabilities, and it recommends investing in two building blocks. The first is a performance-seeking portfolio (PSP), whose objective is to deliver the highest reward per unit of risk, and the second is a liability-hedging portfolio (LHP), which should maximize the absolute correlation with the present value of liabilities. While this approach is often followed in institutional asset-liability management, it can also be used in individual money management by replacing liabilities with goals. The counterpart of the LDI principle in this context is known as goal-based investing (GBI), and the LHP is referred to as the goal-hedging portfolio (GHP).

This section describes the construction of the GHP for a goal that is relevant to almost any individual: to generate income in retirement, so as to replace labor income. Part of the replacement income is provided by social security systems and employer-sponsored pension plans, which make up the first two pillars of pension systems in the World Bank classification, but the general trend towards population aging means that a reduction of social security benefits is likely to occur, whether it takes the form of tighter conditions to claim pensions or pension cuts for retirees. On the pension fund side, traditional defined-benefit pension plans, in which benefits are calculated as a function of an employee's earnings and years of service, have been experiencing a relative decline with respect to defined-contribution arrangements, where benefits are determined by an individual's contributions and the returns on their investments. The growing importance of DC plans is illustrated in Figure 2 with figures borrowed from the OECD report 2016 *Pensions Outlook*. With the exceptions of Iceland and New

Figure 2: Percentage change in the number of members of DB and DC pension arrangements. Source: OECD (2016)

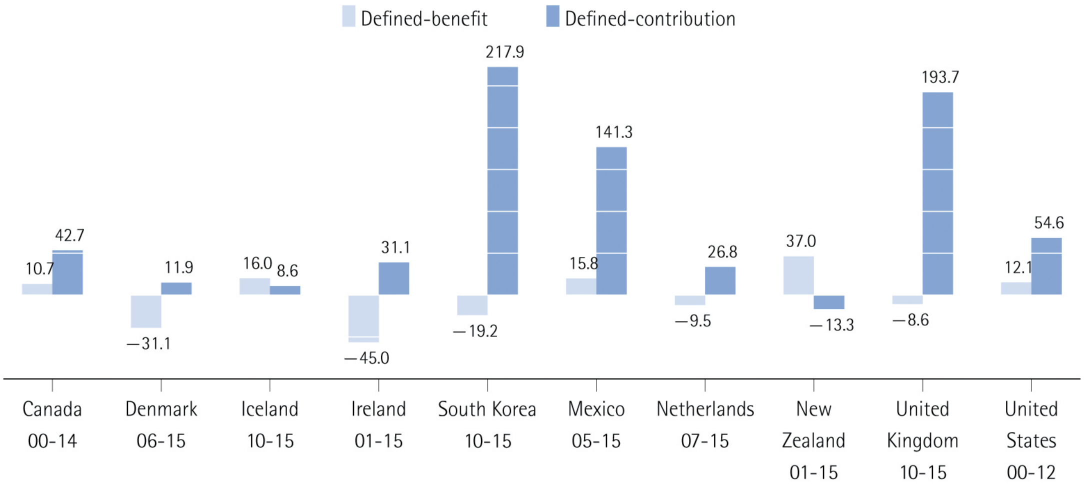

Notes: The period over which the change is measured is country-specific, depending on data availability, and is specified after the country name. Raw data is obtained from the OECD and is available at http://dx.doi.org/10.1787/888933426787.

Zealand, the number of people covered by a DC plan has increased more in the recent past than the number of people enrolled in a DB plan, and in several countries such as Denmark, Ireland, South Korea and the UK, the size of DB plans measured based on the number of members has actually decreased.

Given the expected reduction in social security benefits and the rising importance of DC pension arrangements, individuals are increasingly responsible for making their own savings and investment decisions for retirement. This situation should fuel interest in investment strategies based on the principles of GBI, which combine a PSP and a GHP, the GHP here being defined as a portfolio that produces fixed replacement income in retirement. The construction of the GHP is the focus of this section.

3.1 Goal-Hedging Portfolio in Retirement Investing

3.1.1 Annuities

In principle, annuities are the perfect risk-free asset for individuals seeking a stable income in retirement. The simplest form of contracts, known as lifetime annuities, pay a predefined income for an individual's lifetime, in exchange for a premium. Other contracts pay an income until the passing of the individual or their spouse, whichever comes last, or guarantee income for a minimum amount of time, with continuity of benefits for heirs in the event that the primary beneficiary passes before the term. Some contracts are deferred, meaning that the payments start at a predetermined future date: here, the premium per dollar of income is lower than for an immediate annuity due to the discounting effect. The key selling point of annuities is that they solve the decumulation problem: given a nest egg at retirement, how much should be withdrawn every year so as not to outlive one's savings, and so as to leave as little unspent surplus at the time of death? Individuals making plans for retirement are faced with a dilemma between a low spending rate, insufficient to support their desired lifestyle, and the risk of running short of money. Annuities solve this problem by providing income that is known ex-ante as long as an individual lives.

In spite of these advantages, voluntary annuity purchases are not widespread. According to the US Individual Annuity Sales Surveys conducted by the LIMRA Secure Retirement Institute, sales of fixed immediate annuities in 2018 were $B9.7. A variety of reasons can explain the low demand for annuities; these are reviewed in Benartzi, Previtero and Thaler (2011) and Pashchenko (2013). Some relate to psychological barriers. Individuals are afraid of leaving money on the table if they die earlier than expected, but the risk of making a bad deal in case the event that is a risk from the insurer's standpoint does not occur – here, unexpectedly long life – is inherent to any insurance contract. It is by pooling individual risks that the insurer can pay income to those who live longer than the average population. As a consequence, the annuity market is affected by adverse selection issues, with people endowed with the shortest life expectation being reluctant to annuitize. Annuities are also perceived as relatively complex products, whose price is not easy to check for potential customers. Indeed, annuity prices involve mortality assumptions and ad-hoc adjustments for trends in longevity, and various types of fees charged for increased flexibility or additional guarantees, like a minimum payment period. These fees may be hidden in market quotes, making them difficult to gauge for clients. More observable – for they are documented in annuity contract rules – are the surrender fees charged for exiting annuities, which can reach several percentage points of the invested capital for the first years of the contract.

Overcoming all these obstacles is not a straightforward task, and one can ask in the first place to what extent increasing the demand for annuities is a desirable objective. In this regard, Scott, Watson and Hu (2011) prove a mathematical result that corresponds to a simple intuition: an annuity contract dominates (in the sense of being less expensive) a financial security that pays income for a fixed period only if the writing costs do not outweigh the additional discount due to the application of survival probabilities. In other words, if survival probabilities are close to 100% in the first years of retirement, writing costs must be very low in order to justify the purchase of an annuity in place of a bond ladder that delivers income for this period. It is only when survival probabilities become sufficiently small that the discount from giving up on income distribution after one's death dominates the costs. As pointed out by Benartzi, Previtero and Thaler (2011), it is therefore hard to make a case for period certain annuities, since consumers pay a load for a series of certain payments that requires no insurance against longevity risk and could be obtained from financial securities.

3.1.2 A New Financial Security: The Retirement Bond

Figure 3 shows how the cash flows of an immediate annuity purchased at the retirement date can be decomposed into two series of cash flows, assuming that the individual lives at least for twenty years after he/she retires. The first twenty cash flows are those of a retirement bond, and those that occur thereafter are paid by a "deferred annuity". The deferred payments start twenty years after retirement, which is defined here as "late life". This definition is based on the fact that a further twenty years is approximately the life expectancy of a 65-year-old US individual from the general population, but the 20-year threshold can be modified in other geographic areas to reflect different longevity conditions. We do not claim that the best threshold coincides with life expectancy at retirement. As explained above, it depends on the costs of annuity contracts beyond the fair actuarial price. Life expectancy has the advantage of being a simple measure, which can be easily obtained from public demographic datasets without needing a database of annuity quotes.

The retirement bond is defined as the financial security that pays $1 of replacement income every year, starting at the retirement date, for a predefined period. Thus, it has four characteristics:

3. Efficient Factor Exposure Matching: Improving the Hedging Benefits of the Liability-Hedging Portfolio with Factor Investing

Figure 3: Decomposition of an immediate annuity into a retirement bond and a deferred annuity

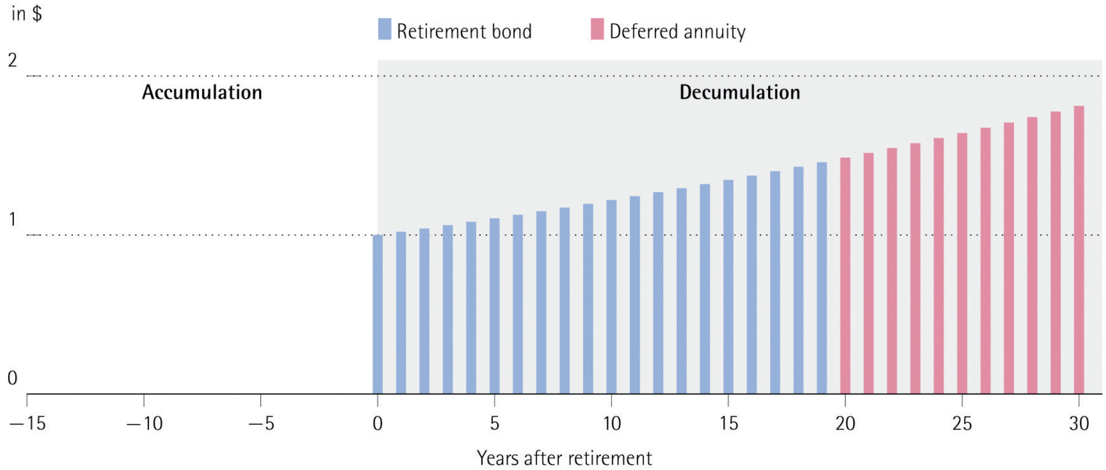

Notes: For the purpose of the illustration, it was assumed that the decumulation period after retirement date is 30 years long, but it is actually uncertain because of longevity risk. Payments are normalized to $1 per year, and 2% annual growth is applied as a cost-of-living adjustment.

1. The geographic area, which determines the discount rates to use;
2. The retirement date;
3. The duration of the period during which it pays income, referred to as the "decumulation period" in what follows;
4. The cost-of-living adjustment.

The last parameter is introduced because cumulative inflation over a typical retirement planning horizon can severely affect the purchasing power of savings: were income constant in nominal terms, the owner of the bond would be able to purchase much less at the end than at the beginning of the decumulation period. The ideal solution to deal with inflation risk would be to have payments indexed on realized inflation, like in French OATi, in UK inflation-linked Gilts and in US Treasury inflation-protected securities. A mid-term approach is to apply a constant growth rate to annual payments, as a provision for expected inflation. This fixed rate represents a constant cost-of-living adjustment (COLA). The advantage of a fixed COLA over full inflation indexation is that the nominal cash flows of the bond are known in advance, so the bond price depends on nominal interest rates, and it can in principle be replicated with fixed-income securities. (To see if this is feasible in practice is the goal of the subsequent analysis.) With inflation-linked payments, real cash flows would have to be discounted at the real interest rates, and inflation-indexed bonds would have to be used for replication purposes. The market for these securities is not as large as for nominal bonds, with fewer security profiles available, so a fixed-income manager would likely prefer to work with nominal bonds. With a fixed annual COLA, bond cash flows have the increasing pattern depicted in Figure 3.

3. Efficient Factor Exposure Matching: Improving the Hedging Benefits of the Liability-Hedging Portfolio with Factor Investing

A series of recent papers have made a case for the issuance of these bonds by referring to them under various names, such as "Bonds for Financial Security" (Muralidhar 2015; Muralidhar, Ohashi and Shin 2016), "Standard of Living indexed, Forward-starting, Income only Securities" (Merton and Muralidhar 2017), or a merge of the two, "BFFS/SeLFIES" (Kobor and Muralidhar 2018).

Synthesizing Retirement Bonds

As the safe asset for retirees and future retirees, retirement bonds are said to have a certain social usefulness, thus justifying their issuance by sovereign states. Muralidhar, Ohashi and Shin (2016) enumerate arguments in favor of these bonds from the issuer's standpoint: they improve retirees' welfare, interest payments are deferred in time, and the cash flow schedule somewhat mirrors that of infrastructure projects, which makes them suitable for funding such projects.

Of all the securities available in financial markets, only standard coupon-paying bonds have predictable cash flows, up to the default risk of the issuer, so they may be the best approximation to retirement bonds. However, there are two important differences between their cash flow schedule and that of retirement bonds. First, these cash flows are not equal in size. Until redemption, the bond pays coupons, generally at an annual or semi-annual frequency, equal to a percentage of the principal, and at the maturity date it repays the principal. This results in a series of cash flows unequal in size, since the last payment dominates the others. Second, coupon payments to the bondholder start no later than twelve or six months (depending on the coupon frequency) after the bond purchase, so they cannot be deferred until the retirement date. As a result, the investor pays for cash flows that occur in the accumulation phase and are not needed. For these reasons, a buy-and-hold position in a coupon-paying bond, even a default-free one, is not a good substitute for a retirement bond.

One can in principle synthesize a retirement bond by constructing a cash-flow matching portfolio, which consists of zero-coupon bonds with staggered maturity dates. The purchased amount of each bond must be such that the payoff equals the cash flow expected on that date. But this method is difficult to implement, for a full set of zero-coupon bonds with suitable maturity dates may not be available. Therefore, a more promising approach is to identify the risk factors that determine the returns on the GHP, and to match these exposures. This is what we turn to in the next section.

Pricing

The retirement bond can be priced by the same no-arbitrage arguments as a regular coupon-paying bond. Its price is the sum of the future cash flows, discounted at the zero-coupon rates of appropriate maturities. Formally, let t denote the current date, T be the retirement date, τ be the number of annual cash flows in decumulation, and π be the annual growth rate. The time unit is one calendar year. The cost-of-living-adjusted cash flow that occurs at date $T + h$, where h is 0, 1, 2, ..., $\tau - 1$, is $[1 + \pi]^{T + h}$. It is discounted at the zero-coupon rate $y_{t, T + h - t}$, which is the prevailing rate at date t for cash flows occurring after $T + h - t$ years. The cash flows to discount are those that have not yet been paid at date t, so the sum is over all cash flows after the current date. By convention, if t is a payment date, the cash flow of this date is excluded, so the bond price drops by the cash flow amount just after date t. So, the no-arbitrage price of the retirement bond at date t is

$$B_t = \sum_{\substack{h=0 \\ h>t-T}}^{\tau-1} [1+\pi]^{T+h} \exp\left(-[T+h-t] y_{t,T+h-t}\right). \tag{3.1}$$

In this equation, we have assumed continuously compounded rates, hence the exponential in the right-hand side. Note that if date t is some point in the accumulation period, $t - T$ is non-positive, so the sum includes all cash flows from the retirement date to the end of the decumulation period. Price drops occur only in decumulation, at the beginning of every year.

A standard coupon-paying bond is also priced by a formula like Equation (3.1), but the cash flows have to be modified: before maturity, they are coupon payments, and at maturity they include both the last coupon and principal redemption.

3.2 Factor-Based Method for the Construction of a Goal-Hedging Portfolio

The roadmap for constructing a GHP by factor matching is as follows. First, specify a set of replicating instruments. Second, identify the set of risk factors that impact the returns on the GHP and those of the replicating securities. Third, measure the exposures of the benchmark and the assets to these factors. Finally, construct a factor-matching portfolio by matching the exposures to selected factors on the asset side (the replicating portfolio) and the liability side (the benchmark). The last operation is to be repeated at every rebalancing date.

Table 3: Risk and return statistics of constant-maturity bonds from January 1986 to January 2019

Maturity (years)	Annual return (%)	Annual volatility (%)	Sharpe ratio	Maximum drawdown (%)
1/12	3.26	0.72	-	0.00
1	4.37	1.14	0.97	0.24
2	4.81	2.02	0.77	1.64
3	5.45	3.07	0.71	3.19
4	6.05	4.12	0.67	4.76
5	6.59	5.14	0.65	6.53
6	7.07	6.12	0.62	8.66
7	7.50	7.08	0.60	10.67
8	7.88	8.01	0.58	12.55
9	8.23	8.93	0.56	14.32
10	8.54	9.85	0.54	16.00
11	8.82	10.75	0.52	17.60
12	9.08	11.64	0.50	19.13
13	9.32	12.51	0.48	20.61
14	9.54	13.37	0.47	22.04

15	9.75	14.20	0.46	23.43
16	9.96	15.01	0.45	24.79
17	10.15	15.79	0.44	26.12
18	10.34	16.55	0.43	27.41
19	10.53	17.29	0.42	28.69
20	10.71	18.02	0.41	29.94
21	10.89	18.74	0.41	31.16
22	11.07	19.46	0.40	32.37
23	11.25	20.20	0.40	33.55
24	11.42	20.96	0.39	35.18
25	11.60	21.76	0.38	37.62
26	11.78	22.61	0.38	40.03
27	11.96	23.53	0.37	42.39
28	12.13	24.54	0.36	44.70
29	12.31	25.65	0.35	46.94
30	12.49	26.87	0.34	49.12

Notes: Monthly returns on constant-maturity bonds are simulated over the period from 1 January 1986 to 1 January 2019, using the US zero coupon rates available from the Federal Reserve website and the secondary market rate on 3-month Treasury bills. The "annual return" is the average annual geometric return, the "annual volatility" is the volatility of log returns, and the Sharpe ratio is equal to the average annual return of the bond, minus that of the one-month bond, divided by the bond volatility.

3.2.1 The Replicating Assets: Constant-Maturity Bonds

To construct the replicating portfolios, we use "constant-maturity" zero-coupon bonds, defined as strategies that roll over zero-coupon bonds with the same maturity. In practice, such roll-over strategies would be implemented with coupon-paying bonds, but using pure discount bonds has the advantage that the returns can be calculated by using the zero-coupon rates series alone, without having to specify coupon rates. The roll-over is done by taking a long position in a pure discount bond of the target maturity, say u, and by selling it after time h, when it has maturity $u - h$, to take a long position in a new bond of maturity u.

We consider monthly roll-overs of bonds with constant maturities equal to one month and ranging from one to thirty years. The input dataset consists of time series of zero-coupon rates with constant maturities ranging from 1 to 30 years, to which we add the secondary market rate on 3-month Treasury bills in order to capture information on the short end of the yield curve. The zero-coupon rates are calculated by the method of Gürkaynak, Sack and Wright (2007), who estimate a Nelson-Siegel-Svensson model of the yield curve. The start date for each series depends on the maturity range of existing Treasury securities. Thus, since the maximum maturity of Treasury bonds was 20 years until 1985, maturities of longer than 20 years are reported only as of 25 November 1985. The dataset is periodically

87 - Source: Federal Reserve Bank of St. Louis, https://fred.stlouisfed.org/series/DGS10.

updated and made available on the Federal Reserve website at https://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html. The 3-month T-bill rate is directly obtained from the Federal Reserve economic database at https://fred.stlouisfed.org/series/TB3MS. Figure 5 shows the rates of maturities at 1, 5, 10 and 20 years. When a zero-coupon rate of a non-integer maturity in years is needed, it is obtained by linear interpolation of the existing rates, and when a rate shorter than three months is needed, it is taken as equal to the 3-month rate.

Table 3 shows statistics on the returns of constant-maturity bonds simulated from January 1986 to January 2019. The average annual return increases monotonically with the maturity, and reaches high values, greater than 10% for maturities of longer than 16 years. But volatilities are high too, at 14.20% for the 15-year bond and even 26.87% for the 30-year one. These long-term bonds also have a risk of substantial losses, with maximum drawdowns of 23.43% for the 15-year one and 49.12% for the 30-year one. The second of these losses was recorded from January 2009 to April 2010, a period in which the 10-year Treasury constant maturity rate rose from 2.46% to 3.89%.[87] Although the average return increases with maturity, volatility increases at a higher rate, so that Sharpe ratios are decreasing: the 1-year bond has a Sharpe ratio close to unity, at 0.97, and the 30-year one has a ratio of 0.34. This result is consistent with a finding by Deguest, Martellini and Milhau (2018b) based on actual Treasury coupon-paying bonds: they report that if no duration constraint is applied, the maximum Sharpe ratio portfolio of Treasuries has an average duration that is much shorter than an equally-weighted portfolio, so it is biased towards short bonds.[88]

For comparison, a US equity index of the 500 largest stocks weighted by capitalization posted a return of 10.24% per year, including reinvestment of dividends, with a volatility of 17.93% and a maximum drawdown of 54.31% over the same sample period.[89] Thus, the equity index has average return and volatility similar to those of the 20-year bond, with a larger maximum drawdown and a slightly lower Sharpe ratio, at 0.39 versus 0.41. We are not arguing that constant-maturity bonds were a better investment than stocks in the past, for they represent the returns on hypothetical portfolios, and the costs of stripping coupon-paying bonds into zero-coupon securities and rolling over the stripped instruments are ignored here. Moreover, these high returns may not occur again going forward because they were driven in part by the decrease in interest rates visible in Figure 5.

Performance Attribution for Constant-Maturity Bonds

Decreasing interest rates mechanically inflate bond returns, but these returns also depend on the level of rates, so it is interesting to try and disentangle the various sources of returns. To this end, Appendix E calculates a decomposition of the monthly logarithmic return on a constant-maturity bond in three terms:

- An "income" term, which is equal to the rate at the beginning of the month, times 1/12. This

88 - See their Exhibit 5.
89 - These figures were obtained by chaining the returns on the Scientific Beta Long-Term US cap-weighted index over the period from January 1986 to January 2016, and the returns on the Scientific Beta US cap-weighted index from January 2016 to January 2019. Data is from the Scientific Beta database.

3. Efficient Factor Exposure Matching: Improving the Hedging Benefits of the Liability-Hedging Portfolio with Factor Investing

Figure 4: Decomposition of average logarithmic returns on constant-maturity bonds for the period from January 1986 to January 2019

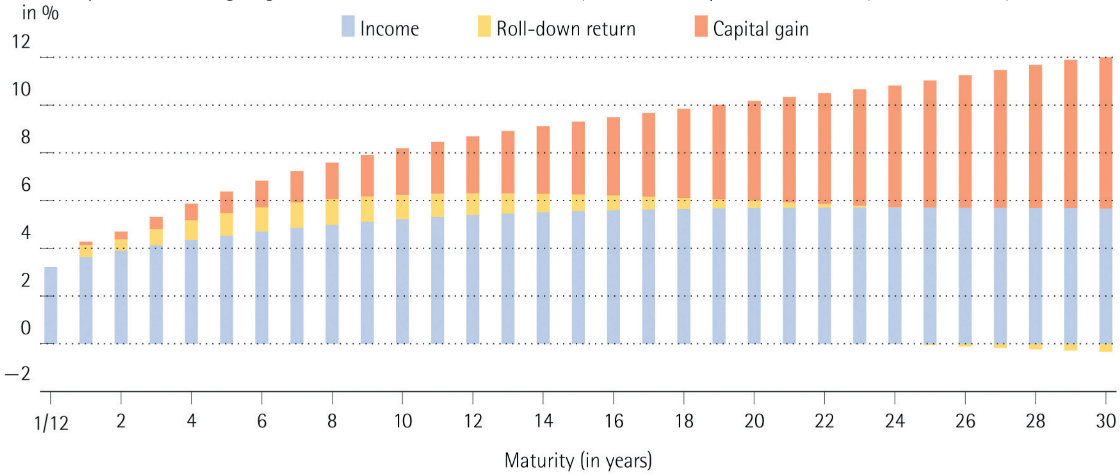

Notes: Average monthly logarithmic returns of constant-maturity bonds are calculated from January 1986 to January 2019, and they are decomposed into three terms: an income term, which reflects the passage of time and is always positive, a contribution of roll-down returns, which corresponds to a shift to the left on the yield curve, and a contribution of capital gains, which reflects changes in the position of the yield curve.

term reflects the passage of time and is always positive;

- A "roll-down return", which corresponds to the replacement of a bond with a bond of longer maturity on the roll over at the end of the month. This term is positive if the bond coming in is less expensive than the bond going out, i.e. if the yield curve is locally increasing near the target maturity at the end of the month. It is negative otherwise;
- A "capital gain" term, which corresponds to the impact of a change in the zero-coupon rate of the target maturity over the month. It contributes positively to the bond return if the interest rate decreases, and negatively otherwise.

The last two terms are zero when the target maturity equals the roll-over period, i.e. one month. For longer maturities, they can be either positive or negative.

Figure 4 shows the average of these three terms from 1986 to 2019. By construction, they add up to the logarithmic return on the bond over this period, and the logarithmic return is less than the geometric return reported in Table 3. The income term is the average zero-coupon rate over the period, so it has a positive contribution for all bonds. The capital gain term is the average change in the rate, multiplied by the maturity minus 1/12. For all but the 1-month bond, the contribution is positive too. The average roll-down return is the average spread between the rate of the target maturity and the rate of the target maturity minus 1/12. As a general rule, it is much lower than the sum of the other two terms. The main observation to take from this figure is that for very long-term bonds, of maturities longer than 20 years, the contribution of decreasing rates to the positive returns was nearly as large as, or even bigger than, the contribution of income.

For instance, the income effect for the 30-year bond was 5.66%, the average roll-down return was −0. 33% and the capital gain effect was 6.44%, leading to a logarithmic return of

$$5.66 - 0.33 + 6.44 = 11.77\%.$$

Suppose that the 30-year interest rate had followed the reverse path of what was observed between 1986 and 2019. The average rate would have been unchanged, but the capital gain term would have changed sign and the roll-down return would have been almost unchanged,[90] so the total return would have been

$$5.66 - 0.33 - 6.44 = -1.10\%.$$

Figure 5: US zero-coupon rates of selected maturities, from July 1961 to January 2019

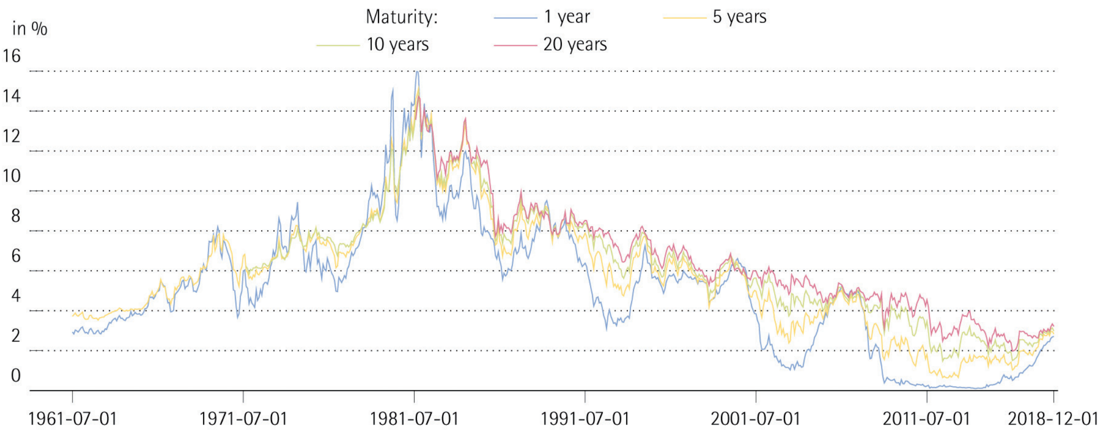

Notes: Data is from the Federal Reserve website, at https://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html.

Figure 6: Exposures of zero-coupon rates to the level, slope and curvature factors obtained by principal component analysis; data from January 1986 to January 2019

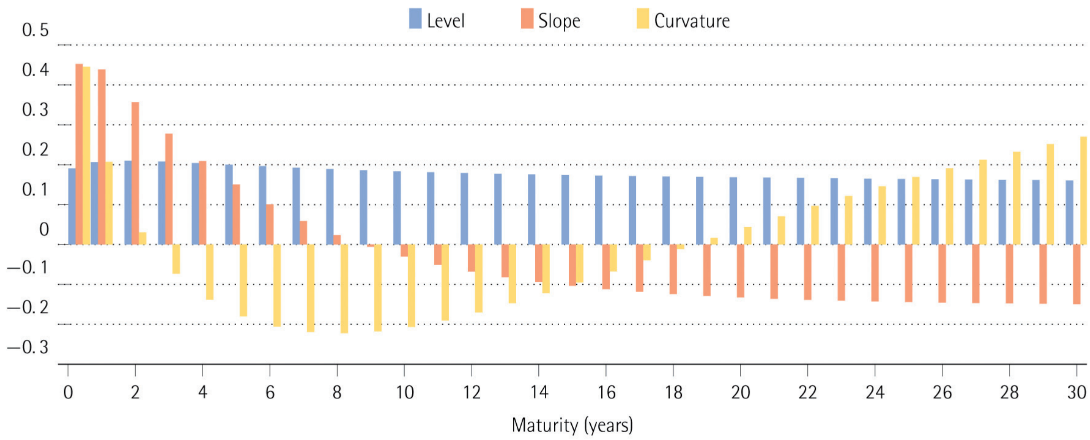

Notes: The level, slope and curvature factors are estimated by performing a principal component analysis of daily US zero-coupon rates over the period from 1 January 1986 to 1 January 2019. The exposures of the zero-coupon rates to the factors are the elements of the first three eigenvectors of the covariance matrix of rates (see Appendix C for details).

90 - Strictly speaking, the roll-down return is slightly modified because it is measured at the end of the month, so reversing the interest rate path amounts to measuring it at the beginning of the month.

The positive return turns into a negative value: this sharp decrease in the average return illustrates the strong impact that interest rate changes have on the performance of long-term bonds. Unless interest rates remain stable on average, the return on such bonds is very different from the average zero-coupon rate of the same maturity, which is the contribution of the income effect.

3.2.2 Risk Factors

Equation (3.1) shows that the risk factors that determine the returns on a retirement bond are those that determine changes in discount rates. To identify the risk factors in bond returns, it therefore suffices to identify the risk factors in nominal interest rates. This approach seems less straightforward than extracting factors directly from bond returns, but it is more appropriate because bond returns are not stationary over time. Indeed, they converge to zero as the security approaches maturity, and their volatility tends to decrease over time as a result of the decreasing maturity. This nonstationarity invalidates the use of many standard statistical techniques: for instance, the input of principal component analysis, which is the covariance matrix, cannot be obtained by taking the sample covariances between bonds.

In this section, we present three standard methods to identify the risk factors in nominal rates, namely principal component analysis, the Nelson-Siegel model and the calculation of explicit factors.

Implicit Risk Factors from Principal Component Analysis

It is well known that within the universe of bonds issued by a single sovereign state, three factors explain most interest rate movements: the level, the slope and the convexity of the yield curve. To illustrate this standard decomposition, we use the same dataset as in the previous section, which comprises the zero-coupon rates of maturities ranging from 1 to 30 years and the rate on 3-month Treasury bills. The strong correlation between the rates of different maturities is evident in Figure 5, pointing to the existence of a level factor.

Figure 6 illustrates the decomposition of the yield curve into three factors by plotting the exposures of the 31 rates to the first three factors from a principal component analysis (PCA). Technically, these betas are the components of the eigenvectors of the covariance matrix of the 31 series of rates, following the formulas written in Appendix C. Their sign and magnitude have no special meaning, as they are the results of normalization conventions in the statistical procedure, but the interesting pattern is the variation of betas across maturities, which conveys the usual interpretation for the three factors. All rates have roughly the same exposure to the first factor, so a change in this factor has about the same impact on all of them: this factor thus represents a level factor. A positive shock on the second factor has a positive impact on rates of maturities of up to 8 years, and a negative impact on those of maturities of greater than 10 years, so it results in a flattening of the curve.

Thus, the second factor can be regarded as a flattening factor around the 10-year maturity. Finally, a positive shock on the third factor lowers the yields of maturities ranging from 3 to 19 years, and it raises the yields outside this range. As a result, a convex yield curve becomes more convex, and a concave curve becomes less concave. In all cases, an increase in this factor implies a more convex curve, so the factor can be regarded as a convexity factor.

It should be noted that while factor values at each point in time are linear combinations of the zero-coupon rates at this date, the coefficients of this combination are estimated from a sample of rates, so they depend on past and future rates. In the example of Figure 6, the sample covers the period from 1 January 1986 to 1 January 2019, so the factor values estimated at any point in this range depend on the entire sample of interest rates up to 1 January 2019. As a result, factor values involve look-ahead bias, except at the last date in the sample, which is 1 January 2019 here. So care must be taken in portfolio construction exercises to avoid introducing this bias in the process.

Factors in the Nelson-Siegel Model

The Nelson-Siegel (NS) model was introduced by Nelson and Siegel (1987) as a model that can capture the three main observed forms of the yield curve, namely humps, S shapes and monotonic curves, with a parsimonious parametrization. The continuously compounded zero-coupon rate of maturity h at date t is modeled as a function of maturity as follows:

$$y_{t,h} = \delta_{0,t} + \delta_{1,t} f\left(\frac{h}{m_{1,t}}\right) + \delta_{2,t}\left[f\left(\frac{h}{m_{1,t}}\right) - \exp\left(-\frac{h}{m_{1,t}}\right)\right], \quad (3.2)$$

$\delta_{0,t}$, $\delta_{1,t}$, $\delta_{2,t}$ and $m_{1,t}$ are four parameters that depend on time but not on maturity, and the function f is given by

$$f(x) = \frac{1 - \exp(-x)}{x}.$$

Each term in the right-hand side of Equation (3.2) has a geometric interpretation. Any change in the parameter $\delta_{0,t}$ results is fully and equally reflected in the yields of all maturities, so it translates the yield curve upwards or downwards. Thus, $\delta_{0,t}$ can be regarded as a level factor. By sending h to infinity in Equation (3.2), we can also see that it is the limit of the zero-coupon rate as the maturity rises towards infinity. The impact of a change in the next parameter, $\delta_{1,t}$, depends on the maturity. Figure 7 shows that it is greater for short than for long maturities, so an increase in $\delta_{1,t}$ leads to a flattening of the curve. Thus, $\delta_{1,t}$ represents a slope factor. This interpretation also becomes clear when one considers that the limit of the yield as the maturity shrinks to zero is $\delta_{0,t} + \delta_{1,t}$, which is to be compared with the limit for an infinite maturity, equal to $\delta_{0,t}$. Thus, $\delta_{1,t}$ is the negative of the term spread, defined here as the excess of the very long-term yield over the instantaneous one. Finally, an increase in $\delta_{2,t}$ produces a more concave curve, since the middle of the yield curve is raised but the short and long ends are hardly impacted, so $\delta_{2,t}$ represents a concavity factor.

Figure 7: Exposures of zero-coupon rates to the factors in the Nelson-Siegel model

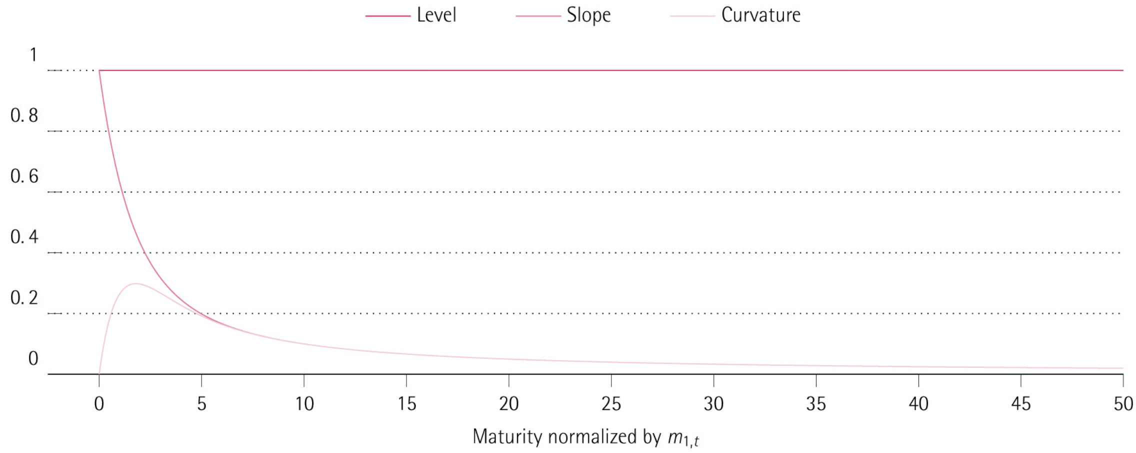

Notes: Factor exposures in the Nelson-Siegel model are measured as the coefficients that multiply the coefficients $\delta_{0,t}$, $\delta_{1,t}$ and $\delta_{2,t}$ in Equation (3.2). The three coefficients are respectively associated with the level, slope and curvature factors.

An extension of the model based on Svensson (1994) adds flexibility to the model by introducing a second curvature term in the right-hand side of Equation (3.2). This term contains a second curvature factor denoted by $\delta_{3,t}$ and a reference maturity $m_{2,t}$. This modification defines the Nelson-Siegel-Svensson (NSS) model, which involves a total of six parameters.

The parameters of the NS and NSS models are estimated by fitting the model to the prices of sovereign bonds. First, a basket of securities is chosen, the model price of each is written as the sum of discounted future cash flows. Parameters are then estimated by minimizing some aggregate measure of the distance, e.g. the sum of squared differences, between model prices and the observed market quotes. Note that since the sum of discounted cash flows is the no-arbitrage price of the security, the market price to match approximately is the price at which the bond trades, namely its dirty price. This procedure is implemented by Gürkaynak, Sack and Wright (2007) for US Treasury notes and bonds and results in time series of estimates for the six parameters of the NSS model, and in time series of estimates of zero-coupon yields of constant maturities ranging from 1 to 30 years.

In principle, the level, slope and concavity factor of the NS or NSS model do not involve the look-ahead bias that is present in principal factors, since they are estimated at each date by fitting the model to the yield curve of the date. Of course, this property no longer holds if estimated factors are revised ex post, e.g. in order to smooth out outliers. In fact, there is strong suspicion of outliers in the two concavity factors, as is evident in Figure 8. In spite of these extreme values, the zero-coupon rates implied by the model do not post any abnormal values, as can be checked in Figure 5.

3. Efficient Factor Exposure Matching: Improving the Hedging Benefits of the Liability-Hedging Portfolio with Factor Investing

Explicit Risk Factors

In fixed income analysis, the slope of the yield curve is commonly measured as the term spread, defined as the difference between a "long" and a "short" rate. Similarly, the "level" of rates is naturally estimated as the cross-sectional mean of rates at a given point in time. Thus, one obvious question is the extent to which these measures are consistent with the factors from PCA or from the NS model. For concavity, there does not seem to be a standard proxy, perhaps because this factor is less used than the other two in fixed-income portfolio analysis. To examine the relationship between explicit factors and the factors introduced above, we define a "level" factor as the mean of the 3-month and 10-year rates, a slope factor as the spread between these two rates, and a concavity factor as the mean of the 1-year and 20-year rates, minus the mean of the 5-year and 10-year rates.

These definitions are somewhat arbitrary, in particular regarding the choice of maturities, and their heuristic justification is that they aim to mimic the pattern of coefficients plotted in Figure 6, without using these coefficients. It is of course impossible to exactly replicate the PCA factors unless the very coefficients of Figure 6 are applied to each maturity, but our goal is to define explicit factors that proxy for PCA factors while relying as little as possible on the results of PCA, so as to avoid inheriting the look-ahead bias involved in PCA factors.

Figure 8: Factors of the Nelson-Siegel-Svensson model from January 1986 to January 2019

(a) Level

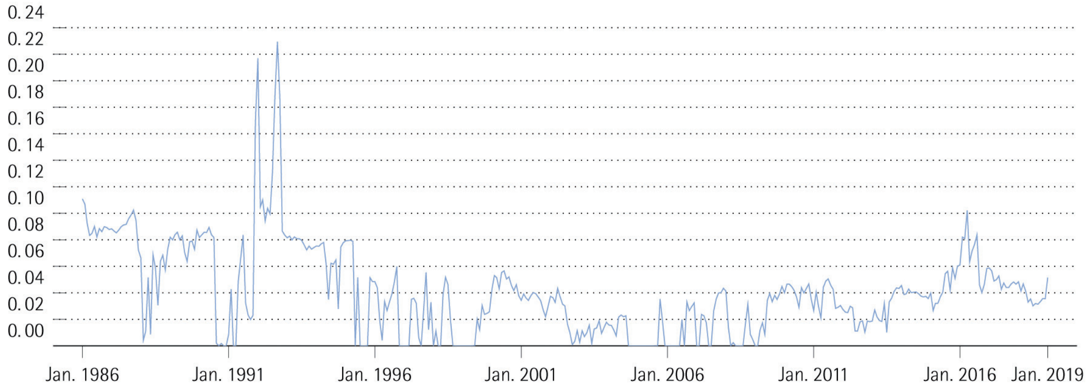

b) Slope

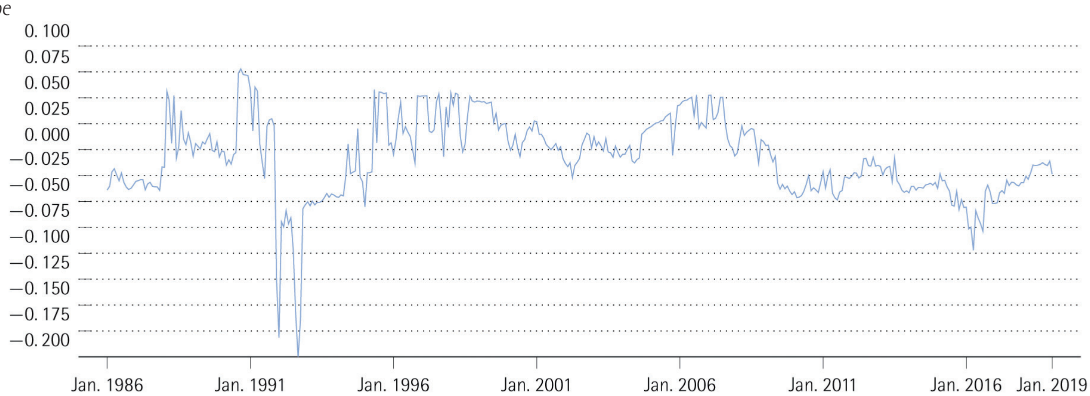

(c) Concavity

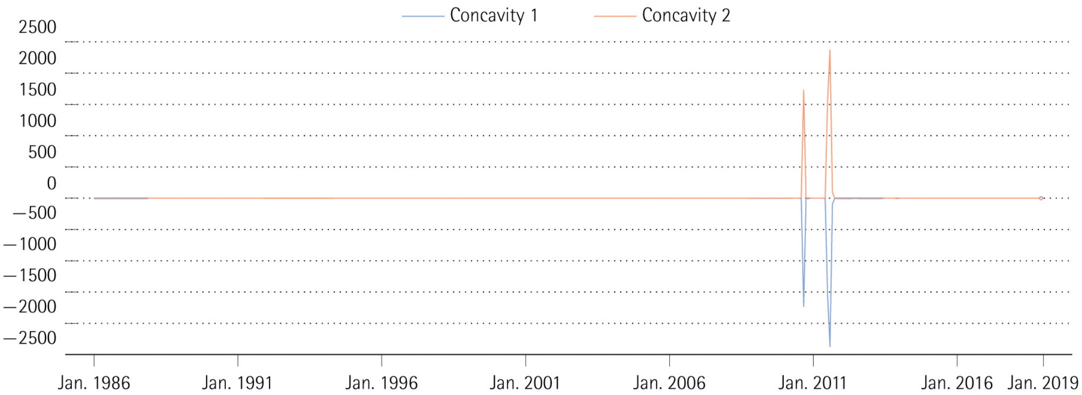

Notes: Factors of the Nelson-Siegel-Svensson model are obtained from the Federal Reserve website.

Thus, the explicit level factor is an equally-weighted average of two rates. One could also take an average of all rates, but in the sample from 1986 to 2019, the correlation between the two metrics is a very high 97.3%, so we favor the less data-intensive definition out of parsimony. The explicit slope factor is a term spread, which is a commonly used indicator for bond return prediction. Clearly, taking the difference between two rates eliminates most of the influence of the level factor, since both rates have high exposures to it. For the concavity factor, the pattern in Figure 6 points to a linear combination of rates with positive coefficients at the short and long ends, and negative loadings in the medium term. Our first attempt was therefore the average of the 1-year and 20-year rates, minus the 10-year rate, but this indicator turned out to have a correlation

Table 4: Correlations between implicit and explicit interest rate risk factors (in %); data from January 1986 to January 2019

	PCA-Lev.	PCA-Slo.	PCA-Con.	NSS-Lev.	NSS-Slo.	NSS-Con.1	NSS-Con.2	EXP-Lev.	EXP-Slo.	EXP-Con.
PCA-Lev.	100.0									
PCA-Slo.	0.0	100.0								
PCA-Con.	0.0	0.0	100.0							
NSS-Lev.	33.5	−23.2	5.1	100.0						
NSS-Slo.	25.0	46.8	2.1	−79.1	100.0					
NSS-Con. 1	5.8	7.8	−8.4	−1.1	5.5	100.0				
NSS-Con. 2	−5.8	−7.8	8.4	1.0	−5.4	−100.0	100.0			
EXP-Lev.	97.6	21.0	2.9	27.0	35.5	6.9	−6.9	100.0		
EXP-Slo.	−7.3	−93.7	−30.8	22.8	−52.0	−5.0	5.0	−28.4	100.0	
EXP-Con.	−17.5	69.4	57.8	−36.3	46.0	−3.0	3.0	−0.2	−83.9	100.0

Notes: Factors beginning with "PCA" are extracted from zero-coupon rates by principal component analysis, and those beginning with "NSS" are obtained by fitting a Nelson-Siegel-Svensson model to the data. NSS factors are obtained from the Federal Reserve website. Finally, factors beginning with "EXP" are explicitly defined factors. EXP - Level is the average of the 3-month rate and the 10-year rate; EXP - Slope is the spread between these two rates; and EXP - Concavity is the mean of the 1-year and 20-year rates, minus the mean of the 5-year and 10-year rates. Data is daily and covers the period from 1 January 1986 to 1 January 2019.

of −96. 7% with the term spread in the sample, and so it provided little additional informational content. We found that replacing the 10-year rate with the average of the 5-year and 10-year rates led to a correlation of −83. 9%, which is still high but is lower than before, suggesting that there is less overlap with the slope factor.

Table 4 summarizes the correlations between the various factors, prefixing the factor names with the name of the approach behind their definition, namely PCA, NSS or EXP – which stands for "explicit". PCA factors are uncorrelated from each other by definition. Explicit factors are more correlated, which confirms that they are imperfect proxies for PCA factors. In particular, an −83. 9% correlation remains between concavity and the term spread. Correlations between the first two PCA factors and their explicit versions are high, at 97.6% for the level and −93. 7% for the slope. The latter is negative because a positive shock on PCA-Slope produces a flatter yield curve, while an increase in EXP-Slope leads to a steeper curve.

The two concavity factors of the NSS model have a strongly negative pairwise correlation, which rounds to −100% in the table. This value, however, should be taken with caution in view of the outliers visible in Panel (c) of Figure 8. To assess the explanatory power of each set of factors, we first regress the 31 series of zero-coupon rates on the level factor of each set, then on the level and slope factors, and finally on the level, slope and concavity factors. For the NSS model, the last stage subdivides into two steps, introducing the two concavity factors sequentially. Figure 9 plots the R-squares from these regressions. Overall, PCA factors have the largest explanatory power: the level factor alone explains at least 81.7% of the variance, and the three factors account for more than 99%. Explicit factors also have good explanatory power, though the three of them jointly explain less than 99% of the variance of the longest rates. With both systems of factors, the marginal gain in R-square from introducing the concavity factor beyond the slope is very modest and never exceeds one percentage point. This result suggests that concavity has limited usefulness as a risk factor when level and slope are already present.

The NSS factors have less explanatory power than the other two sets. A striking feature is that the level factor explains at most 17% of the variance, a maximum reached for the 30-year rate, and also that its explanatory power is even lower for shorter maturities. It is in fact not surprising to find that the R-square of the single-factor regression increases with maturity since the level factor in the NSS model is equal to a hypothetical yield of infinite maturity, and thus by definition is more closely linked to long-term than short-term rates. Introducing the slope factor makes a strong improvement, but the first concavity factor provides almost no marginal benefit. It is the second concavity factor that provides the next substantial improvement, especially for long-term rates. It may seem surprising that this factor has such a marginal contribution with respect to the first concavity factor, given the high correlation between the two – close to −100%, as can be seen in Table 4. For instance,

Figure 9: Percentage of variance of zero-coupon rates explained by factors from January 1986 to January 2019

(a) PCA factors.

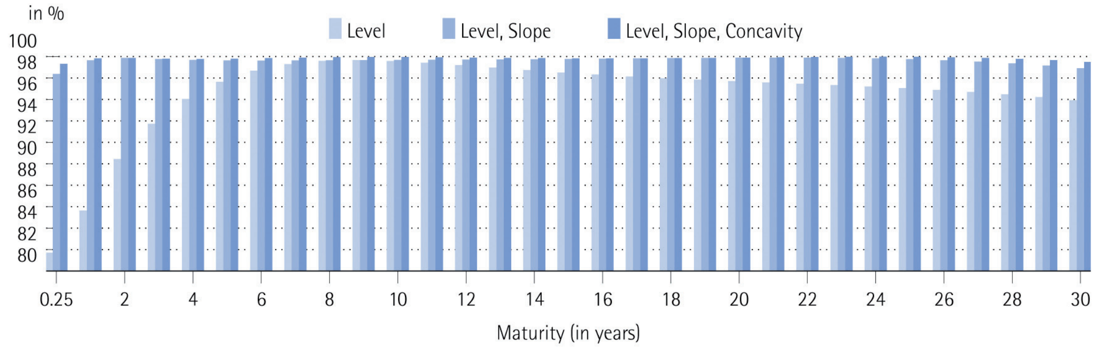

(b) NSS factors.

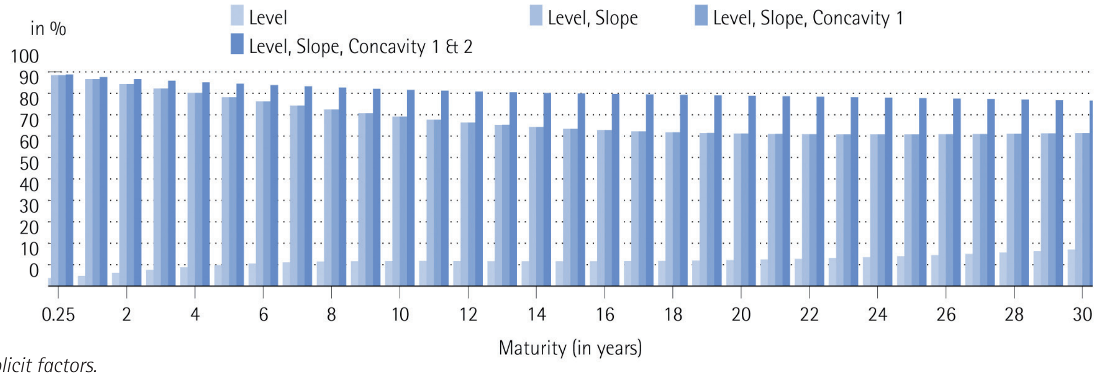

(c) Explicit factors.

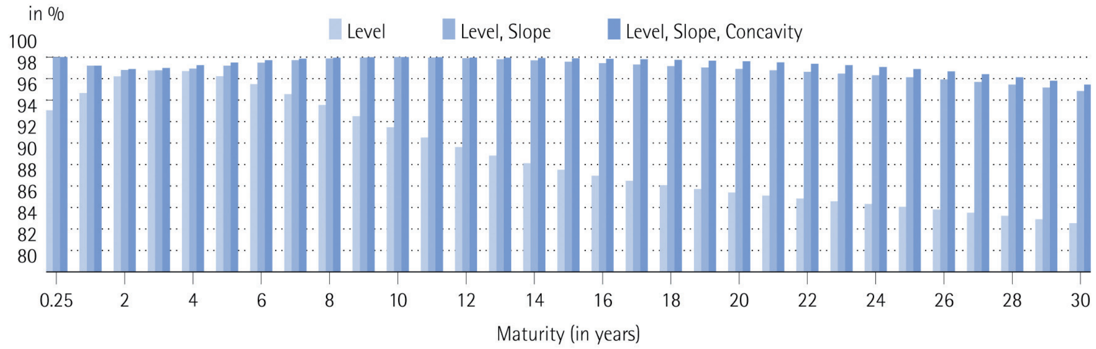

Notes: Zero-coupon rates are regressed on factors obtained by principal component analysis (Panel (a)), on factors from the Nelson-Siegel-Svensson model (Panel (b)), or on explicit factors (Panel (c)). Data is daily and covers the period from 1 January 1986 to 1 January 2019.

the R-square of the regression of the 30-year yield is 76.5% with three factors, and rises to 86.6% with four factors. The reason for this is that while the second concavity factor is almost spanned by the first three factors, i.e. while the residuals of this factor with respect to the other three are very small, they have a non-negligible correlation with the 30-year yield: this correlation is 38.9%, and the gain in R-square equals 15.1%, which is the square of 38.9%. This property is general, and is derived in Appendix D: the gain in R-square from introducing a new explanatory variable in the right-hand side of a regression is equal to the squared correlation between the dependent variable and the residuals of the new variable with respect to the regressors already present.[91]

Conversely, in Panel (a), the PCA concavity factor is uncorrelated from the level and slope factors, but its marginal contribution is small. For the 30-year rate, the gain in R-square is only 0.6%. Repeating the calculation done for the NSS factors enables us to understand why. When regressed on level and slope, the concavity factor produces an R-square of zero because it is uncorrelated with both regressors, so the residuals are equal to the factor itself. The correlation between the concavity factor and the 30-year rate is only 7.6% in the sample, so the squared correlation is 0.6%. One lesson to take away from Panels (a) and (b) is therefore that introducing a new risk factor that is highly correlated with the existing ones may substantially improve the explanatory power of the model, and that conversely, a new factor uncorrelated from the others may be of little help in explaining a variable of interest – an interest rate or a return. What matters is how correlated the residuals of the new factor are with the variable with respect to the existing ones. They need not be large to be highly correlated.

3.2.3 Measuring the Factor Exposures of Assets and Benchmark

The standard approach of regressing bond returns on factors to measure the exposures of bonds is not operational for finite-maturity bonds because of the non-stationarity induced by the finite maturity. To estimate exposures, we thus employ a different technique, by exploiting the functional relationship between bond prices and zero-coupon rates.

General Equations

All factor models described in Appendix 3.2.2 imply that zero-coupon rates are linear functions of one, two or three risk factors. Mathematically, the rate of maturity u at date t is given by

$$y_{t,u} = c(u) + \sum_{k=1}^{K} b_k(u) X_{t,k} + \eta_{t,u}. \qquad (3.3)$$

But the retirement bond price is a non-linear function of discount rates, given by Equation (3.1),[92] and its returns are also non-linear in the factors. Similarly, the arithmetic returns on a constant-maturity bond are non-linear in discount rates, although log returns are linear, and it is the former sort of returns that matters for portfolio construction, because the arithmetic return on the portfolio is the weighted sum of the arithmetic returns on the constituents.

Because of the non-linear relationship between arithmetic bond returns and zero-coupon rates, calculating the sensitivities of returns to changes in the factors requires some algebra and a few simplifying assumptions. Appendix F derives formulas for the retirement bond, which is the

91 - This statistical property has interesting implications in portfolio optimization. See Deguest, Martellini and Milhau (2018a).
92 - This non-linearity would not be eliminated by considering discretely compounded rates as opposed to continuously compounded rates.

benchmark in the replication exercise, and for the constant-maturity bonds, under the following assumptions:

A1. Zero-coupon rates follow an exact factor model, which means that the residuals in Equation (3.3) are zero;

A2. For the constant-maturity bonds, the roll-down return is negligible compared to the income term and the capital gain term. The empirical decomposition of average bond returns in Figure 4 provides some support for this assumption;

A3. Also for constant-maturity bonds, returns are measured between two consecutive roll-over dates, hence over one month;

A4. And again for constant-maturity bonds, monthly arithmetic returns are well approximated as monthly log returns. This assumption relies on a first-order approximation of the logarithm near one: for a small number x, $\log[1 + x]$ is close to x.

Then, it is shown in Appendix F that the sensitivity of the bond price to a change in the k^{th} factor is

$$\gamma^B_{t,k} = -\frac{1}{B_t} \sum_{\substack{i=0 \\ i>t-T}}^{\tau-1} [1+\pi]^{T+i} \exp\left(-[T+i-t] y_{t,T+i-t}\right) [T+i-t] b_k(T+i-t). \quad (3.4)$$

Now consider a portfolio invested in N constant-maturity bonds, with maturities $u_1, \ldots, u_N$ and weights $w_1, \ldots, w_N$.

Appendix F shows that the exposure of its return between two roll-over dates to a change in the k^{th} factor is

$$\gamma^p_{t,k} = -\sum_{i=1}^{N} w_i [u_i - h] b_k(u_i). \quad (3.5)$$

Here, h denotes the length of the roll-over period expressed as a fraction of years, and calculated as the number of days in the month divided by 365.25.

The notations $\gamma^B_{t,k}$ and $\gamma^p_{t,k}$ highlight the twofold dependency of exposures with respect to time and the factor. To be specific about the definition of an "exposure" here, these quantities represent the sensitivity of the return with respect to a change in the k^{th} factor: for a change in the factor equal to unity, the bond return changes by $\gamma^B_{t,k}$ and the portfolio return by $\gamma^p_{t,k}$. To neutralize the effect of a factor change on the excess return of the portfolio with respect to the retirement bond, the sensitivities of the portfolio and the bond must be equal.

Exposures to Principal Factors, Nelson-Siegel-Svensson Factors or Explicit Factors

The general equations can be applied to calculate bond exposures to the factors of a given model, by replacing the coefficients b_k with their model-implied expressions. The first exposure to estimate is the exposure to the level factor, since it is the factor that explains the biggest fraction of interest rate movements.

A "pure" level factor would be a factor such that all zero-coupon rates have the same exposure to it, regardless of their maturity. The NSS level factor falls into this category, as shown in Figure 7, but the PCA level factor only imperfectly matches this definition, since Figure 6 shows that exposures of rates to this factor are not strictly equal across maturities. For a pure level factor, the general expressions for the sensitivities in Equations (3.4) and (3.5) become slightly simplified. For instance, the exposure of the retirement bond can be rewritten as

$$\gamma^{p}_{t,k} = -D^{B}_{t,k} b_k,$$

where $D^{B}_{t,k}$ is the sum of cash flow maturities weighted by their respective sizes and discount factors and b_k is the common exposure of all rates. In the NSS model, b_k is equal to 1 for the level factor, so the bond exposure is the negative of $D^{B}_{t,k}$. Similarly, the level exposure of a portfolio of constant-maturity bonds is minus the weighted sum of the $[u_i - h]$, where the u_i-s are the maturities.

The definition of $D^{B}_{t,k}$ is very close to that of Macaulay duration; the difference lies in the way cash flows are discounted. In Macaulay duration, they are discounted at the yield to maturity of the bond,[93] and in $D^{B}_{t,k}$, they are discounted at the zero-coupon rates of the appropriate maturities. By definition, the yield to maturity is such that the sum of cash flows discounted at the zero-coupon rates equals the sum of cash flows discounted at the yield to maturity, but a difference can arise between Macaulay duration and $D^{B}_{t,k}$. However, we shall see in an example below that this difference is very small. To avoid confusion with the usual notions of duration – dollar, Macaulay and modified durations –, we call $D^{B}_{t,k}$ the "factor duration" of the retirement bond. Note that since continuously compounded rates are used in this paper, Macaulay duration and modified duration are equal.

As a numerical illustration, we calculate the exposures of the retirement bond for an investor who retired on 3 January 2000 and sought 20 replacement income cash flows beginning in January 2000. Thus, the cash flows are paid at the beginning of every calendar year from 2000 to 2019 inclusive. Figure 10 shows the exposures of the bond to the level, slope and concavity factors from the principal component analysis and from the Nelson-Siegel-Svensson model.

For the NSS model, exposures can be calculated at each date based on the dataset obtained from the Federal Reserve website, which provides the values for the parameters $m_{1,t}$ of Equation (3.2) and the parameter $m_{2,t}$ of the NSS model.

For the PCA factors, the exposures are estimated from the parameters displayed in Figure 6. These parameters are estimated over the entire sample period, from January 1986 to January 2019, so they involve look-ahead bias, except at the final date.

For comparison, Panel (a) of Figure 10 also shows the negative of the Macaulay duration of the bond, and it is visually almost indistinguishable from the exposure to NSS-Level. This observation confirms that the factor duration is very close to the Macaulay duration of the bond. Both durations exhibit discontinuities on cash flow dates, i.e. at the start of every year during the decumulation

93 - See Equation (5.6) p. 170 of Martellini, Priaulet and Priaulet (2003).

Figure 10: Exposures of retirement bond to the level, slope and concavity factors from principal component analysis and Nelson-Siegel-Svensson model from January 1986 to January 2019

(a) Level.

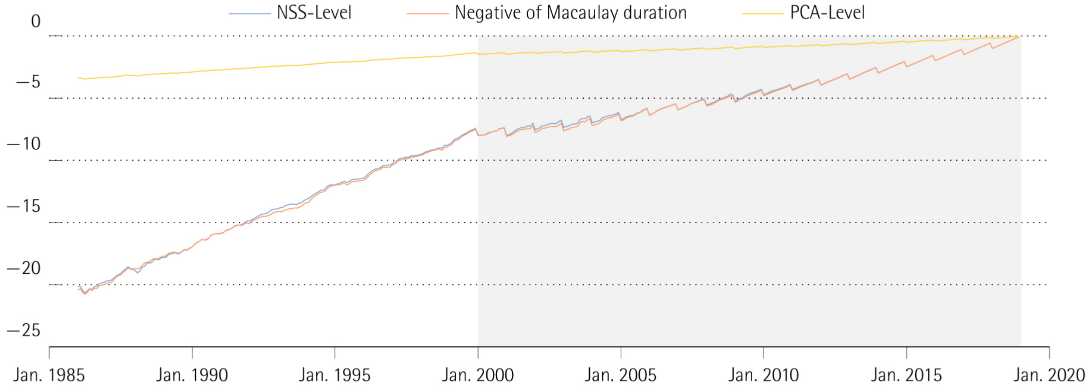

(b) Slope.

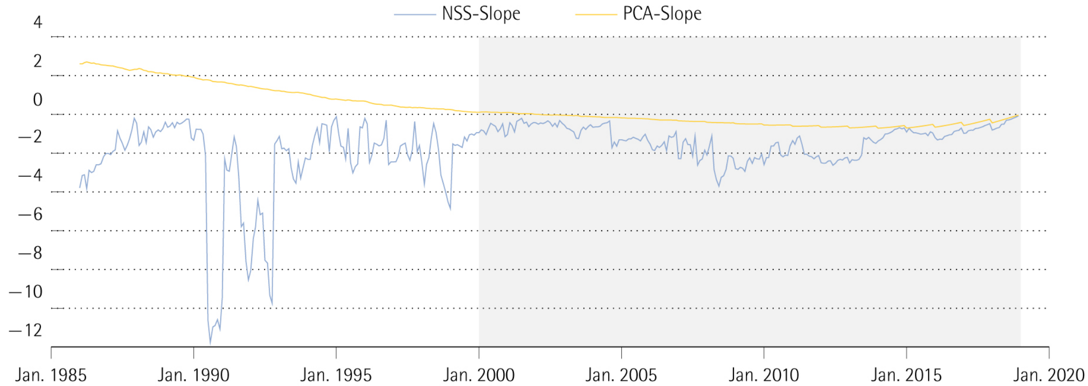

(c) Concavity.

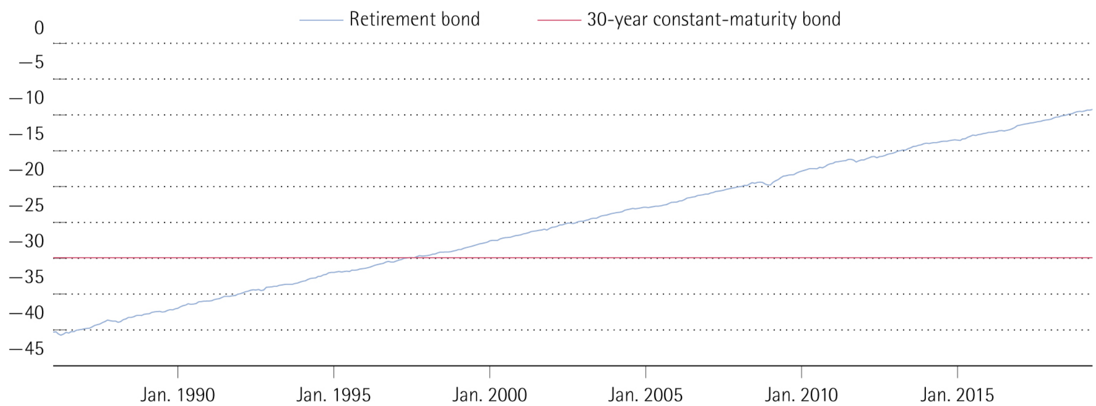

Notes: The retirement bond is for an individual who retired on 3 January 2000 and sought replacement income for 20 years, from 2000 to 2019 inclusive. Exposures to the Nelson-Siegel-Svensson factors are estimated from the parameters published on the Federal Reserve website and estimated by the method of Gürkaynak, Sack and Wright (2007). Exposures to the principal factors are estimated from the parameter values displayed in Figure 6, which are estimated over the entire sample period. Thus, these exposures involve look-ahead bias except at the last sample date. The shaded area is the decumulation period, while the white area on the left is the accumulation phase.

period. These jumps occur because the bond price itself jumps on these dates, while the sum of cash flow maturities weighted by the discounted cash flows is continuous: as a result, the ratio of the two quantities is discontinuous.

Exposure to PCA-Level contains the term structure of the exposures of zero-coupon rates to this factor, a term structure depicted in Figure 6, so it is not strictly proportional to the factor duration of the bond. Nevertheless, it increases almost monotonically during the sample period, ending at zero in January 2019. More generally, all factor exposures converge to zero as the bond approaches the end of the decumulation period.

3.2.4 Constructing Factor-Matching Portfolios

A factor-matching strategy equates the exposures of the replicating portfolio with those of the benchmark for selected factors. Mathematically, the percentage weights w_1, ..., w_N allocated to the constituents are calculated so as to equate the portfolio exposure in Equation (3.5) to the benchmark exposure, given by Equation (3.4). The factor-matching condition is written for every selected factor, which gives K equations if there are K exposures to match. In addition, a budget constraint applies, which states that constituents' weights should add up to 100%.

With $K + 1$ constraints, portfolio weights are completely determined if there are exactly $K + 1$ constituents. With fewer assets, the system of constraints is over-identified, so there is no portfolio that satisfies all equations simultaneously. With more than $K + 1$ constituents, the system is under-identified, meaning that there are an infinite number of combinations of constituents that satisfy the constraints. For simplicity, we consider the case of $K + 1$ constituents. When a single factor-matching constraint is imposed, for the purpose of constructing a level-matching portfolio, this amounts to taking two constituents. When two constraints are imposed, e.g. to construct a level- and slope-matching portfolio, the number of required constituents is three.

Picking Constituents From the Universe

The base universe contains 30 constant-maturity bonds, plus a cash account. Cash is simulated as the value of a money market account in which funds are reinvested every day at some short-term interest rate. The daily interest rate is taken to be the yield on 3-month Treasury bills, obtained from the Federal Reserve economic database. Cash is assumed to have zero exposure to all interest rate risk factors.

Because the universe contains 31 constituents and we need only two to construct a level-matching portfolio, we need some criterion to make a choice. The factor-matching condition suggests one approach. Indeed, the level exposure of the portfolio is the weighted average of the constituents' exposures, and it must be equal to the exposure of the retirement bond. So if the portfolio is long-only, with weights lying between 0 and 100%, the target exposure is comprised between the smallest and largest of the exposures of the assets included in the portfolio. Thus, a necessary condition for the portfolio to have only non-negative weights is to select one constituent with lower exposure than the target, and one with larger exposure.

Since level exposures are negative (see Figure 10) and cash has zero exposure, there is always an asset with greater exposure than the target. But there may be no asset with lower exposure than the target, for the target exposure far from retirement may be lower than the minimum of exposures across all constant-maturity bonds. As can be seen from Equation (3.5), the exposure of a constant-maturity bond at the beginning of a month is $-[u_i - h]$, where u_i is the maturity and h is the length of the month expressed as a fraction of a year. Since the maximum maturity of the bonds is 30 years, the bond with the lowest (i.e., the most negative) level exposure is the 30-year one.

Figure 11 shows the level exposures of the retirement bond for an individual who retires in January 2020 and of the 30-year constant-maturity bond over the period from January 1986 to January 2019, which is the period where interest rate data is available. Until June 1997, the retirement bond has lower exposure than the 30-year bond, so no level-matching portfolio with non-negative weights exists. It is only after this point that there is at least one asset with greater exposure (the cash account) and one with lower exposure. A similar situation will be encountered for any retirement date: the constraint of having a long-only replicating portfolio puts a lower limit on the date at which replication can start.

This restriction has an important financial implication: a level-matching portfolio cannot be constructed for any arbitrarily long time to retirement. In the example of Figure 11, the individual plans to retire in January 2020 and the replication can start in June 1997, 23 years before retirement. Assuming a retirement age of 65, the condition on the start date means that a replicating portfolio can be offered beginning at the age of 42.

Figure 11: Exposures to NSS level factor of retirement bond and constant-maturity bonds from January 1986 to January 2019

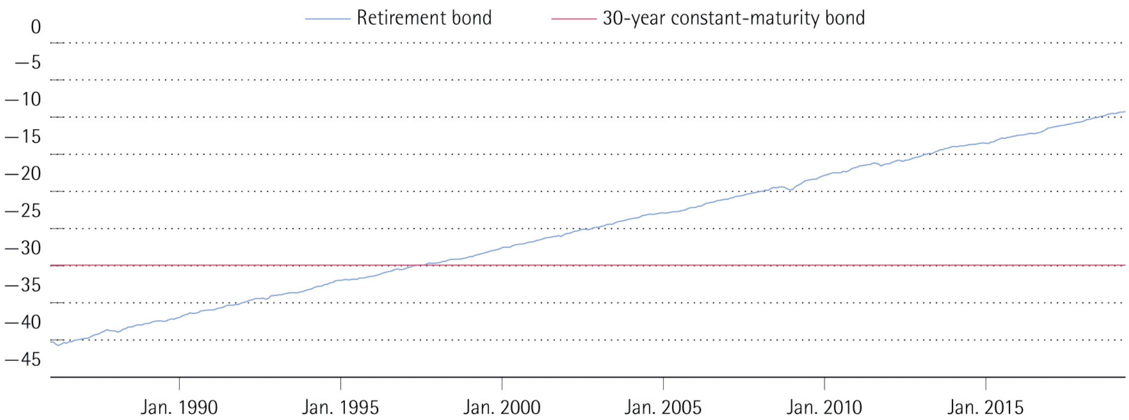

Notes: The retirement bond is for an individual who retired on 3 January 2020 and sought constant replacement income for 20 years, from 2000 to 2019. Exposures to the NSS level factor are calculated with the expressions given in Appendix F. The two lines cross in June 1997, approximately at −29. 9.

3. Efficient Factor Exposure Matching: Improving the Hedging Benefits of the Liability-Hedging Portfolio with Factor Investing

A sufficient (but not necessary) condition for the level exposure of the retirement bond to be greater than the minimum level exposure of potential constituents is to consider a replication period such that the maturity of the last replacement income cash flow never exceeds the longest maturity of the bonds, which is 30 years. Indeed, examination of Equation (3.4) shows that the level exposure of the retirement bond is always greater than minus the maturity of the last cash flow, multiplied by b_k. So if the last cash flow occurs in less than 30 years, the exposure is greater than that of the bond with the longest maturity, and a long-only level-matching portfolio exists. In the subsequent numerical illustrations, we thus let the accumulation phase start on the first date as of which the maturity of the last cash flow is less than 30 years. Since the last cash flow occurs 19 years after the retirement date, this amounts to letting accumulation start 11 years before retirement.

Within the period that begins at the starting date, we select two assets at each rebalancing date. The first is the one with the greatest exposure less than or equal to the benchmark, and the second is the one with the lowest exposure greater than or equal to the benchmark. This selection procedure ensures that the portfolio always contains one bond with greater exposure than the target.

Weighting Constituents

Once constituents are selected from the universe, the factor-matching portfolio is obtained by solving the equation for the weights that equate the factor exposures of the portfolio and the retirement bond and satisfy the budget constraint. When level is the only factor for which exposures are matched, there are two weights to calculate from two equations.

As mentioned previously, the level exposure of a bond is not proportional to its modified duration, so a level-matching portfolio is not a duration-matching portfolio. But because duration matching is a standard form of replication in the fixed-income class, we also test this replication method in addition to level-matching. For this alternative weighting scheme, the selection of the two constituents is based on modified duration, rather than on level exposure: at each rebalancing date, the two assets with the closest durations to the benchmark are chosen, one of which has a shorter duration and the other a longer duration.

3.3 Efficiency of Factor-Based Liability Replication Techniques in Accumulation and Decumulation

3.3.1 Accumulation

Results In Multiple Samples

Table 5 summarizes the results of simulations conducted for various retirement dates. Level-matching portfolios are referred to as “GHP Lev.” and duration-matching portfolios as “GHP Dur.”. In addition to these two replicating strategies, we also present the results of two alternative strategies that do not attempt to match any of the factor exposures of the retirement bond. One is fully invested in the cash account, and the other is invested in the constant-maturity bond with the closest maturity to the average time to replacement income cash flows. Cash flow dates are known in advance, so the time to cash flow for each future date is known too. Taking the average of these horizons across all accumulation dates, which range from the beginning of accumulation, shown in the second column of the table, to the retirement date, generates a value, and the bond with the closest maturity to this value is selected. For all tested strategies, this procedure leads to choosing the 15-year bond.

The quality of replication can be measured through the tracking error, defined as the volatility of excess returns over the benchmark, or as the relative return over the sample period. A relative return of 100% or above means that the replicating strategy preserved or even increased the purchasing power of savings in terms of replacement income, while a value of less than 100% means the purchasing power decreased. The first observation is that both replicating strategies have relative returns close to 100%, ranging from 100.90% to 104.55%. In all samples, their tracking errors are at most equal to 2.08%, and tend to be lower in the most recent ones. For each factor-matching procedure, switching from quarterly to monthly rebalancing slightly reduces the annual tracking error, but the impact is very small, five basis points per year at most.

The alternative strategies display much worse scores. The one that invests in the 15-year bond always has larger tracking errors, ranging from 1.79% to 4.20% per year depending on the period. It also has a bigger dispersion of relative returns across the samples, since these range from 95.72% to 116.13%. Finally, cash is always the worst option to replicate the bond because it has tracking errors greater than 10% per year, and, more importantly, it severely underperforms the benchmark in all samples. The worst relative return is for the individual saver from February 2004 to January 2015, who would have seen a loss of 1 – 51. 37% = 48. 63% in the purchasing power of his/her savings in terms of replacement income by investing in cash. These results echo those of Merton (2014), who points out that the purchasing power of an investment in Treasury bills continuously changes, with upward and downward moves.[94]

A Closer Look at a Specific Sample

To look more closely at one of the samples considered in Table 5, Figure 12 displays the weights of the level-matching strategy and the value of $1 invested in the benchmark or in one of the strategies intended to replicate it for an investor who retired in 2015. At any point in time, the level-matching portfolio is invested in two bonds, chosen to have the closest level exposures to those of the benchmark. Because the target exposure decreases in absolute value as the retirement date approaches, the bond maturities also tend to decrease over time. Thus, in February 2004, the portfolio is invested in the 18-year and 19-year bonds, and just before retirement, it is invested in the 8-year and 9-year bonds.

94 - See the figure on p. 5 of his article.

Table 5: Simulation of level-matching and duration-matching portfolios from the beginning of accumulation to the retirement date or 3 June 2019, whichever comes first

Retirement year	Beginning of accumulation	Strategy	Ann. return (%)	Gross relative return (%)	Ann. volatility (%)	Tracking error (%)	Max. relative drawdown (%)
2000	Jan. 1989	BNC	10.67	-	10.73	-	-
		GHP Lev., Q	10.76	100.90	11.50	2.06	3.61
		GHP Dur., Q	10.78	101.15	11.51	2.08	3.68
		GHP Lev., M	10.74	100.73	11.43	2.05	3.58
		GHP Dur., M	10.76	100.87	11.45	2.07	3.65
		15-year bond	10.98	103.19	12.95	3.90	11.37
		Cash	5.28	57.74	0.21	10.73	52.94
2005	Jan. 1994	BNC	8.69	-	11.51	-	-
		GHP Lev., Q	8.99	103.15	12.32	1.72	2.39
		GHP Dur., Q	9.00	103.26	12.40	1.77	2.49
		GHP Lev., M	8.96	102.81	12.26	1.70	2.29
		GHP Dur., M	8.96	102.85	12.34	1.74	2.37
		15-year bond	9.56	109.17	13.58	3.76	8.87
		Cash	3.93	61.12	0.18	11.51	53.25
2010	Jan. 1999	BNC	6.5	-	11.75	-	-
		GHP Lev., Q	6.97	104.44	12.49	1.45	2.99
		GHP Dur., Q	6.98	104.55	12.59	1.51	3.10
		GHP Lev., M	6.89	103.59	12.41	1.40	2.92
		GHP Dur., M	6.92	103.89	12.52	1.46	3.07
		15-year bond	6.89	103.57	14.22	4.20	11.56
		Cash	2.91	68.20	0.17	11.75	51.31
2015	Feb. 2004	BNC	7.80	-	12.19	-	-
		GHP Lev., Q	8.03	102.29	12.84	1.78	6.54
		GHP Dur., Q	8.12	103.27	13.00	1.84	6.55
		GHP Lev., M	8.01	102.07	12.82	1.78	6.49
		GHP Dur., M	8.09	102.93	12.98	1.82	6.46
		15-year bond	9.06	113.47	14.20	3.53	8.49
		Cash	1.42	51.37	0.14	12.19	53.24
2020	Feb. 2009	BNC	5.27	-	13.36	-	-
		GHP Lev., Q	5.74	104.71	13.78	1.08	1.39
		GHP Dur., Q	5.83	105.69	13.90	1.14	1.46
		GHP Lev., M	5.77	105.03	13.73	1.06	1.32
		GHP Dur., M	5.85	105.86	13.86	1.12	1.34
		15-year bond	6.80	116.13	13.13	2.53	5.57
		Cash	0.46	61.69	0.05	13.36	52.05

3. Efficient Factor Exposure Matching: Improving the Hedging Benefits of the Liability-Hedging Portfolio with Factor Investing

Retirement year	Beginning of accumulation	Strategy	Ann. return (%)	Gross relative return (%)	Ann. volatility (%)	Tracking error (%)	Max. relative drawdown (%)
2025	Jan. 2014	BNC	8.56	-	11.35	-	-
		GHP Lev., Q	8.81	101.25	11.57	0.68	1.84
		GHP Dur., Q	8.87	101.55	11.65	0.71	1.86
		GHP Lev., M	8.76	100.98	11.51	0.66	1.88
		GHP Dur., M	8.82	101.28	11.60	0.69	1.90
		15-year bond	7.69	95.72	10.02	1.79	6.94
		Cash	0.79	66.87	0.07	11.35	36.78

Notes: Retirement takes place on the first day of the year indicated in the first column. Goal-hedging portfolios are rebalanced every month or every quarter. Series of daily returns for the benchmark (BNC) and for goal-hedging portfolios rebalanced every quarter (Q) or every month (M) are simulated over a period that starts on the date indicated in the second column and ending at the retirement date or 3 June 2019, whichever comes first. At each rebalancing date, portfolios "GHP Lev." are weighted so as to have the same level exposure as the retirement bond, and portfolios "GHP Dur." are weighted so as to have the same modified duration. The first date of accumulation in the second column is chosen so as to ensure that the maturity of the last replacement income cash flow does not exceed 30 years, which is the longest maturity of the bonds available in the universe.

Figure 12: Replication strategies in accumulation for an investor retiring in January 2015
(a) Weights of level-matching strategy.

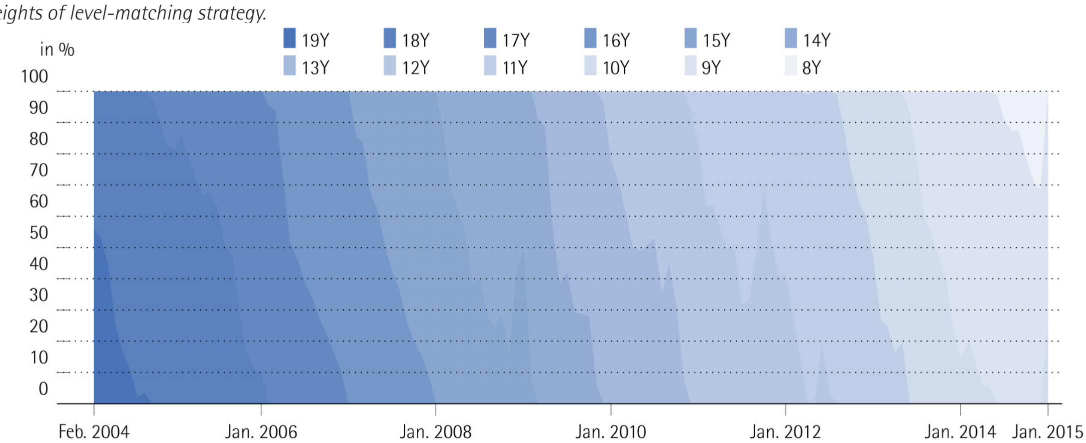

(b) Value of $1 invested in benchmark or in level-matching strategies.

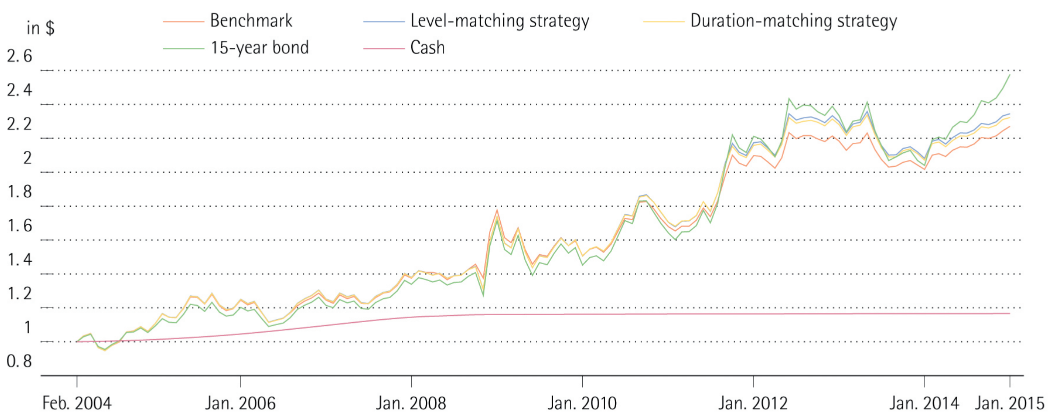

Notes: The benchmark is the price of the retirement bond for an investor who retired in January 2015 and sought 20 years of replacement income. The replicating portfolio is rebalanced every month, and in each month it is invested in two constant-maturity bonds, but these bonds change as the retirement date approaches. At the start of each month, the two bonds with the nearest level exposures to those of the benchmark are selected, by taking one bond with greater exposure and one with lower exposure. Weights are calculated so as to achieve the same level exposure as the benchmark.

The weights of the duration-matching strategy would be very similar, given that the level exposure of a bond is almost proportional to its modified duration.

Panel (b) visually shows that the two matching strategies are close to their benchmark, while the cash portfolio utterly fails to replicate it. The 15-year bond displays a growth rate that is more in line with that of the benchmark, but nevertheless not as close as the two matching portfolios. The analysis in Section 3.2.1 helps to understand these results: unlike a long-term bond, a cash investment does not benefit from the decrease in interest rates, and is in fact penalized by it since its return is simply a function of the level of rates. With short-term interest rates close to zero between November 2008 and January 2015, cash has almost a zero return, while the benchmark still tends to grow as a result of the passage of time (i.e., the decreasing maturity of cash flows) and of the decrease in long-term interest rates. In a period of increasing rates, the underperformance of cash would have likely been less severe, but one should not expect a reliable alignment of cash returns with retirement bond returns.

3.3.2 Decumulation

Level- or Duration-Matching Portfolios in Decumulation

The portfolio construction procedure in decumulation is exactly the same as in accumulation. At each rebalancing date, the two constant-maturity bonds with the nearest level exposure or the nearest modified duration to the benchmark are selected, and they are weighted so as to match its exposure or duration. The main difference with respect to the accumulation phase is that the target exposure or duration jumps at the beginning of every year, when a replacement cash flow is paid. These jumps arise because of the discontinuities in the bond price on cash flow payments. When the portfolio must be rebalanced on a cash flow payment date, we take as a target the exposure or duration of the benchmark just after the payment. This approach seems more natural than to take the exposure or duration just before the payment, because the number of shares of each constituent calculated at the beginning of the period is held constant for the entire period, until the next rebalancing date, so the calculation should be based on a forward-looking measure of the portfolio exposure or duration.

It should be emphasized that the factor-matching portfolios that we consider are self-financed and thus involve no cash withdrawals. The reason for this is that we want to simulate "mass-customized" funds, i.e. those designed for a group of individuals who share the same retirement date and the same cost-of-living adjustment objective, and thus have the same retirement bond. But it is ultimately their responsibility to decide when and how much to withdraw. By simulating the value of a self-financed portfolio, we abstract away from any individual-specific withdrawal schedule. Then, as explained below, we calculate the maximum annual withdrawal rate, which is the maximum percentage of the pre-retirement savings that can be withdrawn every year without leaving room for a shortfall.

Measuring the Quality of Replication: The

Maximum Annual Withdrawal Rate

The second difference with respect to accumulation lies in the choice of the metrics to assess the accuracy of the replication. Indeed, the ultimate test that a good replicating portfolio must pass is whether individuals who invest their savings in it can make all the scheduled withdrawals and whether they are left with a surplus or not. Obviously, ending with a surplus is less of a problem than being unable to sustain the targeted withdrawals for the entire decumulation period, if only because the surplus can be used to fund an annuity purchase for the remaining lifetime after the decumulation period, and because it can be passed on to heirs. However, a final surplus signals inefficient ex-post use of resources: more could have been withdrawn every year without running short of savings before the end of decumulation, but the individual was unaware at the time that he/she could decide how much to withdraw. The retirement bond is precisely intended to solve the problem of knowing at the start of decumulation how much can be withdrawn every year from the pension pot without depleting it before the term, and without discovering a final surplus that would leave the individual with the regret of not spending more. So, a perfect replicating portfolio should allow anyone retiring at the common retirement date to make annual withdrawals equal to the purchasing power of their savings in terms of replacement income at the time of retirement, and including the cost-of-living adjustment.

The tracking error is not a very useful metric to assess whether this goal is achieved, for two reasons. First, it is symmetric and treats upward and downward deviations from the benchmark equally; and second, it is a function of the short-term returns (usually daily, weekly or monthly) on the replicating portfolio and benchmark, and it is unclear how to infer from this aggregate measure whether the strategy was successful or not at serving the replacement income cash flows. The relative return over the decumulation period is not a satisfactory indicator either precisely because it considers only the values of the replicating portfolio and benchmark at the endpoints of the decumulation period, but it disregards what happens on intermediate cash flow dates.

A useful metric in decumulation is the maximum annual withdrawal rate, defined as the maximum percentage of pre-retirement savings that can be withdrawn every year without creating shortfall risk. In fact, this indicator can be calculated for any investment strategy in decumulation, not only one that aims to replicate the retirement bond. It is even possible to calculate it explicitly as a function of the returns of the fund in which the individual invests in decumulation, so that no trial and error is needed to search for the point at which the withdrawal rate becomes too high to be sustainable. Appendix G shows that if withdrawals grow at rate π every year and take place on dates T, $T+1$, $T+2$, ..., $T+\tau-1$, where τ is the total number of cash flows, and the gross fund returns between two dates s and t is denoted by $r_{s,t}$, then the maximum rate is

$$\delta = \frac{1}{\sum_{h=0}^{\tau-1} \frac{[1+\pi]^h}{r_{T+h,T+\tau-1}}}. \tag{3.6}$$

Table 6: Simulation of level-matching and duration-matching portfolios in decumulation

Retirement year	Strategy	Max. annual withdrawal rate (%)	Ratio of benchmark to rate (%)
2000	BNC	8.82	-
	GHP Lev., Q	9.26	95.23
	GHP Dur., Q	9.26	95.17
	GHP Lev., M	9.22	95.65
	GHP Dur., M	9.23	95.58
	1-year bond	6.74	130.85
	Cash	6.13	143.76
1999	BNC	7.78	-
	GHP Lev., Q	8.00	97.17
	GHP Dur., Q	8.00	97.19
	GHP Lev., M	7.99	97.34
	GHP Dur., M	7.99	97.36
	1-year bond	6.86	113.31
	Cash	6.31	123.34
1995	BNC	9.50	-
	GHP Lev., Q	9.77	97.18
	GHP Dur., Q	9.76	97.32
	GHP Lev., M	9.74	97.55
	GHP Dur., M	9.72	97.70
	1-year bond	7.81	121.60
	Cash	7.00	135.66
1990	BNC	9.56	-
	GHP Lev., Q	9.96	96.03
	GHP Dur., Q	9.97	95.90
	GHP Lev., M	9.91	96.49
	GHP Dur., M	9.93	96.35
	1-year bond	8.42	113.62
	Cash	7.45	128.37

Notes: Retirement takes place on the first day of the year indicated in the first column, and the individual makes twenty withdrawals from the pension pot, from the retirement date to nineteen years thereafter. Seven investment strategies for the decumulation phase are considered. The first ("BNC") invests in the theoretical retirement bond. The next four invest in goal-hedging portfolios ("GHPs") rebalanced every quarter or every month in such a way as to achieve the same exposure to the NSS level factor or the same modified duration as the retirement bond. The sixth strategy invests in the 1-year constant-maturity bond, and the last invests in cash. For each strategy, the maximum annual withdrawal rate is the maximum percentage of pre-retirement savings that can be withdrawn every year for twenty years without resulting in a deficit. For the retirement bond, it is equal to the theoretical purchasing power of $100 in terms of replacement income, i.e. the reciprocal of the bond price just before retirement. The last column of the table is the ratio of the withdrawal rate of the bond to that of each strategy.

In the special case where savings were not invested, i.e. held in notes and coins, the annual growth rate would be zero, so the maximum rate would be simply $1/\tau$. So, for a 20-year decumulation period, the individual would be able to withdraw 5% of pre-retirement savings. For a more realistic investment strategy, gross fund returns are expected to be greater than 1, so it is also expected that δ will be greater than $1/\tau$. The problem is that it is a function of all fund returns between the beginning of decumulation and the cash flow dates, so it is not known until the entire path of fund values is discovered. But remarkably, δ turns out to be path-independent when the fund in which savings are invested perfectly replicates the total return index of the retirement bond. The total return on the retirement bond between two dates is defined as the return of a self-financed portfolio fully invested in this asset, so it is the return of the portfolio assuming that any cash flow paid out by the bond between these dates is reinvested in the bond. The property that the withdrawal rate for such a fund is independent from the scenario is proved in Appendix G. As expected, the rate is equal to the reciprocal of the retirement bond price at the beginning of retirement, just before the first withdrawal takes place.

With the retirement bond, the amount that can be withdrawn every year without shortfall risk is known in advance.

Table 6 shows the maximum rate, calculated with Equation (3.6) and based on backtests of the competing strategies. Because we need a complete 20-year decumulation period, the latest retirement date that we consider is January 2000, since the last cash flow for this specification occurs in January 2019, and this paper is being written in 2019. The withdrawal rate of the retirement bond is also shown for reference. Unlike the rates of other strategies, which are only discovered after the last withdrawal, it is known ex-ante, as of the retirement date. So an individual could set a spending plan for retirement by choosing to withdraw an amount equal to accumulated savings multiplied by this rate, but unless they invest in a strategy that perfectly replicates this bond, there is no guarantee that they can sustain this plan for the twenty years of the decumulation period. The last column in the table is intended specifically to tell whether this was eventually possible or not, and if so, whether the individual was left with a final surplus. In detail, we divide the ratio of the retirement bond by the ratio of each strategy, so a ratio of 100% means that the strategy considered allowed the investor to make exactly the same annual withdrawals as the retirement bond, without a final surplus or a shortfall. A ratio greater than 100% indicates that the individual was unable to sustain the spending plan for the entire decumulation period, while a ratio less than 100% means that not only were they able to sustain it, but they would have been able to withdraw even more than scheduled. In other words, they would have ended with a surplus. The ratio in the last column can be loosely seen as a relative return of the bond with respect to each strategy, but a more complex measure than the point-to-point relative return from the start to the end of the 20-year decumulation period since it takes into account multiple intermediate dates.

In addition to the four replicating strategies, we also test a decumulation policy that invests only in cash, and one that invests in the one-year constant-maturity bond. The one-year maturity is selected by following the same rule as in accumulation: we consider the bond with the average level exposure or the duration closest to that of the retirement bond, and for both measures of exposure, this criterion leads to choosing the 1-year bond. As appears from the figures in Table 6, the four replicating strategies have withdrawal rates closer to that of the retirement bond than the one-year bond and the cash account in each sample period. Indeed, their ratios in the last column range from 95.17% to 97.70%, which implies a distance from 100% comprised between 2.30% and 4.83%. For the single bond and the cash portfolio, the distance is much greater, ranging from 13.62% to 43.76%. This confirms that attempts to explicitly match the factor exposures of the retirement bond "pay off" in the decumulation phase: strategies that do not control these exposures are much less effective at mimicking the characteristics of the retirement bond.

Table 6 also shows that in the four samples considered, the replicating strategies have ratios of less than 100%, while the bond and cash strategy have ratios of greater than 100%. This implies that the former strategies sufficiently outperformed the bond, allowing the investor to enjoy a final surplus, while the latter two did not allow the spending plan to be funded as expected. The underperformance of the cash account is especially striking: in the most recent sample, the spending rate permitted by this strategy was 6.13%, much lower than what would have been possible with the retirement bond (8.82%). The corresponding ratio is 8. 82/6. 13 = 143. 8%. To express this difference in dollar amounts, it means that for a $200,000 pension pot at retirement, the maximum annual withdrawal compatible with a cash investment would have been $12,260, versus $17,640 with the bond. This makes a big difference in terms of standard of living.

While it is tempting to conclude that replicating strategies outperform their benchmark and allow an individual to withdraw more than what would be permitted by the retirement bond, we warn the reader that this observation may be too specific to the samples considered here to warrant generalization. We expect these strategies to be systematically closer to the retirement bond than a single constant-maturity bond or a portfolio invested in cash, but not to outperform their benchmark in all market conditions. It should be remembered that the periods considered in these illustrations featured decreases in interest rates, and that rates were mostly much higher than they currently are in 2019. While we have no explanation for why high and decreasing rates would imply a higher performance for exposure-matching strategies than for the retirement bond, these peculiarities should be kept in mind when thinking of how the strategies would behave in different interest rate conditions.

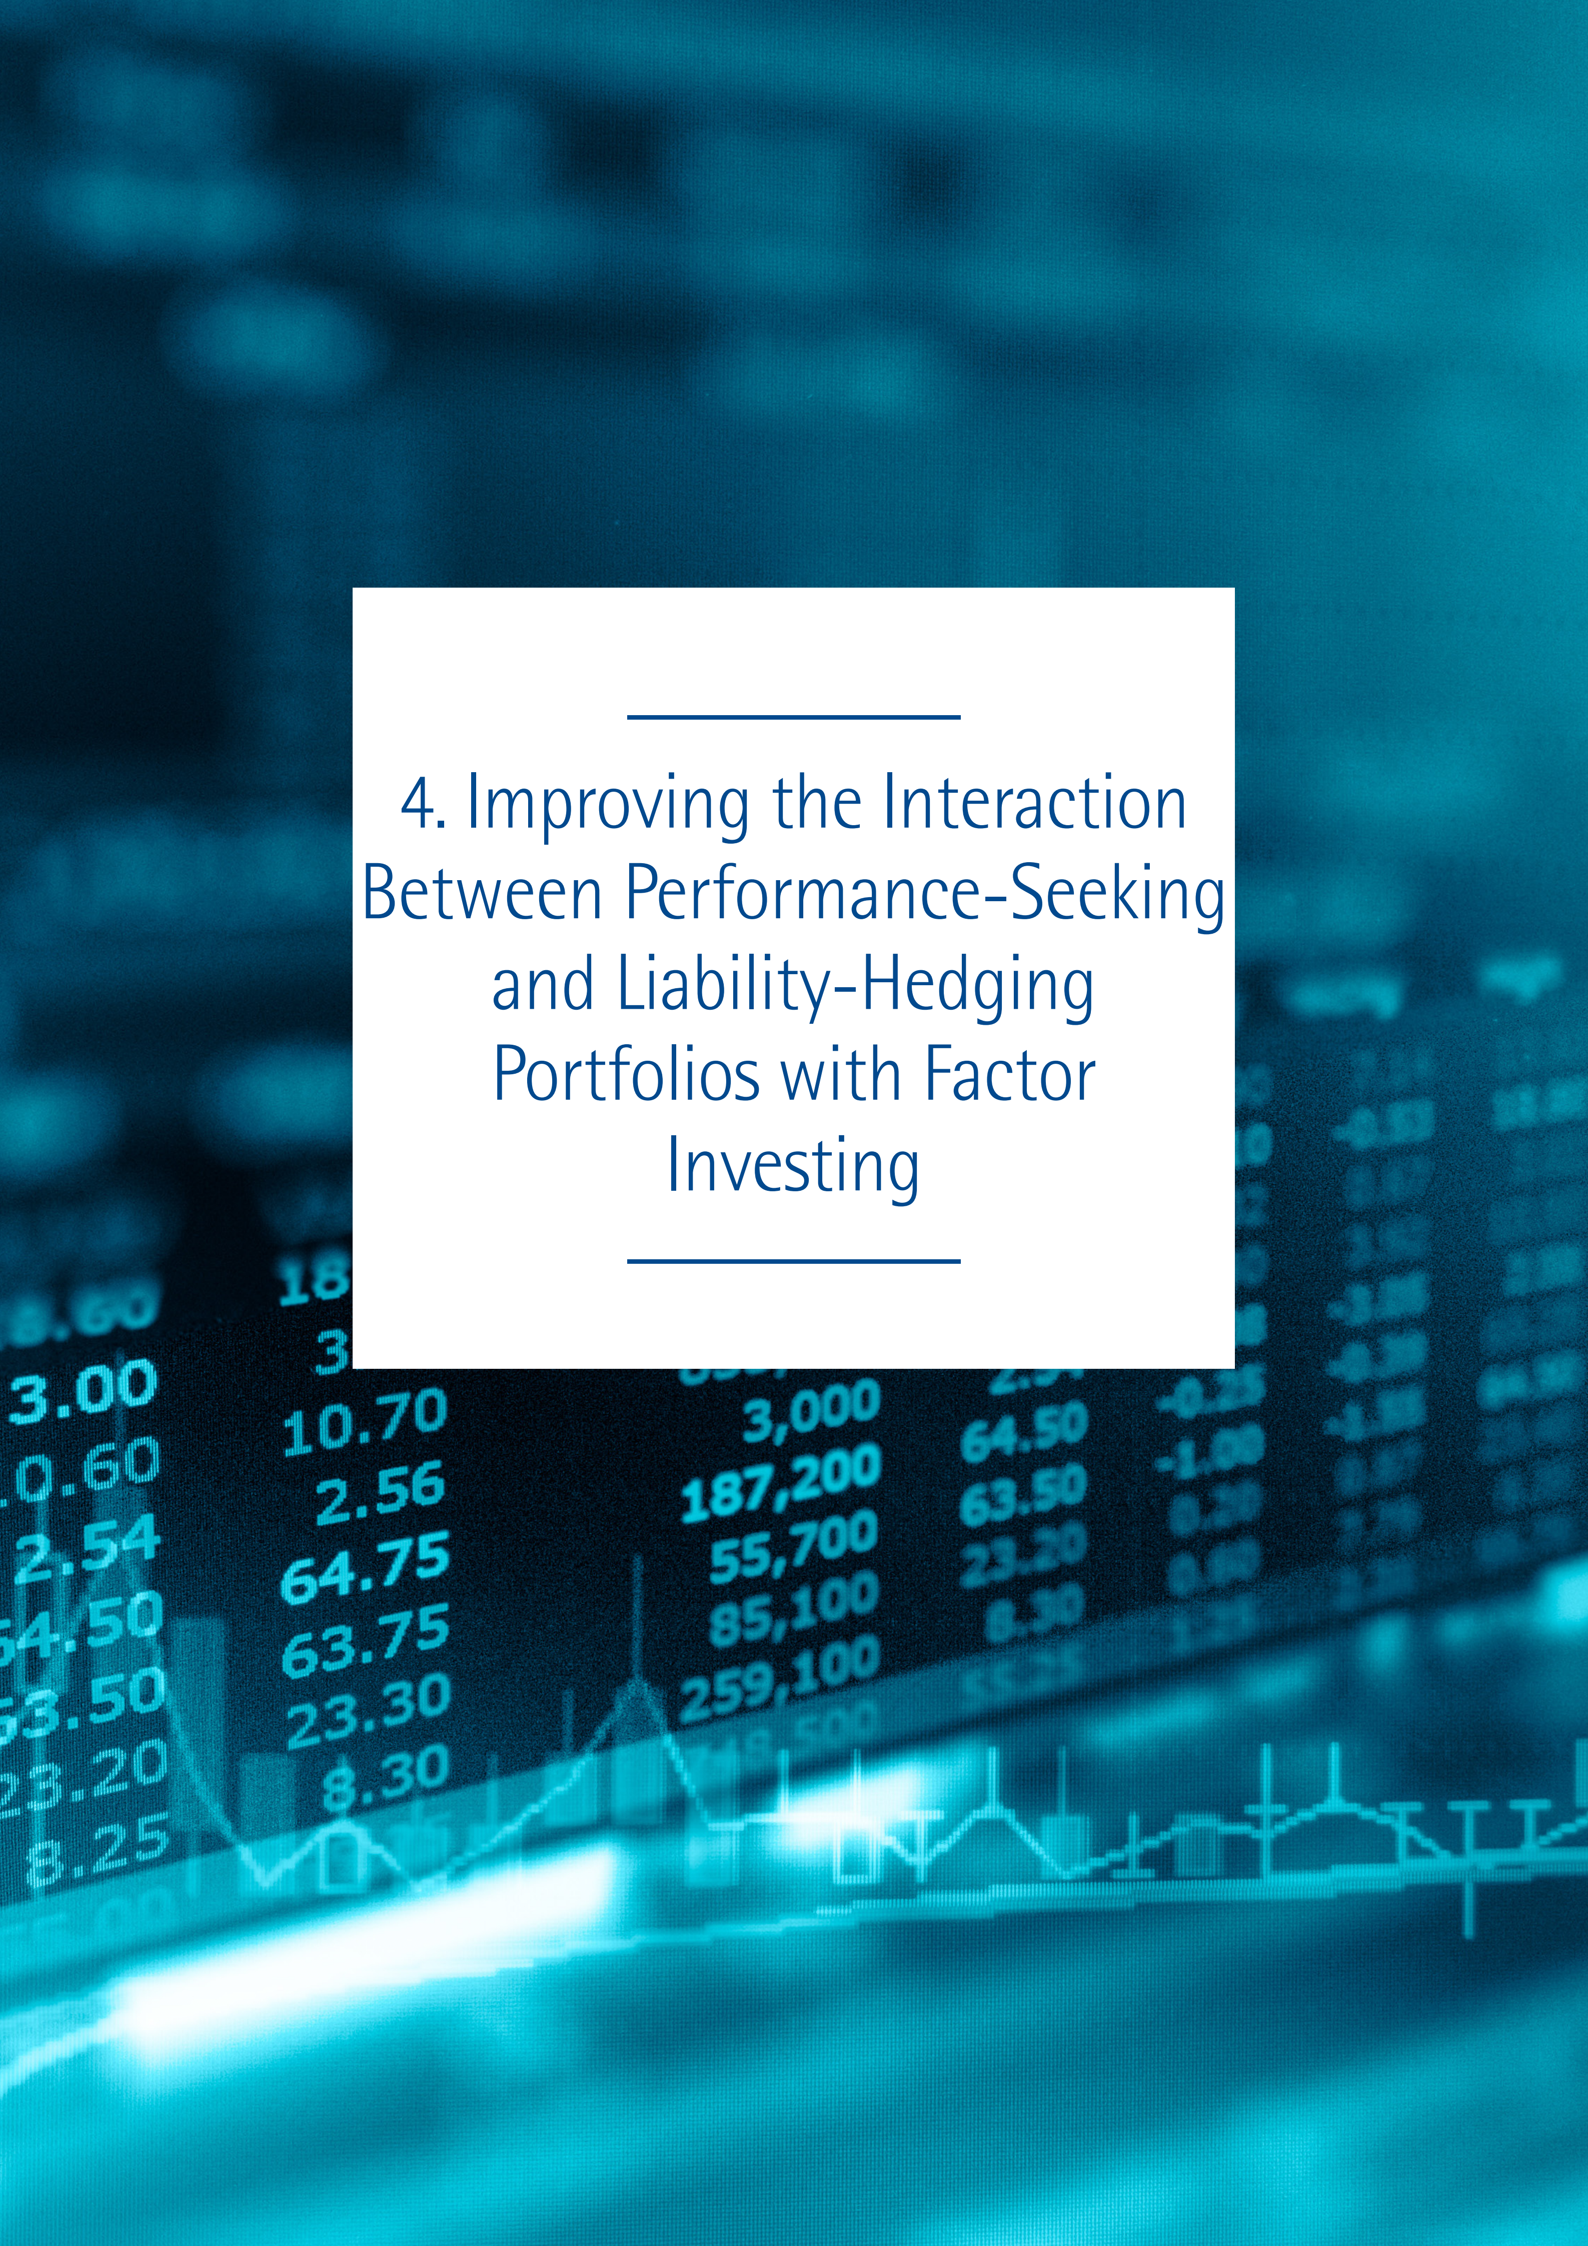

4. Improving the Interaction Between Performance-Seeking and Liability-Hedging Portfolios with Factor Investing

4. Improving the Interaction Between Performance-Seeking and Liability-Hedging Portfolios with Factor Investing

So far, we have focused on the performance-seeking portfolio (PSP) and the liability-hedging portfolio (LHP) in isolation. The liability-driven investing (LDI) paradigm, which has become a relatively standard approach to asset-liability management at pension funds, puts the two pieces together by recommending that pension assets be split between these two building blocks. It establishes a clear functional separation between a portfolio dedicated to the reduction of relative risk with respect to liabilities, and another portfolio whose objective is to outperform liabilities. However, the tracking error of the final mix will not only depend on the hedging qualities of the LHP, but also on those of the PSP, so that picking a PSP with better hedging properties could result in lower tracking error at the level of the blended portfolio. The formal argument is presented by Coqueret, Martellini and Milhau (2017) and is reproduced in Appendix J. This advantage can be dampened if the new PSP has less upside potential, so the gains in term of relative risk should be weighted against the possible loss in performance if such a loss arises.

In this section, we propose a factor framework to analyze these questions. First we introduce a multi-factor model to decompose the risk of assets and liabilities and highlight the overlaps and differences in their risk exposures. Second, we compare different methodologies with respect to their ability to construct equity portfolios with better liability-hedging properties, and we measure the performance gains that can be expected from using them in place of a standard equity index.

4.1 Measuring Factor Exposures in Assets and Liabilities

To find the similarities and differences in the risk exposures of assets and liabilities, we first need to identify a set of common factors that are relevant for both sides of the balance sheet.

Asset Classes and Liabilities

The investment universe contains equities and bonds. The equity class is split across seven portfolios that respectively represent the equity market as a whole and six rewarded equity factors. The broad equity index is a cap-weighted index of the 500 largest US stocks, and the other six portfolios consist respectively of value stocks (stocks with high bookto-market ratio), mid cap stocks, high momentum stocks (stocks with high past twelve-month return), low volatility stocks (stocks with low volatility measured over the past 104 weeks), high profitability stocks (stocks with high gross profit to total assets in past year) or low investment stocks (stocks with low growth in total assets). All these indices are cap-weighted and are borrowed from the US long-term track records of Scientific Beta. In the bond class, we make a distinction between sovereign bonds, represented by the Bloomberg Barclays US Treasury index, and corporate bonds, represented by the Bloomberg Barclays US aggregate Baa corporate index.

The liability process is chosen to represent the value of liabilities of an open-ended pension fund, with no finite horizon in the foreseeable future. If the fund is committed to making fixed payments, or payments rising at a predetermined rate, then the sources of risk that impact its liabilities are the same as those that affect interest rates. Thus, changes in the value of liabilities are represented with the returns on a constant-maturity bond, with a fixed maturity of 10 years. Section 3.2.1 above explains in detail how these returns are simulated out of time series of zero-coupon rates with constant maturities.

Table 7 summarizes the source and the availability period of each series. Equity index values are available every day from 19 June 1970 through 31 December 2016, and bond index values are available beginning in January 1973. However, they are updated only once per month until February 1994, and they become daily series thereafter. Finally, the zero-coupon rate series of 9- and 10-year maturities, which are used in the simulation of the 10-year constant maturity bond, are available from August 1971. So daily records exist for all asset and liability series from 1 March 1994 through 31 December 2016.

Choosing a Factor Model

Clearly, the model must apply to all asset classes present on the asset side, including in particular equities and bonds, and to liabilities, which have a bond-like nature since their value is the sum of discounted future cash flows. Having said this, one could favor the explanatory power of the model in time series, or its ability to explain the cross section of average returns, and these objectives are not necessarily aligned since factors that are good at explaining time-series variance can perform poorly when it comes to explaining differences in average returns. A well-known example is the market factor of equities, which explains a good fraction of common movements across equity portfolios, but is unable to capture most of the patterns in equity returns: this point was explicitly made by Chen, Roll and Ross (1986: 399). While explaining the cross section of average returns is the goal of asset pricing models, it

Table 7: Availability of raw series used for the simulation of assets, liabilities and factors

Raw series	Provider	Start date
Assets		
Scientific Beta cap-weighted index	Scientific Beta	19 June 1970
Scientific Beta indices	Scientific Beta	19 June 1970
Bloomberg Barclays US Treasury index	Datastream	30 January 1973 (monthly) 25 February 1994 (daily)
Bloomberg Barclays US Baa Corporate index	30 January 1973 (monthly)	25 February 1994 (daily)
Liabilities		
9-year and 10-year US zero-coupon rates	Fed website	16 August 1971
Factors		
Moody's Seasoned Baa Corporate Bond Yield	FRED	2 January 1986 (daily) 5 January 1962 (weekly)
10-Year Treasury Constant Maturity Rate	FRED	2 January 1962
3-Month Treasury Bill: Secondary Market Rate	FRED	4 January 1954
Fama-French equity factors (market, value, size)	Ken French's data library	1 July 1963
Fama-French momentum factor	Ken French's data library	3 November 1926
Betting-against-beta factor	Lasse Pedersen's data library	12 January 1930

Notes: This table summarizes the availability period for each raw series. Unless otherwise indicated, all series are daily.

is not the relevant question in asset-liability management. A key step in a sound ALM process is the construction of a liability-hedging portfolio that replicates the returns on liabilities as closely as possible, date by date and scenario by scenario. So, the time-series perspective should be favored, and a good factor model for ALM should capture much of the comovements of assets and liabilities.

To satisfy this requirement, factors that are linked by construction to the returns to explain are good candidates to guarantee good explanatory power.[95] Petkova (2006) introduces a model with factors that are not exclusively defined as portfolio returns or excess returns, but are still explicitly related to financial markets: these factors are the return on a broad equity index and the innovations to the dividend yield, the term spread, the credit spread and the short-term T-bill rate. While they produce decently large R-squares in time-series regressions of equity portfolios sorted on size and book-to-market, ranging from 55% to 88% (see her Table IV), they rank behind the three Fama-French equity factors, which generate R-squares ranging from 80% to 95% (see her Table III). So it appears that using equity factors defined as excess returns on equity portfolios is the most efficient way to achieve large time-series R-squares.

Equity factors, however, are not sufficient for an ALM factor model, as they leave much of the time-series variance of bond returns unexplained. This can be seen from results reported by Fama and French (1993): equity factors explain from 10% to 22% of the variance of government bonds and corporate bonds rated Baa and above, and as much as 33% of the variance of corporate bonds, a value which shows that these bonds represent an intermediate point between equities and fixed-income products. To better model bond returns, one needs to add interest rate-related factors. The approach of introducing additional factors has the advantage that the explanatory power of the model for equity portfolios is not sacrificed. Indeed, the R-square of a regression can never decrease when the number of regressors increases, although the gain depends on the marginal explanatory power of the new regressors with respect to those already present (see the formula in Appendix D). Fama and French (1993) show that the ex-post term premium and default premium, respectively measured as the realized excess return of long-term bonds over T-bills and the excess return of corporate bonds over Treasuries, explain from 79% to 90% of the variance of government bonds and corporate bonds rated Baa and above.

> So our first two interest rate factors are the change in the 3-month Treasury bill rate and…

We could follow the same approach and introduce the realized excess returns on Treasuries and corporate bonds as additional factors, but this would introduce redundancy between the factors and the returns to explain, as Treasuries and corporate bonds are already present in the asset universe. Then, the excess returns on these assets would be perfectly explained by the factors, but the 100% R-square would be a sort of statistical

95 - An alternative approach would be to use macro factors, as simple economic sense suggests that such variables summarize common sources of risk in the returns of stocks and bonds, but it is generally found that they capture a limited fraction of the time-series variance. For instance, Chan, Karceski and Lakonishok (1998) argue that portfolios mimicking macro factors capture less common variance in equities than the market factor or factors based on microeconomic attributes like size and book-to-market, and that only those mimicking the term or the default premium can compete with those micro factors. Flannery and Protopapadakis (2002) argue that macroeconomic news does have an impact on equity markets, but that this impact is not well captured by standard linear regression models with constant coefficients.

artifact due to the property that the dependent variable is a linear combination of the regressors. Moreover, the excess returns on long-term bonds are not a "pure" representation of changes in the level or the slope of the term structure. Indeed, Appendix H shows that the excess returns on long-term zero-coupon bonds aggregate a contribution of the change in the short-term rate and a contribution of the change in the term spread. For these reasons, we depart from the definition of Fama and French's bond-market factors as excess returns and instead favor factors defined directly from the yield curves. In their construction of a factor model to decompose the risk of the funding ratio, Kroon, Wouters and de Carvalho (2017) make a similar choice, and thus deviate from the model of Ransenberg, Hodges and Hunt (2012), in which interest rate-related risks, namely credit risk and real rate risk, are proxied with excess returns. As is well known, there are strong correlations between bond returns and interest rate changes, but the two variables do not have the exact same informational content, precisely because excess returns on long-term bonds depend on changes in the level – as measured by the short-term interest rate – but also on changes in the slope – measured by the spread between a long-term rate and the short-term rate. So our first two interest rate factors are the change in the 3-month Treasury bill rate and the change in the spread between the 10-year Treasury constant maturity rate, both obtained from the research database of the Federal Reserve of St. Louis.[96] However, these two variables may not accurately represent the returns on corporate bonds, which are also impacted by changes in the price of credit risk. Thus, we also use changes in the credit spread, defined as the difference between Moody's seasoned Baa corporate bond yield and the 10-year Treasury rate.[97] Many models described in the literature make use of similar variables (see e.g. Campbell and Viceira (2001); Campbell, Chan and Viceira (2003); Petkova (2006)).

Availability periods for the raw series involved in the construction of factors are shown in Table 7. Prof. French's website provides daily returns on equity factors as of 1 July 1963, and the daily records for the 3-month rate series, the 10-year

Table 8: Correlations and volatilities of factors (in %); data from March 1994 to December 2016

	Mkt-RF	SMB (L/S)	HML (L/S)	Mom (L/S)	BAB (L/S)	LEV	TSP	CSP
Mkt-RF	17.5							
SMB (L/S)	15.3	9.7						
HML (L/S)	−1.8	−2.2	10.3					
Mom (L/S)	−21.5	1.6	−37.9	16.5				
BAB (L/S)	−39.0	−22.0	21.5	27.7	12.0			
LEV	11.8	7.4	−5.6	2.0	8.8	0.7		
TSP	12.5	12.3	9.6	−16.9	−24.2	−46.3	1.0	
CSP	−19.6	−19.4	−4.1	14.4	14.2	−24.2	−49.9	0.8

Notes: Below-diagonal elements are correlations between factors, and diagonal numbers are annualized volatilities, i.e. those of weekly factor values from March 1994 to December 2016 multiplied by $\sqrt{52}$. Mkt-RF = equity market factor; SMB = equity size factor (small-minus-big); HML = equity value factor (high-minus-low); Mom = equity momentum factor; BAB = betting against beta equity factor; LEV = change in interest rate level; TSP = change in term spread; CSP = change in credit spread. Factors marked as "(L/S)" are the excess returns to long-short equity portfolios.

96 - See https://fred.stlouisfed.org/series/DTB3 and https://fred.stlouisfed.org/series/DGS10.
97 - The Baa rate is available at https://fred.stlouisfed.org/series/DBAA.

rate series and the Baa rate series from the FRED start respectively on 1 April 1954, 2 January 1962 and 2 January 1986. However, we prefer to use weekly data because daily data may raise concerns over asynchronous records across series from different sources and over possible interpolation in case genuine observations were not available. This allows in principle for a significant extension of the dataset in the past since the weekly Baa rate series start in January 1962, but our dataset for assets and liabilities starts only in March 1994 because the bond index series are monthly until that date. We therefore construct a dataset of factor values starting in March 1994 in those empirical exercises where the bond series are involved, but when these series are not needed, we start the analysis in August 1971, which is the latest of the start dates of other series according to Table 7. In all cases, the changes in interest rate factors are weekly changes, as are the excess returns on equity factors.

Correlations Between Factors

Table 8 contains the correlations and volatilities of factors. Correlations are moderate since their absolute values range from 0.7% to 49.9%, and many of them are negative. For instance, the correlation between the equity value and momentum factors is –37. 9%, which confirms the findings of Asness, Moskowitz and Pedersen (2013) (see their Tables I and II).

To perform a decomposition of volatility across factors, it may be easier to work with uncorrelated factors, since this avoids the problem of attributing covariance terms. Therefore, we also perform a linear transformation of factors to turn them into uncorrelated variables. The "orthogonalization" method is the minimum linear torsion (MLT) algorithm described in Meucci, Santangelo and Deguest (2015): it replaces the original set of factors with a new set of factors, each of which is a linear combination of the original variables, and the new factors have the properties of being uncorrelated from each other and as close as possible to the original ones. The distance with respect to original factors is measured with the sum of squared tracking errors, so the MLT method selects the linear transformation that minimizes this criterion. Table 9 displays the matrices that map the original factors into the uncorrelated ones, or the uncorrelated factors into the original ones. A remarkable property of both matrices is that they are symmetric (see Appendix B of Martellini and Milhau (2015) for a proof). Each uncorrelated factor is naturally associated with the original factor that it proxies for, so we keep the same names for the new factors. Each uncorrelated factor loads mainly on the corresponding original variable, as can be seen by comparing the magnitudes of diagonal and off-diagonal coefficients in Panel (a), but it also loads on the other factors, with a mixture of positive and negative coefficients, in such a way as to eliminate any correlation.

Explaining the Time Series of Asset and Liability Returns

Table 10 displays the factor exposures of the nine asset classes and the liability process together with the t-statistics, and Figure 13 provides a visual representation of exposures. Because equity portfolios and bond portfolios have very different exposures, they have been plotted separately for better legibility, and liabilities, which have a bond-like nature, have been grouped with bond portfolios. The first remark to be made is that the model explains a large percentage of the time-series variance of each portfolio, with all R-squares but one being greater than 90%. The lowest R-square is that of corporate bonds and is still 70%, and the highest is that of the broad equity index, which is very close to French's momentum exposure, at–0. 06, and that similarly,

Figure 13: Factor exposures of assets and liabilities, and percentage of variance explained by the factor model; data from March 1994 to December 2016

(a) Equity portfolios.

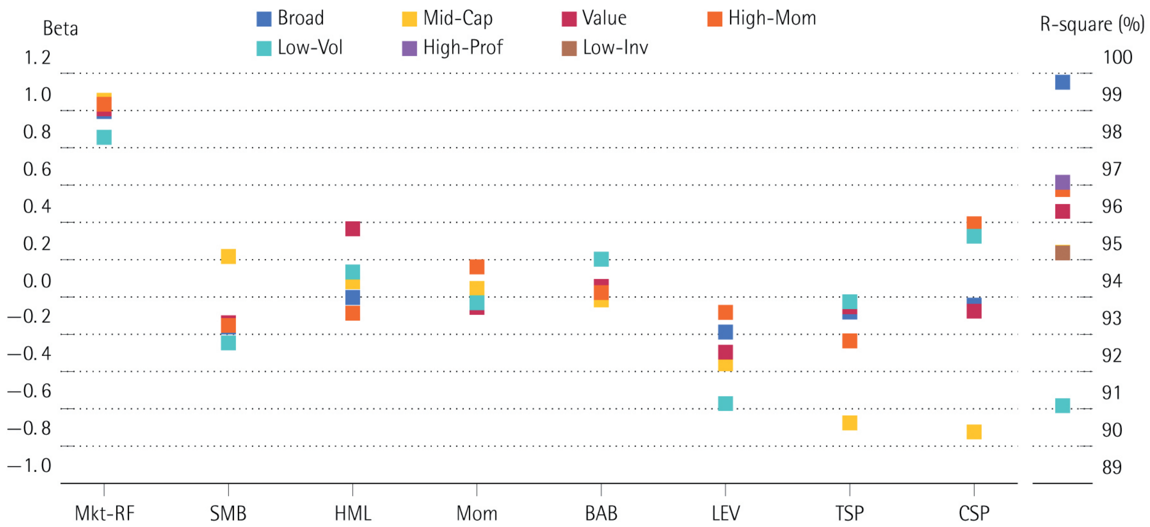

(b) Bond portfolios.

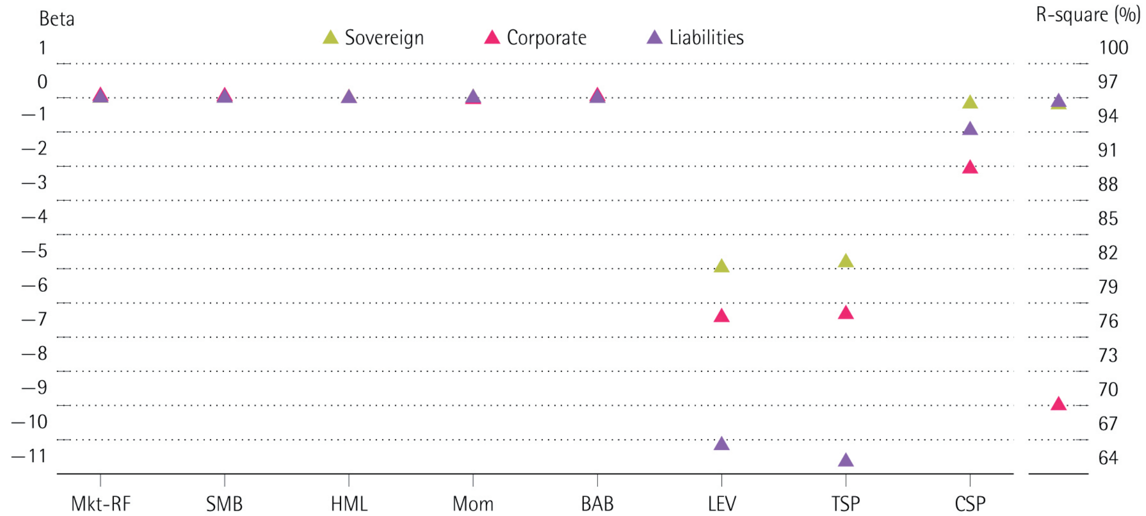

Notes: Weekly returns on assets and liabilities in excess of the risk-free rate are regressed on the eight factors. Data is from 4 March 1994 to 30 December 2016. Datapoints may overlap, so some of them may be hidden by others. The R-square is borrowed from Ken French's dataset. The figure shows the estimated betas and the model R-square. Mkt-RF = equity market factor; SMB = equity size factor (small-minus-big); HML = equity value factor (high-minus-low); Mom = equity momentum factor; BAB = betting against beta equity factor; LEV = change in interest rate level; TSP = change in term spread; CSP = change in credit spread.

the high momentum portfolio has the most negative value exposure, at –0. 09. This is not surprising because the high value selection tends to pick stocks with low past performance, while the high momentum selection focuses on stocks with high past performance. Exposures of equity portfolios to interest rate factors are mostly insignificant, with most of the t-statistics less than 2.0 in absolute value.

Bond portfolios have very different factor exposures from those of equities. Exposures to equity factors are close to zero and most often statistically insignificant. The maximum is at 0.06 for the market beta of corporate bonds, so corporate bonds are more equity-like than sovereign bonds. This is natural given that equities and corporate bonds are claims written by the same entities. This is in line with the results of Fama and French (1993), who find that market betas and the explanatory power of equity factors are higher for Baa-rated and low-grade bonds than for investment-grade bonds (see their Table 6). More generally, the factor exposures of bonds are concentrated among the three bond factors.

All of them of course have negative exposures to changes in the level of rates, and the magnitude

Table 10: Regression of assets and liabilities on factors; data from March 1994 to December 2016

Factor	Mkt-RF	SMB	HML	Mom	BAB	LEV	TSP	CSP
Broad	0.99	–0.16	–0.00	–0.01	0.01	–0.19	–0.08	–0.04
	(629.75)	(–60.43)	(–0.94)	(–4.38)	(2.10)	(–3.74)	(–2.05)	(–0.99)
Mid-Cap	1.06	0.22	0.08	0.05	–0.02	–0.36	–0.68	–0.72
	(134.10)	(16.36)	(5.81)	(5.15)	(–1.22)	(–1.43)	(–3.43)	(–3.27)
Value	1.01	–0.14	0.37	–0.06	0.06	–0.30	–0.05	–0.08
	(154.41)	(–12.58)	(31.63)	(–7.57)	(5.33)	(–1.42)	(–0.32)	(–0.41)
High-Mom	1.03	–0.15	–0.09	0.16	0.02	–0.08	–0.24	0.39
	(177.34)	(–15.55)	(–8.44)	(24.54)	(2.49)	(–0.44)	(–1.61)	(2.39)
Low-Vol	0.86	–0.25	0.13	–0.03	0.20	–0.57	–0.03	0.33
	(103.20)	(–17.62)	(9.12)	(–3.35)	(15.08)	(–2.16)	(–0.12)	(1.39)
High-Prof	0.93	–0.17	–0.25	0.02	–0.01	0.01	0.08	–0.20
	(175.78)	(–18.55)	(–26.43)	(2.76)	(–0.78)	(0.08)	(0.59)	(–1.34)
Low-Inv	0.91	–0.13	0.14	0.01	0.10	–0.36	0.11	–0.23
	(140.84)	(–12.13)	(12.62)	(1.51)	(9.12)	(–1.77)	(0.68)	(–1.28)
Sovereign	–0.00	0.00	–0.00	0.00	0.00	–4.97	–4.81	–0.17
	(–1.53)	(1.30)	(–1.53)	(0.72)	(0.22)	(–97.61)	(–120.51)	(–3.87)
Corporate	0.05	0.04	–0.01	–0.04	0.05	–6.42	–6.33	–2.07
	(9.40)	(4.61)	(–0.98)	(–6.59)	(5.16)	(–34.90)	(–43.86)	(–12.75)
Liabilities	–0.00	–0.00	–0.01	–0.00	–0.01	–10.16	–10.64	–0.95
	(–0.42)	(–0.57)	(–2.30)	(–0.94)	(–1.46)	(–98.13)	(–130.94)	(–10.35)

Notes: Weekly returns on assets and liabilities in excess of the risk-free rate are regressed on factors. Data is from 4 March 1994 to 30 December 2016. t-stats are in parentheses. Mkt-RF = equity market factor; SMB = equity size factor (small-minus-big); HML = equity value factor (high-minus-low); Mom = equity momentum factor; BAB = betting against beta equity factor; LEV = change in interest rate level; TSP = change in term spread; CSP = change in credit spread.

of this exposure is close to the average duration. The US Treasury index and the US Baa corporate bond index have respective average durations of 5.42 and 6.48 years between March 1994 and December 2016, and liabilities are constructed as a sovereign bond index with a constant maturity of 10 years. The corporate bond index also has the most negative exposure to the credit spread factor, while both sovereign bond portfolios are less sensitive to this factor. Finally, it can be noted that exposures to changes in the slope measured as the term spread are not very different from level exposures. A mathematical argument is presented in Appendix H: the result is established for zero-coupon bonds for simplicity, but it gives a sense of how portfolios of coupon-paying bonds behave.

4.2 Decomposing the Risk of Equity-Bond Strategies

Typical strategies used in ALM combine equities and bonds. While bonds are mainly used in the liability-hedging portfolio, for their large correlation with liabilities, and equities are primarily intended as performance-seeking assets, it is important to remember that in principle, equities can also enter the LHP, just as bonds are eligible for the PSP. The separation theorem for ALM (see e.g. Martellini and Milhau (2012)) implies that the relevant distinction is between building blocks with well-defined objectives, respectively the maximization of the Sharpe ratio and the maximization of the correlation with liabilities, as opposed to between asset classes. Any asset class can enter any building block, even though the hedging portfolio ends up being dominated by bonds. In this section, we consider arbitrary portfolios invested in equities and/or bonds so as to show how the choice of an allocation impacts the risk-return profile, and we analyze in particular the risk of each strategy to perform an attribution to the eight factors.

The six portfolio strategies are as follows:

• A portfolio fully invested in a broad US equity index weighted by capitalization;

• An equity-only portfolio with 50% in the broad index and 10% in each of the six long-only equity factor indices (mid cap, value, high

Table 11: Summary statistics of ALM strategies, from March 1994 to December 2016

	Equ.-broad	Equ.-fac	60-40 Equ.-Bnd	40-60 Equ.-Bnd	Low-Vol	Bnd
Annualized return (%)	9.30	9.66	8.55	7.79	9.45	5.92
Volatility (%)	17.33	16.99	10.02	6.98	15.01	4.81
Sharpe ratio	0.39	0.42	0.60	0.76	0.46	0.71
Maximum drawdown (%)	53.78	53.12	33.86	22.71	49.09	8.11
Tracking error (%)	21.32	21.01	13.92	10.60	18.88	5.33
Information ratio	0.08	0.09	0.06	0.01	0.09	–0.33
Maximum relative drawdown (%)	77.38	74.56	59.88	50.83	68.24	39.88

Notes: Statistics are based on weekly returns from 4 March 1994 to 30 December 2016. Volatilities and tracking errors are annualized. Annualized return is geometric. The Sharpe ratio is the difference between the annualized returns on the strategy and the risk-free asset, divided by the volatility. The value of the risk-free asset is the continuously compounded interest rate on 3-month Treasury bills. The information ratio is the difference between the annualized returns on the strategy and the liability process, divided by the tracking error. The maximum relative drawdown is the maximum drawdown in the funding ratio, equal to assets divided by liabilities.

momentum, low volatility, high profitability and low investment). The portfolio is rebalanced every quarter towards these target weights;

- A portfolio containing 60% of the previous equity portfolio and 40% of bonds, split with equal weights between sovereign and corporate bonds;
- A portfolio containing 40% of equities and 60% of bonds, with the same composition for each building block as before;
- A strategy replicating the returns on the low volatility index;
- A bond-only portfolio with 50% of sovereign bonds and 50% of corporate bonds, again rebalanced quarterly.

Risk and Return of Strategies

Weekly returns are simulated from March 1994 through December 2016, and Table 11 displays standard risk and return metrics. With an annual return of 5.92%, bonds underperformed equities but they had much lower volatility, so they posted a greater Sharpe ratio, at 0.71 versus 0.39 for the broad equity portfolio and 0.46 for low volatility stocks. Unsurprisingly, the bond portfolio had the lowest tracking error with respect to liabilities, at 5.33%, since these have a bond-like nature. Note that no attempt was made here to align the duration of the bond portfolio with that of liabilities, so the 5.33% tracking error could be further improved. However, the information ratio of bonds is negative, meaning that the sovereign-corporate bond portfolio underperformed the 10-year constant maturity bond. This underperformance occurred because the two bond indices on the asset side had shorter durations on average than liabilities, and interest rates followed a decreasing trend.

To outperform liabilities, one had to introduce stocks in the asset mix, but this comes at the cost of increased tracking error with respect to liabilities since the broad index displayed a relative standard deviation of 21.32%. A large tracking error penalizes the information ratio, so low volatility stocks are interesting from this perspective since they retain some of the good performance of equities while having a lower tracking error than a broad index. In this sample, they delivered a tracking error of 18.90%, leading to an information ratio of 0.09. The property of low volatility stocks of being better bond trackers than the general population of stocks has been noted by Coqueret, Martellini and Milhau (2017) (see their Exhibits 1 and 2). Moreover, in a period where the low volatility anomaly shows up, investors who substitute these stocks for a broad index benefit from increased performance in addition to lower tracking error, and the combination of both effects produces a superior information ratio. Such was the case in the sample period, since low volatility equities outperformed the market, at 9.45% versus 9.30% per year.

More generally, introducing factors in the equity bucket led to better returns in this period, since the equity portfolio containing the broad index and the six factor indices posted an annual return of 9.66%, versus 9.30% for equities in general. This addition also resulted in slightly lower volatility and tracking error, at respectively 16.99% and 21.01%, which represents a reduction of about 30 basis points with respect to the broad index. But as expected, all equity-only portfolios have larger relative risk, whether measured as the tracking error or as the maximum relative drawdown.

4. Improving the Interaction Between Performance-Seeking and Liability-Hedging Portfolios with Factor Investing

Risk Decomposition Across Factors

The contribution of a constituent to the risk of a portfolio is defined by Maillard, Roncalli and Teiletche (2010), but this definition requires adaptation here for two reasons. First, it applies to the ex-ante volatility of a portfolio, while we are interested in decomposing the realized volatility. Second, a portfolio is invested in assets, and its return is the weighted sum of the asset returns. Here, we want to make a decomposition across risk factors, but the portfolio is not physically invested in these factors, if only because some of them are not investable, e.g. those that are represented by interest rate changes. To decompose the risk of strategies across the eight factors, we thus regress their excess returns over the risk-free rate on factors, which leads to a decomposition of excess returns of the form

$$\tilde{r} = c + \beta_1 F_1 + \cdots + \beta_8 F_8 + \varepsilon, \qquad (4.1)$$

where ε denotes the residual. Because residuals are uncorrelated from factors, we obtain the usual split of variance across "systematic" variance explained by factors and "idiosyncratic" variance:

$$\sigma^2 = \sum_{i,j=1}^{8} \beta_i \beta_j \mathbb{C}\mathrm{ov}[F_i, F_j] + \sigma_\varepsilon^2.$$

Based on this equation, it is shown in Appendix I that the contribution of the ith factor to volatility can be defined as

$$q_i = \frac{\beta_i \mathbb{C}\mathrm{ov}[r, F_i]}{\sigma}, \qquad (4.2)$$

and the contribution of idiosyncratic risk as

$$q_\varepsilon = \sigma[1 - R^2],$$

where R^2 is the R-square of the regression (4.1). As shown in the Appendix, these nine contributions (eight factors plus the residual) add up to the volatility of the portfolio, so the contributions divided by volatility add up to 100%.

The definition of factor contributions is formally identical to the definition of asset contributions to volatility in Maillard, Roncalli and Teiletche (2010), with the factors replacing assets and the portfolio exposures replacing asset weights. The second difference between a decomposition across assets and one across factors is that the latter involves a contribution of specific risk, since the factors do not explain 100% of the time-series variance. But if the factors are chosen to have good explanatory power in the time series, this contribution should be small. The final point to be made is that the decomposition procedure can be adapted to the decomposition of relative risk with respect to liabilities, i.e. the tracking error. To perform this decomposition, one simply estimates the exposures of excess returns with respect to liabilities on the factors through a time-series regression.

Table 12 displays the relative contributions to absolute and relative risks. Some of them are negative, which is explained by looking at Equation (4.2): a factor contributes negatively to volatility or the tracking error if the portfolio has negative exposure to a factor with which it covaries positively, or positive exposure to a factor with which it covaries negatively. But overall, the negative contributions present in Table 12 are small in magnitude. In all portfolios containing some equity, the biggest contribution is that of the market factor. This results from the combination of relatively large betas, reported in Table 13, and a large factor volatility, at 17.5% per year (Table 8): in regressions of absolute returns, i.e., excess returns over the risk-free rate, betas range from 0.86 to 0.99 for equity-only portfolios, and from 0.40 to 0.59 for equity-bond portfolios, and they are not substantially modified when returns are taken in excess of liability returns.

Betas of absolute returns with respect to the three interest rate factors are larger in size, but these factors have much lower volatilities than the market (see Table 8), so they explain a much lower fraction of risk, which does not exceed 7.86% for portfolios containing some equities, while the market factor alone accounts for at least 89.89%. Overall, the market factor explains most of the absolute risk of portfolios containing equities, so the typical policy portfolios that allocate 40% or 60% of their assets to equities have a much less balanced risk allocation than

Table 12: Relative contributions of factors to risk of ALM strategies (in %), from March 1994 to December 2016
(a) Absolute risk: volatility of excess returns of assets over risk-free rate

Factor	Equ.-broad	Equ.-fac	60-40 Equ.-Bnd	40-60 Equ.-Bnd	Low-Vol	Bnd
Mkt-RF	100.39	101.12	100.35	89.89	91.68	−1.10
SMB	−0.56	−0.56	−0.34	−0.10	0.84	−0.60
HML	0.00	0.00	−0.01	−0.01	1.17	0.06
Mom	0.17	−0.11	0.02	0.24	0.84	−0.63
BAB	−0.14	−0.68	−0.70	−0.82	−2.97	1.05
Total equity factors	99.86	99.76	99.32	89.20	91.55	−1.21
LEV	−0.08	−0.10	−1.14	−0.36	−0.23	19.09
TSP	−0.06	−0.07	0.44	8.80	−0.01	80.85
CSP	0.04	0.05	0.31	−0.58	−0.22	−10.84
Total interest rate factors	−0.10	−0.12	−0.39	7.86	−0.47	89.10
Unexplained	0.24	0.36	1.06	2.94	8.92	12.10
Total	100.00	100.00	100.00	100.00	100.00	100.00

(b) Relative risk: volatility of excess returns of assets over liabilities

Factor	Equ.-broad	Equ.-fac	60-40 Equ.-Bnd	40-60 Equ.-Bnd	Low-Vol	Bnd
Mkt-RF	73.98	73.93	62.95	51.72	66.38	2.46
SMB	−0.95	−0.88	−0.79	−0.62	−0.64	1.07
HML	0.01	0.07	0.08	0.09	1.13	0.12
Mom	0.09	−0.17	−0.08	0.11	0.67	1.15
BAB	−0.28	−0.74	−0.80	−0.98	−3.19	−1.24
Total equity factors	72.85	72.20	61.35	50.32	64.34	3.56
LEV	6.08	6.23	7.61	8.77	6.29	10.29
TSP	21.82	22.24	30.62	38.92	25.65	70.07
CSP	−1.55	−1.56	−1.20	−0.83	−2.43	1.68
Total interest rate factors	26.34	26.91	37.04	46.87	29.51	82.04
Unexplained	0.81	0.89	1.61	2.82	6.15	14.40
Total	100.00	100.00	100.00	100.00	100.00	100.00

Notes: Figures have been rounded, so that totals reported in the table may slightly differ from those calculated by summing terms.

their dollar allocation. This result is reminiscent of the example given by Qian (2005) of a 60-40 stock-bond portfolio, in which most portfolio risk comes from equities.

In excess returns, equity factors including notably the market factor are still dominant, but to a lesser extent than in absolute returns, and interest rate factors explain from 26.34% to 46.87% of the risk of portfolios containing equities. This bigger fraction comes mainly from the term spread, which explains about 20% to 40% of the volatility of excess returns. Panel (b) of Table 13 reveals that although term spread betas are slightly larger than level betas, they are roughly similar, as was already observed at the constituent level in Figure 13, so their higher contributions can be attributed to a larger covariance with excess returns. As a matter of fact, excess returns over liabilities correlate more with changes in the term spread than with changes in the level. Correlations (not shown in the tables) range from 17.2% to 19.7% for the level, and from 42.9% to 71.9% for the slope.

4.3 Equity Portfolios With Improved Liability-Hedging Properties

Our objective here is to construct equity portfolios with better liability-hedging properties than a broad cap-weighted equity index, which is a standard benchmark for most pension funds.

Table 13: Betas of ALM strategies, measured with data from March 1994 to December 2016
(a) Regressions of excess returns over the risk-free rate

Factor	Equ.-broad	Equ.-fac	60-40 Equ.-Bnd	40-60 Equ.-Bnd	Low-Vol	Bnd
Mkt-RF	0.99	0.98	0.59	0.40	0.86	0.03
SMB	−0.16	−0.14	−0.07	−0.04	−0.25	0.02
HML	−0.00	0.02	0.01	0.00	0.13	−0.01
Mom	−0.01	0.01	−0.00	−0.01	−0.03	−0.02
BAB	0.01	0.03	0.02	0.02	0.20	0.02
LEV	−0.19	−0.22	−2.43	−3.52	−0.57	−5.68
TSP	−0.08	−0.10	−2.32	−3.41	−0.03	−5.56
CSP	−0.04	−0.06	−0.55	−0.76	0.33	−1.11

(b) Regressions of excess returns over liabilities

Factor	Equ.-broad	Equ.-fac	60-40 Equ.-Bnd	40-60 Equ.-Bnd	Low-Vol	Bnd
Mkt-RF	1.00	0.98	0.59	0.40	0.86	0.03
SMB	−0.16	−0.13	−0.07	−0.04	−0.24	0.03
HML	0.01	0.04	0.02	0.02	0.15	0.01
Mom	−0.00	0.01	0.00	−0.00	−0.03	−0.02
BAB	0.01	0.03	0.03	0.03	0.21	0.03
LEV	9.97	9.95	7.74	6.65	9.59	4.48
TSP	10.56	10.55	8.32	7.23	10.62	5.08
CSP	0.90	0.89	0.40	0.19	1.27	−0.16

Notes: Betas are estimated by regressing weekly excess returns from 4 March 1994 to 30 December 2016 on the eight factors.

4. Improving the Interaction Between Performance-Seeking and Liability-Hedging Portfolios with Factor Investing

Coqueret, Martellini and Milhau (2017) address this question at the level of individual stocks, by searching for the most "liability-friendly" stocks, when the source of risk in liabilities is the level of interest rates, and they find that those with low volatility and/or high dividend yield stand out for their above-average hedging properties. They also show that minimizing the portfolio volatility instead of weighting stocks by capitalization further reduces the tracking error of the equity portfolio with respect to liabilities. In this section, we also seek to improve the liability-hedging properties of the equity portfolio used as the performance building block, but our constituents are equity portfolios rather than individual stocks, each portfolio being chosen to provide exposure to a rewarded factor.

Construction Methods

We test three construction methods:

- The first portfolio is calculated by minimizing the distance between the exposures of the portfolio and those of liabilities to the eight factors. Portfolio exposures are obtained by weighting the constituents' exposures, and the distance is defined as the sum of the eight squared differences in exposures. Other notions of distance could be used, which would result in different portfolios, but this particular one has the advantage of involving a quadratic and convex optimization program, which is fast to solve;
- The second portfolio minimizes the systematic part of the tracking error with respect to liabilities. The decomposition of the tracking error is obtained from the decomposition of returns into a systematic part (explained by the factors) and an idiosyncratic term, as per Equation (4.1). The excess return of the portfolio with respect to liabilities is a linear combination of the factors, plus an uncorrelated residual, so the tracking error can be written as

$$\tilde{\sigma}^2 = \sum_{i,j=1}^{8}[\beta_{pi} - \beta_{Li}][\beta_{pj} - \beta_{Lj}]\mathbb{Cov}[F_i, F_j] + \tilde{\sigma}_{\varepsilon}^2,$$

where β_{pi} and β_{Li} are the respective betas of the portfolio and liabilities with respect to the i^{th} factor, and $\tilde{\sigma}^2$ is the variance of the idiosyncratic term. To construct the second portfolio, we minimize the first term in the right-hand side. This is equivalent to minimizing the distance between exposures, as in the first weighting scheme, if factors were uncorrelated and had equal variances, but here, factors are correlated and have unequal variances, so the two methods are expected to give different portfolios;[100]
- The "max ENB portfolio" maximizes the ENB, defined as the reciprocal of the sum of squared relative factor contributions to tracking error. Unlike in Section 2.3, where the ENB was taken with respect to uncorrelated factors extracted by statistical analysis from a covariance matrix, we consider the raw factors here, skipping the transformation into uncorrelated variables. As a consequence, the relative factor contributions do not add up to 100%, as part of the portfolio risk cannot be attributed to factors.

We estimate the parameters of each strategy, namely the covariances and the betas, over the full sample, i.e. from September 1971 to December 2016. Portfolios are then rebalanced every quarter (on the first Friday of each quarter) towards these weights, so they are fixed-mix strategies. The backtests therefore involve look-ahead bias by construction.

100 -In unreported results, we also constructed a minimum total tracking error portfolio and found it to be very close to the minimum systematic error portfolio in terms of all risk and return indicators. This result is, of course, not a surprise since the eight-factor model explains at least 90% of the variance of each equity index, leaving only a small amount of residual risk.

4. Improving the Interaction Between Performance-Seeking and Liability-Hedging Portfolios with Factor Investing

In addition to the aforementioned weighting schemes, we also consider portfolios subject to tracking error constraints with respect to the broad index, with tracking error caps at 2%, 4% or 6% per year. We exclude the broad index from the underlying universe of these strategies. As appears from Table 2 above, the tracking errors of the four individual factors range from 4.05% to 5.75%. Thus, the 6% constraint is most often not binding, while the 2% one is too tight at some dates, especially after imposing non-negativity constraints. The compatibility of the tracking error constraint can be formally checked at each rebalancing date by first minimizing the tracking error under long-only constraints: if the minimum is greater than 2%, then no portfolio under 2% can be constructed, and in this situation, we bypass the optimization step, which would be a distance minimization, an ENB maximization or a tracking error minimization, and we instead use the portfolio that minimizes the tracking error with respect to the broad index.

In the analysis of these equity strategies, we do not need the bond indices, so the sample period can be extended back before 1994. All interest series – the equity series and the zero-coupon rate series – begin in the middle of August 1971, so we begin the analysis on 1 September 1971.

Table 14 contains risk and performance indicators for the various strategies. The first observation is that in this sample, minimizing the systematic tracking error with respect to liabilities and maximizing the ENB without a tracking error constraint with respect to the cap-weighted index are two equivalent objectives. This is confirmed by looking at the weights in Figure 14. Second, it seems that some of the portfolios violate the tracking error constraint: for instance, the minimum distance portfolio with a 2% constraint has a tracking error of 2.01%, and the minimum tracking error portfolios with a 4% constraint have tracking errors (with respect to liabilities) of 4.04%. This is not the result of undue tolerance with respect to the constraint in the optimizer: in fact, the minimum tracking error attainable in the sample by mixing the six factors is 1.12%, so it is lower than 2%. With the 2% and 4% caps, the portfolio returned by the optimizer has an ex-ante tracking error exactly equal to the cap, meaning that the constraint is binding. But portfolios are rebalanced every quarter, while ex-post tracking errors are calculated from weekly returns, so these ex-post errors differ from the ex-ante ones, and the constraint may appear to be violated in some cases.

If one takes the ex-post tracking error with respect to liabilities as the quantitative measure of the hedging ability, the minimum systematic tracking error portfolio and the maximum ENB portfolios, which coincide in fact with the portfolio minimizing the tracking error with respect to the cap-weighted index, have the best scores: they have a tracking error of 17.32% with respect to liabilities, which represents a substantial reduction with respect to the 19.16% of the cap-weighted index. Distance minimization lags behind, since the reduction in the initial 19.16% is at most 40 basis points. As a general rule, constraining the portfolio to have lower tracking error with respect to the broad index than a cap implies a loss of optimality for the criterion considered in the form of a larger distance, a lower ENB or a larger tracking error.

Table 14: Simulation of equity strategies with improved liability-hedging properties with data from September 1971 to December 2016

	Equ.-broad	Min distance				Min systematic tracking error			
Target TE w.r.t. Equ.-broad (%)	–	None	2	4	6	None	2	4	6
Annual return (%)	10.15	11.84	11.17	11.60	11.84	10.47	10.86	10.56	10.47
Volatility (%)	16.39	17.73	15.88	16.52	17.73	14.61	15.37	14.75	14.61
Sharpe ratio	0.32	0.39	0.39	0.40	0.39	0.38	0.39	0.38	0.38
Max drawdown (%)	53.78	58.28	50.82	55.70	58.28	49.09	49.23	48.41	49.09
Relative to liabilities									
Tracking error (%)	19.16	20.59	18.76	19.40	20.59	17.32	18.19	17.51	17.32
Information ratio	0.09	0.16	0.14	0.16	0.16	0.11	0.13	0.12	0.11
Max relative drawdown (%)	77.38	67.90	68.34	65.92	67.90	68.24	70.51	68.74	68.24
Relative to broad index									
Tracking error (%)	-	5.72	2.01	3.98	5.72	4.84	2.05	4.04	4.84
Information ratio	-	0.30	0.51	0.36	0.30	0.07	0.35	0.10	0.07
Max relative drawdown (%)	-	28.57	13.53	25.05	28.57	38.39	15.00	32.07	38.39
Distance exposures	13.85	13.06	13.56	13.24	13.06	13.71	13.80	13.74	13.71
ENB in TE	1.84	2.01	1.89	1.92	2.01	2.13	1.91	2.06	2.13

		Max ENB			
Target TE w.r.t. Equ.-broad (%)		None	2	4	6
Annual return (%)		10.47	10.69	10.52	10.47
Volatility (%)		14.61	16.07	14.77	14.61
Sharpe ratio		0.38	0.36	0.38	0.38
Max drawdown (%)		49.09	47.20	47.95	49.09
Relative to liabilities					
Tracking error (%)		17.32	18.92	17.52	17.32
Information ratio		0.11	0.12	0.11	0.11
Max relative drawdown (%)		68.24	73.63	69.08	68.24
Relative to broad index					
Tracking error (%)		4.84	2.00	4.04	4.84
Information ratio		0.07	0.27	0.09	0.07
Max relative drawdown (%)		38.39	8.43	31.62	38.39
Distance exposures		13.71	3.92	13.76	13.71
ENB in TE		2.13	1.99	2.07	2.13

Notes: Portfolios are rebalanced on the first Friday of each quarter from 3 September 1971 to 30 December 2016. The first portfolio is fully invested in the broad cap-weighted index, and the alternative portfolios are invested in six equity factor indices (mid cap, value, high momentum, low volatility, high profitability and low investment), plus possibly the broad index. Only those that are not subject to a tracking error constraint with respect to the broad index are allowed to contain that index. For the others, the maximum tracking error with respect to the broad index is set to 2%, 4% or 6% per year. Each alternative portfolio has a fixed-mix allocation calculated over the entire sample, from 1971 to 2016. The minimum distance portfolio minimizes the distance between the portfolio exposures to the eight factors and those of liabilities. The minimum systematic tracking error portfolio minimizes the systematic part of the tracking error (i.e. the part explained by the eight factors), and the maximum ENB portfolio maximizes the effective number of bets in tracking error. Some of the portfolios are subject to a tracking error constraint with respect to the broad index: the target is 2%, 4% or 6% per year. Statistics are based on weekly returns, and volatilities and tracking errors are in annual terms. The last two rows display the Euclidean distance between the exposures of each portfolio and those of liabilities, and the ENB in tracking error.

Figure 14: Rebalancing weights of equity strategies with improved liability-hedging properties, estimated with data from September 1971 to December 2016

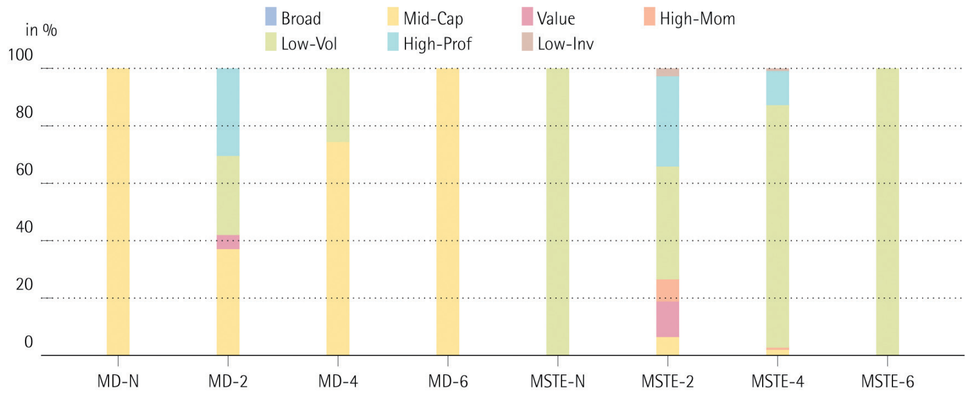

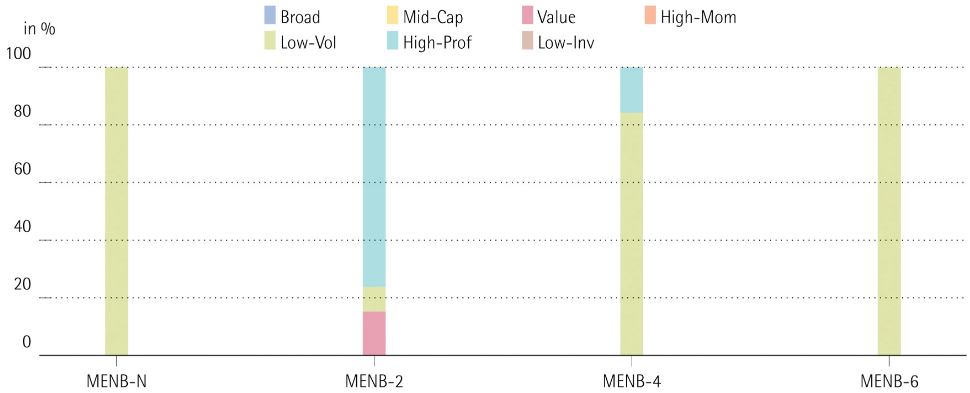

Notes: The investment universe of each strategy includes the mid cap, value, high momentum, low volatility, high profitability and low investment factor indices, and the universe of strategies with no tracking error constraint with respect to the broad index also includes the broad index itself. MD portfolios minimize the distance between the factor exposures of a portfolio and those of liabilities. MSTE portfolios minimize the ex-ante tracking error with respect to liabilities, i.e. the part of the tracking error that is explained by the eight ALM factors. MENB portfolios maximize the effective number of uncorrelated bets (ENB) in tracking error. Some of the portfolios are subject to a tracking error constraint with respect to the broad index. The maximum tracking error is 2%, 4% or 6% per year and is indicated in the portfolio name.

Figure 14 shows the rebalancing weights of strategies with no tracking error constraint. The low volatility index largely dominates the allocations that minimize tracking error. This result is in line with the findings of Coqueret, Martellini and Milhau (2017), who show that selecting low volatility stocks is an effective way to identify those with good liability-hedging properties. This happens because a portfolio of the six factors is more volatile than liabilities, so, as Coqueret, Martellini and Milhau (2017) mathematically show, a decrease in volatility implies a decrease in tracking error.[101] In the presence of long-only constraints, investing fully in the low volatility index minimizes volatility and, in this case, minimizes the tracking error too. Note, however, that although they are two related objectives, minimizing tracking error and minimizing volatility are not mathematically equivalent, because the tracking error also depends on the correlation between the portfolio and liabilities.

101 - This property would not hold if the asset portfolio was less volatile than liabilities. In that case, it is an increase in volatility that would imply a lower tracking error.

To summarize the previous results, can we rank the three portfolio construction methods on their ability to deliver an equity portfolio with better liability-hedging properties? If they are sorted on their ex-post tracking error, which is a standard metric, distance minimization does not prove to be an effective way of constructing a more liability-friendly portfolio, as it is dominated by the tracking error minimization schemes. In this sample, the lowest tracking errors are attained by maximizing ENB or minimizing the systematic tracking error.

4.4 Benefits of Equity Portfolios with Improved Hedging Properties in Liability-Driven Investing Strategies

We now want to evaluate the benefits, if any, of choosing a PSP with improved liability-hedging properties in a liabilitydriven investing (LDI) strategy. The intuition is explained by Coqueret, Martellini and Milhau (2017): the tracking error of the portfolio with respect to liabilities depends on the respective liability-hedging abilities of the PSP and LHP. A PSP generally has larger tracking error than an LHP since it is not designed to hedge liabilities, but using a PSP with less relative risk allows the manager to allocate a bigger fraction of assets to this building block without increasing relative risk at the portfolio level. If the PSP with the lower tracking error has expected return at the same level or above the PSP with the higher error, the change in allocation will result in higher expected return for the LDI strategy, but if it has lower performance potential, the effect for the LDI strategy is ambiguous: does a larger allocation to an equity portfolio with lower returns lead to higher or lower average returns? The answer depends on how large the loss in expected return within the PSP is compared to the increase in allocation. From the numbers in Table 14, it appears that in our sample, all alternative PSPs outperform the broad index, so we expect the LDI strategies using these PSPs to have unambiguously higher average returns than the one invested in the broad index. Our specific objective in what follows is to provide a quantitative measure of the gain in expected return.

Comparison of LDI Strategies with Identical Equity Allocations

Before we turn to the analysis of risk-matched strategies, defined as LDI strategies with allocations adjusted to equate some measure of relative risk with respect to liabilities, we consider a simple substitution exercise in which the broad index is replaced with one of the alternative PSPs and the allocation at the building block level is left untouched. Results are shown in Table 15. While the percentage allocation to the PSP in each LDI strategy is a constant 40%, the ex-post tracking error of a strategy is not exactly equal to 40% of the tracking error of a portfolio because it is calculated with weekly returns, while portfolios are rebalanced every quarter. Nevertheless, the ex-post error of each strategy is close to 40% of that of its PSP. In any case, the dispersion of the tracking errors of the strategies is lower than that of the PSPs. Tracking errors range from 17.32% to 20.59% for PSPs, and from 6.83% to 8.01% for strategies. Similarly, mixing each PSP with the constant-maturity bond implies a narrowing of the range of maximum relative drawdowns. The worst losses relative to liabilities range from 65.92% to 77.38% for PSPs, and from 33.25% to 40.79% for the strategies in which they are included. So LDI strategies are less different from each other in terms of relative risk than the various PSPs are. In the next series of illustrations, we re-calculate each one's allocation to equities so as to make them strictly equivalent in terms of relative risk.

It is also seen in Table 15 that none of the LDI strategies using an alternative equity portfolio underperforms the strategy invested in the broad index, and that most of them actually outperform it. Because all strategies have the same LHP and the same relative allocation to their building blocks, this is the consequence of the outperformance of each PSP with respect to the broad index. The gain in annual performance ranges from 0 to 80 basis points, a return spread that is relatively modest but translates into sizable differences in cumulative returns over the entire sample period, which is longer than 40 years. Since we face an asset-liability management problem here, the relative return of assets with respect to liabilities is a more relevant metric than the return on assets. It is equivalent to the change in the funding ratio, defined as the ratio of assets to liabilities. With the default PSP, namely the broad index, the funding ratio increases by 64.80% from 1971 to 2016. Here, the best performing portfolio is the one that minimizes the distance with respect to liability exposures without a tracking error constraint. The annual return of the corresponding LDI strategy is 80 basis points above that of the strategy that uses the broad index, and the funding ratio increases by 129.62% as opposed to 64.80%. In other words, starting from a funding ratio of 100% in 1971, the alternative PSP leads to a ratio of 229.62% in 2016, versus 164.80% for the broad index. In contrast, the portfolio that minimizes the tracking error with respect to liabilities has the lowest annual return in the sample, at 10.47% per year, and the final funding ratio after 45 years is 165.03%, barely above 164.80%. But the LDI strategy that uses this PSP has lower relative risk than the one invested in the broad index, both in terms of tracking error and maximum relative drawdown, so one could increase the allocation to equities beyond 40% so as to attain the exact same level of relative risk, and the gain in funding ratio would be more substantial, as we show below.

To use the terminology of the mean-variance theory, a trade-off between (relative) risk and return arises between two LDI strategies if one of them has larger tracking error and higher expected return than the other. Here, most of the LDI strategies invested in alternative PSPs dominate the one using the broad index in terms of mean-variance as they have both lower tracking error and higher expected returns. A trade-off arises only for some of the minimum distance portfolios, which lead to higher errors and higher relative returns. In these cases, it is unclear ex ante whether adjusting the allocation to equities to make the relative risk identical to that of the strategy invested in the broad index will lead to higher or lower average return than this strategy. Measuring the change in performance is the focus of what follows.

Risk-Matched LDI Strategies

So far, we have considered LDI strategies with identical allocation to equities, but different levels of relative risk. To compare them from a mean-variance standpoint (relative to liabilities),

Table 15: Simulation of liability-driven investing strategies with different performance-seeking portfolios, with data from September 1971 to December 2016

Performance-seeking portfolio	Equ.-broad	Min distance				Min systematic tracking error			
Target TE w.r.t. Equ.-broad (%)	–	None	2	4	6	None	2	4	6
PSP									
Annual return (%)	10.15	11.84	11.17	11.60	11.84	10.47	10.86	10.56	10.47
Tracking error (%)	19.16	20.59	18.76	19.40	20.59	17.32	18.19	17.51	17.32
Max relative drawdown (%)	77.38	67.90	68.34	65.92	67.90	68.24	70.51	68.74	68.24
LDI strategy									
Allocation to PSP (%)	40.00	40.00	40.00	40.00	40.00	40.00	40.00	40.00	40.00
Annual return (%)	9.72	10.52	10.10	10.33	10.52	9.72	9.93	9.77	9.72
Volatility (%)	9.46	9.65	9.30	9.41	9.65	9.15	9.24	9.16	9.15
Cumulative relative return (%)	64.80	129.62	93.50	112.40	129.62	65.03	79.97	68.26	65.03
Tracking error (%)	7.51	8.01	7.36	7.58	8.01	6.83	7.16	6.90	6.83
Information ratio	0.16	0.25	0.22	0.24	0.25	0.18	0.20	0.18	0.18
Max relative drawdown (%)	40.79	33.25	32.80	32.19	33.25	33.48	35.11	33.92	33.48

		Max ENB			
Target TE w.r.t. Equ.-broad (%)		None	2	4	6
PSP					
Annual return (%)		10.47	10.69	10.52	10.47
Tracking error (%)		17.32	18.92	17.52	17.32
Max relative drawdown (%)		68.24	73.63	69.08	68.24
LDI strategy					
Allocation to PSP (%)		40.00	40.00	40.00	40.00
Annual return (%)		9.72	9.91	9.75	9.72
Volatility (%)		9.15	9.37	9.17	9.15
Cumulative relative return (%)		65.03	78.72	67.07	65.03
Tracking error (%)		6.83	7.44	6.90	6.83
Information ratio		0.18	0.19	0.18	0.18
Max relative drawdown (%)		33.48	37.60	34.22	33.48

Notes: Portfolios are rebalanced on the first Friday of every quarter from 3 September 1971 to 30 December 2016. All liability-driven investing strategies invest 40% of their assets in an equity portfolio, and the remaining 60% in the 10-year constant-maturity bond, and they differ through the composition of their equity building block, which is either the broad cap-weighted index or an alternative portfolio invested in six equity factor indices and possibly the broad index. The six factors are mid cap, value, high momentum, low volatility, high profitability and low investment, and the broad index is only included when no tracking error constraint with respect to the broad index is imposed. When this constraint is present, the cap on the annual tracking error is 2%, 4% or 6%. Minimum distance portfolios minimize the distance between their factor exposures and those of liabilities, and maximum ENB portfolios maximize the effective number of uncorrelated bets (ENB) in tracking error. The other portfolios minimize either the systematic part of the tracking error, i.e. the part explained by the eight factors. The PSPs are constructed with parameter values estimated over the period from September 1971 to December 2016.

one would have to specify some risk aversion parameter, but such parameters are notoriously hard to estimate, so we use a parameter-free approach which involves comparing strategies with the same level of relative risk. Specifically, we take as a reference point the LDI strategy invested in the broad cap-weighted index and the 10-year bond, which plays the role of a perfect liability-hedging portfolio, with respective weights of 40% and 60%. As shown in Table 15, it has a tracking error of 7.51% and a maximum relative drawdown of 40.79% in the period from September 1971 to December 2016. For each alternative PSP, we calculate the allocation to equities that makes the tracking error or the maximum relative drawdown of an LDI strategy invested in this PSP and the perfect LHP equal to that of the reference strategy. This is done by trial and error: we simulate fixed-mix strategies invested in the two building blocks and rebalanced every quarter, and we search for the allocation that gives the same relative risk as the reference strategy over the sample period. The risk-matching allocations therefore involve look-ahead bias by construction.

Table 16 displays the results for the risk-matched strategies. Whenever the alternative PSP has lower relative risk than the broad index, the risk-matching allocation to equities in the alternative LDI strategy is greater than 40%. This is the case for all alternative PSPs when the risk measure is the maximum relative drawdown, and for all but three minimum distance portfolios when risk is measured as the tracking error. But in all cases, the risk-matching allocations to the PSP are higher with the relative drawdown criterion. In other words, an investor willing to accept the same drawdown risk as with the reference strategy can allocate more to the alternative PSP than one who targets the same level of tracking error. The downside for the former investor is that by matching the maximum relative drawdown, they face significantly larger tracking errors than the latter. The tracking error criterion can be regarded as more conservative because by construction it involves the same level of tracking error as with the reference strategy but also happens to lead to lower drawdowns.

The higher allocation to equities is expected to result in higher expected return, but as explained previously, the change in the performance of the LDI strategy is a function of both the new equity allocation and the return spread between the two PSPs. The first observation to take from Table 16 is that for all alternative PSPs, a performance gain is achieved, which can be measured either on an annual basis by looking at the standalone performance of the LDI strategy, or on a cumulative basis by looking at the total relative return of the strategy with respect to liabilities. As an illustration of the remark at the beginning of this paragraph, it is not the largest reduction in tracking error that necessarily leads to the greatest performance gain. In Table 16, the LDI strategy with the unconstrained minimum tracking error portfolio and a matched tracking error has 43.99% in equities, which is the largest share across the array of strategies in this table, and its final funding ratio is 171.29%, which is greater than the 164.80% achieved with the reference strategy, albeit not a huge gain. In contrast, the LDI strategy with the unconstrained minimum distance portfolio and a matched tracking error has only 37.57% in equities because this portfolio has greater tracking error than the broad index, but it leads to a funding ratio of 220.61%, which represents a gain, and a much bigger one than with the minimum tracking error PSP. This is because the minimum distance portfolio has the largest annual return in the sample, at 11.84% versus 10.15% for the broad index and 10.47% for the minimum tracking error portfolio.

4. Improving the Interaction Between Performance-Seeking and Liability-Hedging Portfolios with Factor Investing

Using Non-Cap-Weighted Equity Factor Indices

As explained in Section 2.2, it is possible to combine exposure to a rewarded factor and improved diversification of idiosyncratic risks with respect to cap weighting by constructing a factor index with an alternative weighting scheme. So far, we have used cap-weighted factor indices, but the long-term track records of the Scientific Beta database also contain versions of these factor indices with a non-capitalization-based weighting scheme. Coqueret, Martellini and Milhau (2017) show that more substantial decreases in tracking error with respect to the broad cap-weighted index can be achieved by deviating from cap-weighting indexation, and especially by adopting a variance minimization approach or by weighting stocks by the inverse of their volatility. This result is the motivation for the last section of this paper, in which we examine how sensitive our results are to a change in the weighting method of the equity factor indices.

Table 16: Simulation of risk-matched liability-driven investing strategies with data from September 1971 to December 2016

Performance-seeking portfolio	Equ.-broad	Min distance				Min systematic tracking error			
Target TE w.r.t. Equ.-broad (%)	–	None	2	4	6	None	2	4	6
PSP									
Annual return (%)	10.15	11.84	11.17	11.60	11.84	10.47	10.86	10.56	10.47
Tracking error (%)	19.16	20.59	18.76	19.40	20.59	17.32	18.19	17.51	17.32
Max relative drawdown (%)	77.38	67.90	68.34	65.92	67.90	68.24	70.51	68.74	68.24
LDI strategy with matched tracking error									
Allocation to PSP (%)	40.00	37.57	40.83	39.66	37.57	43.99	41.98	43.54	43.99
Annual return (%)	9.72	10.42	10.13	10.32	10.42	9.81	9.98	9.85	9.81
Volatility (%)	9.46	9.53	9.32	9.39	9.53	9.25	9.29	9.25	9.25
Cumulative relative return (%)	64.80	120.61	95.56	111.33	120.61	71.29	84.02	74.19	71.29
Tracking error (%)	7.51	7.51	7.51	7.51	7.51	7.51	7.51	7.51	7.51
Information ratio	0.16	0.25	0.22	0.24	0.25	0.17	0.20	0.18	0.17
Max relative drawdown (%)	40.79	31.48	33.40	31.95	31.48	36.35	36.60	36.50	36.35
LDI strategy with matched relative maximum drawdown									
Allocation to PSP (%)	40.00	50.87	51.32	52.88	50.87	50.40	47.72	49.67	50.40
Annual return (%)	9.72	10.92	10.43	10.76	10.92	9.94	10.13	9.98	9.94
Volatility (%)	9.46	10.43	9.86	10.18	10.43	9.50	9.53	9.49	9.50
Cumulative relative return (%)	64.80	169.99	121.27	152.86	169.99	80.97	95.58	84.12	80.97
Tracking error (%)	7.51	10.22	9.47	10.06	10.22	8.62	8.55	8.58	8.62
Information ratio	0.16	0.24	0.20	0.22	0.24	0.17	0.19	0.17	0.17
Max relative drawdown (%)	40.79	40.79	40.79	40.79	40.79	40.79	40.79	40.79	40.79

		Max ENB			
Target TE w.r.t. Equ.-broad (%)		None	2	4	6
PSP					
Annual return (%)		10.47	10.69	10.52	10.47
Tracking error (%)		17.32	18.92	17.52	17.32
Max relative drawdown (%)		68.24	73.63	69.08	68.24
LDI strategy with matched tracking error					
Allocation to PSP (%)		43.99	40.41	43.52	43.99
Annual return (%)		9.81	9.92	9.83	9.81
Volatility (%)		9.25	9.38	9.25	9.25
Cumulative relative return (%)		71.29	79.52	72.81	71.29
Tracking error (%)		7.51	7.51	7.51	7.51
Information ratio		0.17	0.19	0.18	0.17
Max relative drawdown (%)		36.35	37.92	36.79	36.35
LDI strategy with matched relative maximum drawdown					
Allocation to PSP (%)		50.40	44.09	49.19	50.40
Annual return (%)		9.94	10.02	9.95	9.94
Volatility (%)		9.50	9.53	9.47	9.50
Cumulative relative return (%)		80.97	86.58	81.75	80.97
Tracking error (%)		8.62	8.20	8.50	8.62
Information ratio		0.17	0.18	0.17	0.17
Max relative drawdown (%)		40.79	40.79	40.79	40.79

Notes: Portfolios are rebalanced on the first Friday of every quarter from 3 September 1971 to 30 December 2016. The reference liability-driven investing strategy invests 40% of assets in the broad equity index, and the remaining 60% in the 10-year constant-maturity bond. The other strategies use alternative equity portfolios as their performance-seeking building block, and the 10-year bond as the liability-hedging portfolio. They are invested in six equity factor indices (mid cap, value, high momentum, low volatility, high profitability and low investment), and possibly the broad index. Only those that are not subject to a tracking error constraint with respect to the broad index are allowed to contain this constituent, and in the others, the tracking error is capped to 2%, 4% or 6% per year. Minimum distance portfolios minimize the distance between their factor exposures and those of liabilities, and maximum ENB portfolios maximize the effective number of uncorrelated bets (ENB) in tracking error. The other portfolios minimize either the systematic part of the tracking error, i.e. the part explained by the eight factors. In each of the liability-driven investing strategies using a PSP other than the broad index, the allocation to the PSP is calculated over the full sample period so as to match the tracking error of the strategy that uses the broad index.

Table 17: Simulation of equity strategies with improved liability-hedging properties and made up of minimum variance factor indices, with data from September 1971 to December 2016

	Equ.-broad	Min distance				Min systematic tracking error			
Target TE w.r.t. Equ.-broad (%)	–	None	2	4	6	None	2	4	6
Annual return (%)	10.15	12.97	12.90	12.90	13.05	12.65	12.90	12.90	12.79
Volatility (%)	16.39	14.91	14.99	14.99	14.31	13.19	14.99	14.99	13.62
Sharpe ratio	0.32	0.54	0.53	0.53	0.57	0.59	0.53	0.53	0.58
Max drawdown (%)	53.78	52.47	44.96	44.96	50.58	42.99	44.96	44.96	43.32
Relative to liabilities									
Tracking error (%)	19.16	18.02	17.94	17.94	17.38	16.12	17.94	17.94	16.57
Information ratio	0.09	0.25	0.24	0.24	0.26	0.26	0.24	0.24	0.26
Max relative drawdown (%)	77.38	63.44	57.57	57.57	61.95	55.94	57.57	57.57	56.29
Relative to broad index									
Tracking error (%)	–	6.67 4.44	4.44	6.00	7.17	4.44	4.44	6.01	
Information ratio	–	0.42	0.62	0.62	0.48	0.35	0.62	0.62	0.44
Max relative drawdown (%)	–	46.05	27.67	27.67	42.93	46.68	27.67	27.67	39.72
Distance exposures	13.85	12.72	13.45	13.45	12.82	12.80	13.45	13.45	13.02
ENB in TE	1.84	2.42	2.08	2.08	2.29	2.58	2.08	2.08	2.31

		Max ENB			
Target TE w.r.t. Equ.-broad (%)		None	2	4	6
Annual return (%)		12.65	12.90	12.90	12.64
Volatility (%)		13.19	14.99	14.99	13.61
Sharpe ratio		0.59	0.53	0.53	0.57
Max drawdown (%)		42.99	44.96	44.96	42.81
Relative to liabilities					
Tracking error (%)		16.12	17.94	17.94	16.61
Information ratio		0.26	0.24	0.24	0.25
Max relative drawdown (%)		55.94	57.57	57.57	55.84
Relative to broad index					
Tracking error (%)		7.17	4.44	4.44	6.01
Information ratio		0.35	0.62	0.62	0.41
Max relative drawdown (%)		46.68	27.67	27.67	40.98
Distance exposures		12.80	13.45	13.45	13.04
ENB in TE		2.58	2.08	2.08	2.36

Notes: See the caption of Table 14 for a description of the various strategies. The difference between the two tables is that in this one the equity factor indices are minimum variance portfolios of their respective investment universes, while they are weighted by capitalization in Table 14.

Table 18: Simulation of risk-matched liability-driven investing strategies with out-of-sample equity portfolios made up of minimum variance factor indices with data from September 1971 to December 2016

Performance-seeking portfolio	Equ.-broad	Min distance				Min systematic tracking error			
Target TE w.r.t. Equ.-broad (%)	–	None	2	4	6	None	2	4	6
PSP									
Annual return (%)	10.15	12.97	12.90	12.90	13.05	12.65	12.90	12.90	12.79
Tracking error (%)	19.16	18.02	17.94	17.94	17.38	16.12	17.94	17.94	16.57
Max relative drawdown	77.38	63.44	57.57	57.57	61.95	55.94	57.57	57.57	56.29
LDI strategy with matched tracking error									
Allocation to PSP (%)	40.00	42.57	42.62	42.62	44.01	47.17	42.62	42.62	45.96
Annual return (%)	9.72	10.94	10.88	10.88	10.99	10.84	10.88	10.88	10.88
Volatility (%)	9.46	9.06	9.17	9.17	9.00	8.90	9.17	9.17	8.96
Cumulative relative return (%)	64.80	173.14	165.48	165.48	177.85	161.50	165.48	165.48	166.10
Tracking error (%)	7.51	7.51	7.51	7.51	7.51	7.51	7.51	7.51	7.51
Information ratio	0.16	0.32	0.31	0.31	0.33	0.31	0.31	0.31	0.32
Max relative drawdown (%)	40.79	32.35	28.55	28.55	32.37	30.34	28.55	28.55	29.82
LDI strategy with matched maximum relative drawdown									
Allocation to PSP (%)	40.00	55.96	64.47	64.47	57.84	66.75	64.47	64.47	66.28
Annual return (%)	9.72	11.55	11.81	11.81	11.61	11.61	11.81	11.81	11.72
Volatility (%)	9.46	9.72	10.54	10.54	9.64	9.83	10.54	10.54	10.01
Cumulative relative return (%)	64.80	249.75	288.60	288.60	258.83	258.45	288.60	288.60	273.80
Tracking error (%)	7.51	9.91	11.43	11.43	9.91	10.67	11.43	11.43	10.88
Information ratio	0.16	0.31	0.29	0.29	0.31	0.29	0.29	0.29	0.29
Max relative drawdown (%)	40.79	40.79	40.79	40.79	40.79	40.79	40.79	40.79	40.79

		Max ENB			
Target TE w.r.t. Equ.-broad (%)		None	2	4	6
PSP					
Annual return (%)		12.65	12.90	12.90	12.64
Tracking error (%)		16.12	17.94	17.94	16.61
Max relative drawdown		55.94	57.57	57.57	55.84
LDI strategy with matched tracking error					
Allocation to PSP (%)		47.17	42.62	42.62	45.86
Annual return (%)		10.84	10.88	10.88	10.82
Volatility (%)		8.90	9.17	9.17	8.94
Cumulative relative return (%)		161.50	165.48	165.48	159.23
Tracking error (%)		7.51	7.51	7.51	7.51
Information ratio		0.31	0.31	0.31	0.31
Max relative drawdown (%)		30.34	28.55	28.55	29.48
LDI strategy with matched maximum relative drawdown					
Allocation to PSP (%)		66.75	64.47	64.47	67.01
Annual return (%)		11.61	11.81	11.81	11.65
Volatility (%)		9.83	10.54	10.54	10.04
Cumulative relative return (%)		258.45	288.60	288.60	264.05
Tracking error (%)		10.67	11.43	11.43	11.03
Information ratio		0.29	0.29	0.29	0.28
Max relative drawdown (%)		40.79	40.79	40.79	40.79

Notes: See the caption of Table 16 for a description of the various strategies. The difference between the two tables is that in this one the equity factor indices are minimum variance portfolios of their respective investment universes, while they are weighted by capitalization in Table 16.

4. Improving the Interaction Between Performance-Seeking and Liability-Hedging Portfolios with Factor Investing

In Table 17, we repeat an exercise previously done with cap-weighted indices. Specifically, we construct alternative PSPs with improved liability-hedging properties by combining the six equity factors, plus the broad index unless a cap is set on the tracking error of the portfolio with respect to this index. The weighting methods are the same as in the previous illustrations and respectively involve minimizing the distance between the factor exposures of an equity portfolio and those of liabilities, minimizing the systematic tracking error and maximizing the ENB in tracking error (with respect to liabilities). Some of these portfolios are subject to a maximum tracking error constraint with respect to the broad cap-weighted index, the cap being set to 2%, 4% or 6% per year. The parameters required for the portfolio construction, namely the covariances of constituents or factors, and the betas of constituents with respect to factors, are re-estimated at the start of every quarter from the previous two years of weekly data. This means the portfolios are constructed out of sample, with no look-ahead bias.

The figures in Table 17 are to be compared with those of Table 14, where the factors were cap-weighted. The first observation is that all the alternative portfolios combining the six minimum variance factors outperform those made with the cap-weighted factors, because in this sample, the minimum variance indices outperform their cap-weighted counterparts: this outperformance was already documented in Table 1 over the period from June 1970 to December 2016, and it is also observed in the shorter sample considered here. The second observation is that tracking errors with respect to the broad index are systematically larger than with cap-weighted factors. While the 2% cap was only exceeded by a few basis points with these factors, it is largely breached here, by 244 basis points. In fact, the three weighting methods appear to give exactly the same portfolios at each rebalancing date when they are subject to a 2% tracking error constraint. This observation shows that the 2% constraint can in fact never be satisfied as the minimum tracking error (with respect to the broad index) of a portfolio of the six minimum variance factors is greater than 2%. Wherever the constraint cannot be satisfied, we use the portfolio minimizing the tracking error with respect to the broad index in place of the portfolio that would minimize the distance or the systematic tracking error with respect to liabilities, or maximizing the ENB, and this choice explains why the portfolios subject to the 2% constraint are all equivalent. These higher tracking errors can be attributed to the differences in weighting schemes: minimum variance portfolios are more different from the broad cap-weighted index than cap-weighted portfolios are.

Minimum variance factor indices lead to lower tracking errors with respect to liabilities than cap-weighted indices.

Another difference between these factor indices and the cap-weighted ones lies in their tracking errors with respect to liabilities. With the exception of the minimum systematic tracking error portfolios subject to a 2% constraint, all alternative PSPs have lower tracking errors than when factors are cap-weighted. For the unconstrained minimum systematic tracking error portfolio, which is the one with the lowest ex-post error, the error with cap-weighted constituents was 17.32%, and 16.12% with the minimum variance ones. This result echoes the aforementioned finding of Coqueret, Martellini and Milhau (2017) that minimizing variance within a given stock selection leads to a more liability-friendly equity portfolio. In all cases, except for the unconstrained minimum distance portfolio, a reduction in the maximum relative drawdown is also observed.

In Table 18, we look at LDI strategies to invest in one of the alternative PSPs and in the perfect LHP, with an allocation to the PSP calibrated to achieve the exact same level of relative risk as the reference strategy invested in the broad index and the LHP. The calculation of the risk-matching allocation is done over the period from September 1971 to December 2016, and the relative risk measure is taken to be the tracking error or the maximum relative drawdown. The figures in this table are to be compared with those in Table 16. Because the reduction in tracking error is larger when minimum variance constituents are used, the corresponding increase in the equity allocation with respect to the reference strategy is also larger. For instance, the allocation to the minimum systematic tracking error portfolio that gives the same tracking error as the reference strategy is 47.17%, while it was 43.99% when the portfolio was invested in cap-weighted factors. The higher average returns for PSPs and the larger equity allocations combine to produce higher funding ratios. With the minimum systematic tracking error portfolio, the final funding ratio is 261.50%, assuming a ratio of 100% in September 1971, versus 164.80% when the broad index is employed as the PSP. With cap-weighted factor indices, the same weighting method would have delivered a final funding ratio of 171.29%, so the gain from switching to minimum variance indices is very substantial. Similar observations can be made for the other risk measure, which is the maximum relative drawdown. In all cases, the allocation to the PSP is higher here than it was with cap-weighted indices, so the performance gains are magnified.

4. Improving the Interaction Between Performance-Seeking and Liability-Hedging Portfolios with Factor Investing

5. Conclusion

5. Conclusion

Asset management practices have seen important changes in the first two decades of this century, with the development of factor investing and liability-driven investing. This paper examines how the two approaches can be combined, and more specifically explains how a factor framework can be applied at each stage of the liability-driven investing process.

Interestingly, both factor investing and liability-driven investing are tightly connected to advances in financial theory, which can be positively interpreted with the view that the gap between theory and practice is not as wide as it used to be. Factor investing relies on a large body of empirical research on the determinants of expected returns, research that has so far mostly been conducted in the equity class but has been progressing in other classes. Indeed, the goal of these empirical studies is to identify characteristics, possibly including risk exposures, that have an impact on average returns, so these attributes can then be used as criteria to sort securities into groups with different expected returns, and to construct long-only strategies with expected returns above the market. Provided the sort is based on public information and both the selection of securities and their weighting is done according to systematic and transparent rules, this process leads to passive strategies that do not involve active views and are not equivalent to tracking the performance of a broad cap-weighted index. As such, factor investing has blurred the traditional frontier between passive investment, traditionally regarded as equivalent to replicating a cap-weighted index, and active management, a term that used to apply to any investment process involving deviations from the broad market index. Similarly, liability-driven investing also has a theoretical justification, based on the fund separation theorems arising in portfolio optimization problems. These results tell us that instead of trying to achieve performance and good hedging of their liabilities with a single portfolio, investors are better off splitting their assets between two building blocks with well-defined roles, namely a performance-seeking portfolio and a liability-hedging portfolio. This paradigm is thus challenging the concept of policy portfolio, which has long been dominant in institutional money management.

While factor investing and liability-driven investing relate to two separate strands of the academic literature, a case can be made for combining these approaches. Indeed, each of the three steps of a liability-driven investing process, namely the construction of a performance-seeking portfolio, that of a liability-hedging portfolio and an allocation to these building blocks, can be addressed by adopting a factor perspective. Our paper introduces a framework for doing this. It begins with a review of the possible definitions of a "factor" in practice: some of them focus on the risk of assets and strategies, and others on their expected profitability; some focus on explaining the cross sectional differences between securities while others are more concerned with explaining the time series changes in their risk and return properties. In the equity class, a number of candidate factors have been proposed, but only a handful have withstood robustness checks and are economically plausible. We show how they can be combined to deliver substantial improvements over the risk-return characteristics of the broad cap-weighted index. At the hedging stage, the factor lens is also useful since it involves the identification of the risk factors that have an impact on the present value of liabilities, and calculating the exposures to these factors of liabilities and the securities used as replicating instruments. We illustrate this approach by considering the replication of a fictitious security called the "retirement bond" designed to provide retirees with stable replacement income for a period equal to their life expectancy without the need for annuities. This problem is formally identical to the problem that would be faced by a defined-benefit pension fund that is committed to making fixed payments for a clearly defined period.

The last part of this paper considers the allocation stage. First, it introduces a factor model to analyze the overlap between the risk exposures of assets and liabilities and uses it to decompose the risk of equity-bond strategies, both in absolute terms and relative to liabilities. This risk decomposition is based on the same formulas as the decomposition of a portfolio's volatility across its constituents, which is at the heart of the concept of risk parity in asset allocation, but it extends them to account for the presence of residual idiosyncratic risk in the portfolio. This extension is necessary because even a factor model with good explanatory power does not account for 100% of the time variation in the returns on assets and liabilities. The factor model can also serve as a tool to measure the degree of similarity between a portfolio and liabilities, by calculating the distance between their respective factor exposures. In particular, one can search for a performance-seeking portfolio with closer factor exposures to those of liabilities than the broad cap-weighted index. While it seems to violate the fund separation principle, the search for such equity portfolios is in fact justified by the observation that investors' welfare depends both on the hedging properties of the performance-seeking portfolio and – as expected – on its performance properties and the hedging qualities of the liability-hedging portfolio.

Our work could be extended in several directions. First, the replication problem could be studied in the context of inflation-linked replacement income, so as to introduce other risk factors in addition to the level of interest rates. Martellini, Milhau and Tarelli (2014b) have done work in this direction by considering the replication of an index of inflation-linked bonds, and a critical step is the estimation of the exposures to real rates and to expected inflation. It remains to be seen how these techniques would perform for the replication of a finite-maturity bond with a long duration, like a retirement bond that pays out replacement income indexed on realized inflation. Given that the discounted value of liabilities has non-linear exposures to the interest rate risk factors, it may also be useful to test if matching derivatives of a higher order than 1 leads to a more accurate replication. On the performance-seeking portfolio side, our paper focuses on equities, since it is the asset class in which the factor approach has become the most widely accepted and in which products providing exposures to selected factors are now available. However, a series of recent papers by Maeso, Martellini and Rebonato (2019a, 2019b, 2019c) suggest that certain factors are also present in bond markets, including the level factor in particular.

5. Conclusion

To fully take advantage of this factor, the authors point that it is important to take time variation in exposures into account. Incorporating these techniques into the design of a multi-asset performance-seeking portfolio would be another practically relevant area for further research.

6. Appendices

6. Appendices

A. Defining Factors

The notion of factor in portfolio construction and asset allocation is polysemic, so the first step towards the setup of a factor investing framework is to clarify the various definitions. In this section, we follow the classification proposed by Martellini and Milhau (2018c), who identify four categories of factors. A factor can be an asset pricing factor, in the sense of asset pricing theory, and, as such, it explains or helps explain the differences in expected returns across assets. Such factors can be exploited by following systematic investment rules that give rise to positive expected returns. This leads to a broader definition often adopted in investment practice, by which a factor is a profitable passive strategy. Third, a factor can refer to a common source of risk in the returns of a large set of securities. The last notion of a factor is that of a state variable that has some forecasting power for returns, and, as such, can also be used to construct profitable strategies.

A.1 Definition 1: Factors as the Sources of Risk Premia

The most academic definition of a factor comes from asset pricing theory. An asset pricing factor is some financial or economic variable that helps explain the differences in expected returns across assets. Factor models postulate that the risk premium on a security *i*, which is its expected return in excess of the risk-free interest rate, is the sum of its exposures to a set of *K* factors, weighted by the premia of the factors. Mathematically, the expected excess return is given by

$$\mu_i = \beta_{i1}\Lambda_1 + \cdots + \beta_{iK}\Lambda_K, \tag{A.1}$$

where β_{ik} is the beta with respect to factor k and Λ_k is the factor premium. The factor premium is defined as the expected excess return of a portfolio with a unit beta with respect to this factor, and zero betas with respect to the others. It can also be interpreted as the change in expected return per unit of change in the beta. With this definition, a factor premium can be positive or negative. It is positive if investors require an additional reward (in the form of higher expected return) for bearing additional exposure to the factor, and negative otherwise.

In a single-factor model, the exposures to a unique factor explain the risk premia of all securities. Many of the standard asset pricing models fall into this category. In consumption-based models with time-separable preferences, a representative agent maximizes the sum of expected utilities from consumption at different future dates. These models are analyzed in detail in the reference textbook by Cochrane (2005), and an important conclusion is that the marginal utility from current consumption is a pricing factor. In other words, the expected excess return can be written as in Equation (A.1), with K equal to 1 and the beta with respect to marginal utility in the right-hand side. The models also establish that the single factor premium is negative, so a positive beta with respect to marginal utility results in negative expected return. This result intuitively makes sense because an asset that covaries positively with marginal utility tends to pay off more in scenarios where marginal utility is higher, i.e. when one additional dollar is most needed. Thus, this asset is attractive for the purpose of hedging against the risk of being far from satiation, where satiation corresponding to a state where the marginal utility of increasing consumption is low. Having this property, the asset does not deserve a positive premium. On the other hand, assets that pay off more when marginal utility is low are less appealing for hedging purposes, so they must promise a higher reward.

Theoretical Pricing Factors in Models Based on Individual Optimization
The Capital Asset Pricing Model (CAPM) of Sharpe (1964) and Lintner (1965) and the Intertemporal CAPM (ICAPM) of Merton (1973) derive representations of expected returns in the form of Equation (A.1) by assuming that each market participant maximizes some utility function and that the demand for securities equals supply. Assumptions on individual behavior differ between the two models and lead to a single-factor representation in the CAPM, and to a multi-factor one in the ICAPM.

The CAPM assumes market equilibrium and homogenous expectations and horizons across investors. Provided investors do not rebalance their portfolios and seek to maximize the expected return penalized by the variance of their portfolio, the expected excess return of any risky asset is proportional to its beta with respect to the market portfolio, defined as the portfolio in which securities are held in proportion to their market capitalization. Formally, it holds that

$$\mu_i = \mu_{mkt}\beta_i, \tag{A.2}$$

where μ_{mkt} is the expected excess return of the market portfolio. So the return of the market portfolio acts as a pricing factor. This result can be regarded as a special case of a more general finding that the return of the maximum Sharpe ratio portfolio is a pricing factor (see Appendix B for a proof).

Merton (1973) relaxes the CAPM's assumptions that investors have identical horizons and that they do not rebalance their portfolios, and postulates that each of them maximizes the expected utility from future consumption, somewhat in the manner of consumption-based models, and also from any remaining wealth after the end of the consumption cycle. The model is cast in continuous time, so the investment period is divided into infinitesimal trading periods. The article shows that the conditional expected excess return of an asset at the beginning of each short period is given in the form of equation (A.1), and the factors in the right-hand side are the "state variables" that determine investment opportunities at this time. Investment opportunities are described by the risk-free interest rate, the risk premia on risky assets, their volatilities and their correlations. It should be noted that in this model, Equation (A.1) holds conditionally, i.e. with conditional expected returns in the left-hand side and conditional betas and factor premia in the right-hand side.

The intuition is that changes in state variables can be good or bad news for investors, depending on how they impact their marginal utility of wealth. Suppose for the purposes of illustration that the level of interest rates is one of these state variables, and that a decrease in the level is associated with bad economic news that make the marginal utility of consumption rise.[102] Bond prices increase when the level decreases, so bonds are useful for investors to hedge against an unfavorable shock on interest rates. As a result, investors purchase bonds for hedging purposes, so the model predicts that they have a lower expected return than that implied by their covariance with the market portfolio. Conversely, an asset that investors do not want to hold for hedging purposes would deserve a larger premium than that implied by their market exposure.

Breeden (1979) shows that Merton's multi-factor model can in fact be rewritten as a single-factor model in which the factor is the marginal utility of consumption. This representation is compatible with the static CAPM in the limit case where there is a single period. Indeed, if there is only one period, the final wealth is entirely consumed, so future consumption equals future wealth. Breeden's model then predicts that the expected excess return of an asset is determined by its covariance with future wealth. Under the market clearing condition, the return on wealth equals the return on the market portfolio, so the expected excess return is eventually determined by the covariance with the market portfolio, and the CAPM's prediction is recovered.

Campbell (1993) develops an alternative version of the ICAPM, by altering some of Merton's assumptions. In particular, the form of the utility function is different: Merton's investors maximize the sum of expected utilities from consumption at all future dates, while Campbell assumes a "recursive" utility function, in which welfare is a more complex function of future expected utilities. The pricing equation that he derives says that the expected excess return of an asset depends on its covariance with the market and on its covariance with changes in the expectations of future market returns. Thus, an asset that grows in value when expectations are revised downwards, i.e. when investment opportunities become less promising, deserves a lower premium than what is implied by the market exposure, since this asset is helpful for hedging against unfavorable changes in the opportunity set.

A summarized formal presentation of these models (CAPM, ICAPMs and consumption-based models) can be found in Section 2 of Martellini and Milhau (2015).

Theoretical Pricing Factors from Statistical Analysis

The other big category of multi-factor models is that of the Arbitrage Pricing Theory of Ross (1976). Unlike the CAPM and the ICAPM, the APT does not assume some form of market equilibrium and relies solely on a statistical decomposition of security returns. Each return is expressed as the sum of a "systematic part" driven by exposures to a limited number of factors, and an "idiosyncratic" part not explained by factor exposures. The restrictions imposed by the model are of a statistical nature: idiosyncratic returns must be uncorrelated across assets, and their variance must be "not too large".

102 - Formally, this condition means that the second-order derivative of the indirect utility function with respect to wealth and the level of rates is negative.

Then, as the variance of idiosyncratic returns decreases to zero, a factor representation of the form (A.1) is obtained.

From Asset Pricing Factors to Profitable Strategies
Knowing a set of pricing factors is interesting from a fundamental standpoint when it comes to understanding why expected returns differ across assets, but it is also useful for investment purposes. Suppose that two portfolios have roughly the same betas with respect to factors 2, 3, ..., K and have different betas with respect to factor 1. Then, based on Equation (A.1), the spread of expected returns is approximately

$$\mu_{P1} - \mu_{P2} \approx [\beta_{P1} - \beta_{P2}]\Lambda_1.$$

(Subscripts P denote portfolios, whose expected returns and betas are the weighted sums of the expected returns and the betas of the constituents.) So, if the first factor has a positive premium, it is the portfolio more exposed to this factor that has the highest expected return. Thus, a zero dollar strategy with positive expected return can be designed by taking a long position in the more exposed portfolio and a short position in the other. Things work the other way round if the factor premium is negative.

This procedure also works for investors who wish or have to be long only. For them, the benchmark to beat is some broad index, often thought of as a proxy for the market portfolio, so the goal is to find a portfolio with factor exposures that imply higher expected returns than the benchmark. In the CAPM world, the only way to accomplish this is by having a market beta greater than unity, since the market portfolio has a beta of 1. In a multi-factor setting, there are as many degrees of freedom as factors.

Designing long-short strategies with positive expected return by sorting portfolios on betas is easier than doing so by sorting directly on expected returns to the extent that betas are easier to estimate than expected returns. This is most often the case, as sample means are very imprecise estimates of expected returns, which cannot be improved by increasing the sampling frequency of observations (see Appendix A of Merton (1980)). In contrast, risk parameter estimates get more accurate as data frequency increases: Merton (1980) provides a mathematical justification, and Jagannathan and Ma (2003) illustrate this fact in the context of covariance matrix estimation. That said, estimation of betas is not always a straightforward task: beta estimates are more accurate at the level of a portfolio than at the level of individual securities, which is one of the motivations behind the classical Fama-MacBeth procedure applied in empirical asset pricing (Fama and MacBeth 1973), and betas vary over time, which complicates statistical estimation.

A.2 Definition 2: Factors as Profitable Passive Strategies

The second definition of factors is often used by practitioners and is inspired by the observation in the previous section that asset pricing factors can be used to design strategies that are profitable on average. With this definition, a factor is seen as a passive long-short strategy with a positive expected return, or a passive long-only strategy with greater expected return than a given benchmark, generally taken to represent the asset class considered. For US equities, a standard benchmark is the S&P 500 index, while it could be the Euro Stoxx 50 for European equities, and so on.

An important element of this definition is that the strategy must be passive in the sense that it should not attempt to take advantage of pricing anomalies discovered by the use of non-public information or proprietary techniques. On the other hand, it can be rebalanced, although the act of rebalancing is sometimes seen as an "active" decision.

This restriction is satisfied by the aforementioned strategies that go long a portfolio more exposed to a positively rewarded factor, and short a portfolio less exposed to that same factor. Indeed, the sorting criterion is the beta with respect to the factor, which can be estimated from public data made up of portfolio returns and factor values. Where Definition 2 is broader than Definition 1 is in the choice of the sorting criterion. In Definition 1, differences in expected returns across securities are entirely explained by the differences in their betas with respect to so-called "pricing factors", but Definition 2 acknowledges that some expected return spreads can be explained other than by invoking differences in risk. Formally, Definition 2 can be stated by writing that expected returns on securities are given by a formula of the same form as Equation (A.1), but with the betas re-interpreted as more general "characteristics" known to be systematically associated with higher expected returns. Examples of such characteristics include betas with respect to a factor, but also the firm capitalization, the ratio of the book value to the market value of equities, the price-earnings ratio or the dividend yield, among many other attributes known to produce dispersion in returns. For instance, Daniel and Titman (1997) argue that the historical outperformance of small stocks over large stocks, and value stocks over growth stocks, is not the reward for bearing larger undiversifiable risk, and that it is the characteristics, not the risk exposures, that are for some reason associated with expected returns.

A possible explanation that they suggest is the existence of behavioral and cognitive biases. More generally, behavioral explanations are the other main category of economic justifications for observed differences in average returns, beyond differences in risk. They put forward systematic pricing errors made by investors, which lead them to undervalue some securities and to overbid others and which arise because of behavioral and cognitive biases such as under- or over-reaction, the tendency to extrapolate and the inclination to follow the "herd". Certain features of financial markets may also help explain why some securities perform better than others: examples include the gradual diffusion of information and the unequal attention given by the media to different securities (see Hong and

Stein's (1999) explanation of momentum and reversal in equity markets, or agency effects). For these anomalies to subsist in spite of investors' attempts to exploit them, there must be "limits to arbitrage", which prevent traders from capturing them completely.

A.3 Definition 3: Factors as Common Sources of Risk

Factors can also be understood as common sources of risk that affect a large set of constituents. As a starting point, the return of each individual security or asset class is expressed as the sum of a systematic part, which is a sum of contributions from different factors, and an idiosyncratic part, which represents the asset's specific behavior. Mathematically, this decomposition is written as

$$\tilde{r}_i = c_i + \beta_{i1} F_1 + \cdots + \beta_{iK} F_K + \varepsilon_i, \tag{A.3}$$

where $\tilde{r}_i$ denotes the excess return of an asset i over the risk-free interest rate, F_1, ..., F_K are the factors and ε_i is the idiosyncratic return, uncorrelated from all factors. The volatility of the residual is also known as the idiosyncratic volatility. The ratio of idiosyncratic variance to the total asset variance is the percentage of variance explained by the factors, or the R-square of a regression of asset returns on factors. Adding risk factors mechanically increases the explanatory power of the model, and Appendix D derives a mathematical expression for the gain in R-square by introducing new regressors: when a new factor is added, the change in R-square is the squared correlation between r_i and the residuals of the new factor with respect to those already present in the model. So, informally, the R-square does not increase if the correlation between the new factor and the return to explain is entirely explained by the correlation of this factor and the previous ones.

Using Equation (A.3), each asset variance can be decomposed as the sum of a "systematic variance", which is explained by the factor exposures, and an idiosyncratic variance. Formally,

$$\sigma_i^2 = \sum_{k,l=1}^{K} \beta_{ik}\beta_{jl}\mathbb{C}\mathbf{ov}[F_k, F_l] + \sigma_{\varepsilon,i}^2,$$

where $\sigma_{\varepsilon,i}$ is the idiosyncratic volatility. Similarly, cross-asset covariances can be decomposed into a sum of contributions arising from factor exposures, and the covariance between the idiosyncratic returns.

Examples of Explicit Risk Factors

A standard example of a common risk factor is the "market factor" in equities listed on a given stock exchange. It affects most returns in the same direction since most stocks have a positive beta, although some stocks are more exposed than others, which is expressed by saying that they are more "defensive". In a fixed-income universe consisting of bonds from the same issuer, like the set of all bonds issued by a sovereign state, a common risk factor is also present: it is the level of interest rates, which has the same directional effect on all bonds since all bond prices decrease when interest rates rise. While a rise in interest rates implies a drop in all bond prices, some bonds are more impacted by a given interest rate movement: indeed, the interest rate exposure depends on the bond duration, so long-duration bonds are more exposed than short-duration ones.[103]

If the securities at hand have common characteristics that imply some degree of commonality in returns, one can form explicit risk factors by taking some weighted average of their returns. For instance, the market factor in stocks of a given country is often taken to be a cap-weighted index of those stocks that make up for the largest part of the aggregate capitalization. It is also likely that all securities of a given industry are affected by the same economic news, so industry factors defined as the returns of cap-weighted or equally-weighted portfolios of all stocks within an industry are candidate common risk factors.

A standard method to estimate the betas with respect to explicit factors is to run a linear regression of asset returns on factors, but this may be inappropriate in some cases, e.g. if returns are not stationary. This situation occurs when assets are finite-maturity bonds, because returns converge to zero as the bond approaches maturity. Another approach must therefore be taken by establishing a functional, rather than statistical, link between the bond return and the factors. More generally, it can also be employed whenever the prices of assets to explain are naturally expressed as functions of risk factors, so that the knowledge of the function that maps factors into prices can be exploited to find the exposures of the return with respect to the factors. Section 3.2.3 applies this method to calculate the exposures of bond returns with respect to the three standard fixed-income factors, namely level, slope and concavity.

Implicit Risk Factors with Unknown Exposures

The risk factors constructed by examining the common characteristics of securities do not necessarily fully explain returns, as some other risk factors may be associated with less evident commonalities: a non-zero residual therefore remains in the right-hand side of Equation (A.3). Thus, statistical methods to extract implicit factors can be very useful to identify a complete set of risk factors, i.e. which completely explain the observed returns. These methods can also be especially helpful when common characteristics are less obvious, as in multi-class universes, where notions of country and industry are not always applicable.

The first of these methods is principal component analysis, which aims to extract uncorrelated factors expressed as linear combinations of returns. The factors are sorted by decreasing order of volatility, and this order also reflects the contribution to explaining common variance of returns. This procedure is routinely applied, but the principal factors are not always easy to interpret economically. The second method, known as minimum linear torsion, was introduced by Meucci, Santangelo and Deguest (2015), and also aims to extract uncorrelated factors defined as linear combinations of returns, but

103 - More generally, the existence of a "market" or "level" factor, defined as a factor with respect to which all securities of a given universe are positively exposed, can be mathematically established with the Perron-Frobenius theorem: this result implies that if a covariance matrix has only positive entries (i.e. if all securities are positively correlated), then there exists an eigenvector with positive entries associated with the largest eigenvalue. The first principal factor in a principal component analysis is precisely defined as the factor associated with the largest eigenvalue, and the elements of the first eigenvector are the loadings of the securities on this factor. The first factor is a linear combination of returns with coefficients given by the elements of the first eigenvector. Thus, all securities load positively on the first principal factor, and the factor is itself a combination with positive coefficients of all returns. In the case of bonds, the first principal factor can be interpreted as minus the level of interest rates.

among all factor sets that comply with this rule, it selects the one that minimizes the distance with respect to the original returns. In other words, it finds the least distorting linear transformation of returns into factors. In both procedures, asset exposures to factors are estimated jointly with factor values. Both methods are illustrated in a multi-class universe by Martellini and Milhau (2018a), and in the equity class in Section 2.3 below.

Implicit Risk Factors with Known Exposures

A second category of implicit risk factors is that used in risk models like those of Barra. In these models, exposures are treated as known quantities and factor returns are estimated in each period by regressing the cross section of asset returns on the exposures.[104] Some of these exposures are associated with selected attributes of stocks, including accounting attributes and market data in particular. In this case, attributes are standardized by subtracting the cross-sectional mean and dividing by the standard deviation, so as to align them on a common scale. Industry and country factors are also present in these models. The simplest approach for calculating the industry exposures of each stock is to assign a 0-1 score to each industry, depending on the industry classification of the company. Another method, which recognizes that firms can allocate part of their activities to different industries, is to set each industry exposure equal to the percentage of the firm's activity (e.g. measured by earnings) in this area.

A model similar to Barra's, in which the volatility of a stock is expressed as a function of its characteristics, is used by Haugen and Baker (1996). Specifically, the model includes variables related to risk or liquidity, accounting ratios comparing the stock price to fundamental variables, various measures of the firm's profitability, and "technical factors" that summarize the stock's price history. In total, the model has 41 variables.[105] Monthly factor values are estimated by running cross-sectional regressions of stock returns on exposures, taken equal to observable characteristics with some normalization to smooth out differences in scales. The authors use the model to estimate expected returns, expressed as linear combinations of expected factor values weighted by characteristics.

Use of Risk Factors

An important application of risk factors is to the estimation of risk parameters, including the covariance matrix, in large-scale universes. Barra risk models are precisely employed for this purpose. To see why the use of a factor model alleviates the estimation burden with respect to the case where no factor structure is assumed, let us count the number of independent parameters in both cases. For N assets, there are $N[N-1]/2$ pairwise correlations and N volatilities to estimate, hence a total of $N[N+1]/2$ parameters. This number is a quadratic function of N, so it is approximately multiplied by 4 when the universe size is doubled. With a factor model as in Equation (A.3), there are NK betas to estimate, plus the $K[K+1]/2$ correlations and volatilities of factors, and the covariances of residuals. If residuals are assumed to be uncorrelated across assets, there are N idiosyncratic volatilities to estimate, so the total

104 - See p. 27 of MSCI (2007).
105 - The complete list of the 41 characteristics is given in the paper's Appendix.

number of parameters is a linear function of N. For large universes, the factor model considerably reduces the number of parameters to estimate.

Tackling the dimensionality curse was the motivation of Sharpe (1963) when he introduced the single-factor model. This model is described by Equation (A.3) and has a unique factor, which Sharpe takes to be a "stock market index" of the NYSE. It is used to estimate the expected returns and the covariances of N stocks. Assuming that idiosyncratic returns are uncorrelated across securities, the total number of parameters to estimate is given by

$$\underbrace{N}_{\text{betas}} + \underbrace{N}_{\substack{\text{idiosyncratic}\\\text{volatilities}}} + \underbrace{N}_{\text{intercepts}} + \underbrace{1}_{\substack{\text{factor}\\\text{premium}}} + \underbrace{1}_{\substack{\text{factor}\\\text{volatility}}} = 3N + 2,$$

while the number of parameters to estimate if returns are not generated by a factor model is

$$\underbrace{\frac{N[N-1]}{2}}_{\text{correlations}} + \underbrace{N}_{\text{volatilities}} + \underbrace{N}_{\substack{\text{expected}\\\text{returns}}} = \frac{N[N+3]}{2}.$$

For $N = 500$ securities, the number of required estimates decreases from 125,750 to 1,502.

The identification of common sources of risk is also a key step in the construction of replicating portfolios for a benchmark by using factor-matching techniques (see Section 3 for an example). Indeed, provided the factors have sufficiently large explanatory power – meaning that the residuals are sufficiently small –, constructing a portfolio with the same factor exposures as the benchmark will result in low tracking error, defined as the volatility of the return spread between the portfolio and the benchmark.

A.4 Definition 4: Factors as Time-Series Predictors of Returns

A fourth understanding of a factor is as a state variable that contributes to the description of market conditions and investment opportunities. This definition focuses on the time series variation of returns, as opposed to the cross-sectional variation in average returns.

Time-Varying Volatilities and Expected Returns

As pointed out by Merton (1973), interest rates are a straightforward example of time-varying investment opportunities. For volatilities and expected returns, evidence of time variation is more indirect, as these quantities are not readily observable.

Time variation in volatilities is revealed by the volatility clustering phenomenon: Mandelbrot (1967) noted that there are periods in which large positive or negative changes accumulate, while at other times returns are smaller in magnitude. In statistical terms, squared returns display positive autocorrelation. This behavior cannot be replicated if returns are independently and identically distributed, but it can be captured by assuming that volatility varies over time and is predictable with past returns, as in the ARCH model of Engle (1982), or with a mixture of past returns and past volatilities, as in the GARCH of Bollerslev (1986).

For expected returns, evidence of time variation coincides with evidence of predictability. Part of the excess returns on long-term bonds can be predicted with variables related to the slope of the term structure. Fama and French (1989) use the term spread together with the default spread or the dividend yield of stocks and find that the percentage of variance explained depends on the prediction horizon, increasing up to one year, then decreasing and finally increasing again from three to four years.[106] Fama (1986) and Fama and Bliss (1987) regress future bond returns on the current forward-spot spread. In these predictive regressions, slopes are close to 1 and significantly different from zero, meaning that the spread does have predictive power, and the R-square ranges from approximately 2% to 15%,[107] showing that a big fraction of the time variation of return is still not captured. In these regressions, the maturity of the forward rate used in the right-hand side is matched with the bond maturity to predict, but Cochrane and Piazzesi (2005) show that a single factor constructed from the forward rates of various maturities, which they label "the tent factor", explains an even greater fraction of the variance.[108] All these results invalidate the expectation hypothesis, which says that the term premium is zero, or equivalently that long-term bonds and short-term bonds have the same expected return. Cieslak and Povala (2015) introduce a "cycle factor" that aggregates the transitory components in the yields of various maturities: for each yield, the transitory part is measured by subtracting the part of the yield that is explained by "trend inflation", calculated as a moving average of past inflation rates.

On the equity side, dividend-price ratios are standard predictors, but their predictive power increases markedly with the horizon over which returns are measured:[109] according to the results of Fama and French (1988),[110] the percentage of variance explained by the dividend yield is approximately 1% to 3% for monthly and quarterly returns, but it rises to much higher levels, from 13% up to 49%, when returns are measured over four years. A similar pattern is found by Fama and French (1989) for the dividend-price ratio, and it is different from the pattern of R-squares in regressions of investment-grade bond returns: bond R-squares are lower at the two- and three-year horizons than at the one-year horizon. Goyal and Welch (2003) emphasize that predictability is weak at the annual horizon, and they argue that the dividend-price ratio has no forecasting ability for market returns under five years. Another salient fact from predictability studies using dividend-price ratios is a lack of robustness with respect to the sample period, as pointed out by Welch and Goyal (2008) in a review of prediction models.

106 - See their Table 2.
107 - See Table 2 in Fama (1986) and Table 1 in Fama and Bliss (1987).
108 - Compare the R-squares in their Table 1 with the R-squares of Fama-Bliss regressions in Table 2.
109 - Here, the dividend-price ratio is defined as the sum of past dividends divided by the current price, and the dividend yield as the sum of past dividends over the price at the start of the period. However, there does not seem to be established and consistent terminology across studies: the "dividend yield" in Fama and French (1989) is calculated by dividing dividends by the current price, while Fama and French (1988) calculate it with reference to the past price.
110 - See their Table 3.

But all evidence for predictability does not rely on dividend-price ratios. From a theoretical standpoint, the use of these ratios can be motivated through accounting equations expressing returns as functions of capital gains and dividends, as done in Campbell and Shiller (1988). A similar reasoning but based on a different accounting identity, namely the intertemporal budget constraint of a representative agent, leads Lettau and Ludvigson (2001a) to another predictor, which is the aggregate consumption-wealth ratio. The authors show that an empirically measurable proxy for this quantity is more significantly related to future quarterly returns than the dividend-price ratio. Menzly, Santos and Veronesi (2004) argue that the limited forecasting power of the dividend yield is due to the blurring effects of time variation in risk preferences due to habit formation and time variation in expected dividend growth, and they suggest introducing the consumption-price ratio along with the dividend-price ratio in predictive regressions of industry portfolios, showing that for many of these portfolios, the latter variable is significant and its inclusion leads to a substantial gain in R-square with respect to the regression using only the dividend yield.[111]

Implications for Portfolio Optimization

Time variation in investment opportunities has implications for portfolio choice, which were first explored by Merton (1973), and then in a number of subsequent papers focusing on the effects of stochastic risk premia, volatilities and correlations. Merton (1992) extends the 1973 paper by providing a general mathematical approach to portfolio optimization in the presence of an arbitrary number of stochastic state variables. Changing market conditions have an impact on optimal portfolios for agents who invest over multiple periods and want to maximize their expected welfare at the term of these periods. Indeed, the optimal investment decision at each point in time for them is obtained by maximizing the expected welfare of the next rebalancing date, not the expected utility of wealth at the next date. But welfare at the next date depends on business conditions prevailing at that date and also on the amount of wealth, because good investment opportunities – e.g. high expected returns and low volatilities – result in higher welfare. So the investor has to anticipate the impact that unexpected changes in the state of the economy can have on their future welfare by purchasing assets that deliver higher returns when business conditions deteriorate than when they improve. Thus, fixed-rate bonds are useful to hedge against unfavorable changes in interest rates since their value increases when rates decrease, and equities can be used as a hedge against the risk of a fall in their own risk premium since the risk premium is countercyclical and tends to be lower in bull than in bear markets.

From State Variables to Pricing Factors

The Intertemporal Capital Asset Pricing Model (ICAPM) of Merton (1973) connects the time series and the cross section perspectives by showing that the state variables that describe investment opportunities at each point in time act as pricing factors, in addition to the market portfolio of the static CAPM. Specifically, the conditional expected excess return on an asset is a linear combination

111 - See their Table 5.

of its betas with respect to the market portfolio and the state variables, with coefficients equal to the premia of these factors. A more detailed description of the ICAPM and the versions of Breeden (1979) and Campbell (1993) is given in Appendix A.1, and a formal presentation can be found in Section 2 of Martellini and Milhau (2015).

Put briefly, the pricing implications of the ICAPM are different from those of the CAPM – which assigns no role to state variables describing business conditions – because the former model predicts that agents want to hedge against unfavorable changes in investment opportunities: the hedging motive impacts the demand for each security, so that in equilibrium each expected return depends on the security's ability to provide a hedge against such changes.

B. Efficient Portfolio Return as Pricing Factor

Consider a universe made up of N risky securities and one risk-free asset. The N risky excess returns over the risk-free rate are stacked in a random vector $\mathbf{R}$, with expectation $\boldsymbol{\mu}$ and covariance matrix Σ. Assets are assumed to be non-redundant, so the matrix Σ is non-singular.

A portfolio is described by the vector $\mathbf{w}$ of percentage weights allocated to the risky assets, and one minus the sum of the elements of $\mathbf{w}$ is the weight of the risk-free asset. Cauchy-Schwarz inequality dictates that

$$\mathbf{w}'\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} \leq \sqrt{\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w}} \times \sqrt{\boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}},$$

such that

$$\frac{\mathbf{w}'\boldsymbol{\mu}}{\sqrt{\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w}}} \leq \sqrt{\boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}.$$

Both inequalities become equalities when $\mathbf{w}$ is proportional to $\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$, so the maximum Sharpe ratio (MSR) portfolio of risky assets is

$$\mathbf{w}_{MSR} = x\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu},$$

where the constant x is adjusted in such a way that the weights of the MSR portfolio in risky assets add up to 1.

The expected return and the variance of the MSR portfolio are given by

$$\mu_{MSR} = \mathbf{w}'_{MSR}\boldsymbol{\mu} = x\boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu},$$
$$\sigma^2_{MSR} = \mathbf{w}'_{MSR}\boldsymbol{\Sigma}\mathbf{w}_{MSR} = x^2\boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu},$$

so the beta of a security with respect to the MSR portfolio is

$$\beta_i = \frac{\mathbb{Cov}[R_{MSR}, R_i]}{\sigma^2_{MSR}} = \frac{\mathbf{w}'_{MSR}\boldsymbol{\Sigma}\mathbf{e}_i}{x^2\boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}},$$

where $\mathbf{e}_i$ is the vector with i^{th} element equal to 1, and the others equal zero.

It follows that

$$\begin{aligned}\beta_i \mu_{MSR} &= \frac{1}{x}\mathbf{w}'_{MSR}\mathbf{\Sigma}\mathbf{e}_i \\ &= \boldsymbol{\mu}'\mathbf{e}_i \\ &= \mu_i.\end{aligned}$$

Therefore, the return of the MSR portfolio is a pricing factor.

C. Mathematics of Principal Component Analysis

Consider a rectangular matrix of asset returns $\mathbf{R}$ with T rows and N columns, where T is the number of dates and N is the universe size. Let $\mathbf{1}$ be a $T \times 1$ vector filled with ones. The sample means of returns are given by

$$\mathbf{m} = \frac{1}{T}\mathbf{R}'\mathbf{1},$$

and the sample covariance matrix is

$$\mathbf{\Sigma} = \frac{1}{T-1}[\mathbf{R} - \mathbf{1}\mathbf{m}']'[\mathbf{R} - \mathbf{1}\mathbf{m}'].$$

When the columns of $\mathbf{R}$ are independent and sampled from the same distribution, Σ is an unbiased estimate of the covariance matrix of this distribution.

Principal component analysis returns N principal factors by diagonalizing the covariance matrix. The diagonal form is

$$\mathbf{\Sigma} = \mathbf{P}\mathbf{\Sigma}_F\mathbf{P}',$$

where $\mathbf{P}$ is an orthogonal matrix and Σ_F is a diagonal matrix, with diagonal coefficients sorted in decreasing order. The matrix of factor values is defined as

$$\mathbf{F} = [\mathbf{R} - \mathbf{1}\mathbf{m}']\mathbf{P},$$

so the factors have zero mean and their covariance matrix is Σ_F.

Returns can be recovered from factors through

$$\mathbf{R} = \mathbf{1}\mathbf{m}' + \mathbf{F}\mathbf{P}',$$

so the N betas of the N returns with respect to factor k are the elements of the k^{th} row of $\mathbf{P}'$, or equivalently the elements of the k^{th} column of $\mathbf{P}$.

D. Gain in R-Square from Adding Regressors

Consider a random variable y, which is the dependent variable. The explanatory variables are stacked in two random vectors $\mathbf{x}_1$ and $\mathbf{x}_2$, with respective lengths K_1 and K_2. To simplify notations, we assume that all variables have zero mean, which has no impact on the calculation of R-squares.

The first regression has only $\mathbf{x}_1$ in the right-hand side, and the second uses both $\mathbf{x}_1$ and $\mathbf{x}_2$. Betas and residuals are defined in the following equations:

$$y = \boldsymbol{\gamma}_1' \mathbf{x}_1 + \varepsilon_1, \tag{D.1}$$

$$y = \boldsymbol{\beta}_1' \mathbf{x}_1 + \boldsymbol{\beta}_2' \mathbf{x}_2 + \varepsilon_2. \tag{D.2}$$

Let $\mathbf{x}$ denote the full vector of regressors and β the concatenation of β_1 and β_2. The betas in the second equation are given by

$$\boldsymbol{\beta} = \mathbb{V}[\mathbf{x}]^{-1} \mathbb{C}\mathrm{ov}[\mathbf{x}, y],$$

and the covariance matrix of $\mathbf{x}$ can be written as

$$\mathbb{V}[\mathbf{x}] = \begin{bmatrix} \boldsymbol{\Sigma}_1 & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{12}' & \boldsymbol{\Sigma}_2 \end{bmatrix},$$

where Σ_i is the covariance matrix of $\mathbf{x}_i$, and Σ_{12} is the K_1 x K_2 covariance matrix of $\mathbf{x}_1$ and $\mathbf{x}_2$. The vector of covariances between explanatory variables $\mathbf{x}$ and y can be decomposed as

$$\mathbb{C}\mathrm{ov}[\mathbf{x}, y] = \begin{bmatrix} \mathbf{C}_1 \\ \mathbf{C}_2 \end{bmatrix},$$

where $\mathbf{C}_1$ is the covariance vector of $\mathbf{x}_1$ and y, and $\mathbf{C}_2$ the covariance vector between $\mathbf{x}_2$ and y.

The blockwise inversion matrix formula gives

$$\mathbb{V}[\mathbf{x}]^{-1} = \begin{bmatrix} \boldsymbol{\Sigma}_1^{-1} + \boldsymbol{\Sigma}_1^{-1}\boldsymbol{\Sigma}_{12}\mathbf{M}\boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1} & -\boldsymbol{\Sigma}_1^{-1}\boldsymbol{\Sigma}_{12}\mathbf{M} \\ -\mathbf{M}\boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1} & \mathbf{M} \end{bmatrix},$$

where

$$\mathbf{M} = [\boldsymbol{\Sigma}_2 - \boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\boldsymbol{\Sigma}_{12}]^{-1},$$

such that

$$\boldsymbol{\beta} = \begin{bmatrix} \boldsymbol{\Sigma}_1^{-1}\mathbf{C}_1 - \boldsymbol{\Sigma}_1^{-1}\boldsymbol{\Sigma}_{12}\mathbf{M}\mathbf{v} \\ \mathbf{M}\mathbf{v} \end{bmatrix},$$

where

$$\mathbf{v} = \mathbf{C}_2 - \boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\mathbf{C}_1.$$

The residuals in regression (D.2) are given by

$$\begin{aligned}\varepsilon_2 &= y - \boldsymbol{\beta}'\mathbf{x}\\ &= y - \mathbf{C}_1'\boldsymbol{\Sigma}_1^{-1}\mathbf{x}_1 + \mathbf{v}'\mathbf{M}\boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\mathbf{x}_1 - \mathbf{v}'\mathbf{M}\mathbf{x}_2\\ &= \varepsilon_1 + \mathbf{v}'\mathbf{M}\boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\mathbf{x}_1 - \mathbf{v}'\mathbf{M}\mathbf{x}_2.\end{aligned}$$

The residuals in regression (D.1) are uncorrelated from $\mathbf{x}_1$, so the variance of ε_2 is

$$\sigma_{\varepsilon,2}^2 = \sigma_{\varepsilon,1}^2 + \mathbf{v}'\mathbf{M}\boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\boldsymbol{\Sigma}_{12}\mathbf{M}\mathbf{v} + \mathbf{v}'\mathbf{M}\boldsymbol{\Sigma}_2\mathbf{M}\mathbf{v} - 2\mathbf{v}'\mathbf{M}\boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\boldsymbol{\Sigma}_{12}\mathbf{M}\mathbf{v} - 2\mathbf{v}'\mathbf{M}\mathbb{C}\mathrm{ov}[\varepsilon_1, \mathbf{x}_2].$$

We have

$$\begin{aligned}\mathbb{C}\mathrm{ov}[\varepsilon_1, \mathbf{x}_2] &= \mathbb{C}\mathrm{ov}[y - \boldsymbol{\beta}_1'\mathbf{x}_1, \mathbf{x}_2]\\ &= \mathbf{C}_2 - \boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\mathbf{C}_1\\ &= \mathbf{v}.\end{aligned}$$

Hence

$$\begin{aligned}\sigma_{\varepsilon,2}^2 &= \sigma_{\varepsilon,1}^2 - \mathbf{v}'\mathbf{M}\boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\boldsymbol{\Sigma}_{12}\mathbf{M}\mathbf{v} + \mathbf{v}'\mathbf{M}\boldsymbol{\Sigma}_2\mathbf{M}\mathbf{v} - 2\mathbf{v}'\mathbf{M}\mathbf{v}\\ &= \sigma_{\varepsilon,1}^2 + \mathbf{v}'\mathbf{M}[\boldsymbol{\Sigma}_2 - \boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\boldsymbol{\Sigma}_{12}]\mathbf{M}\mathbf{v} - 2\mathbf{v}'\mathbf{M}\mathbf{v}\\ &= \sigma_{\varepsilon,1}^2 + \mathbf{v}'\mathbf{M}\mathbf{M}^{-1}\mathbf{M}\mathbf{v} - 2\mathbf{v}'\mathbf{M}\mathbf{v}\\ &= \sigma_{\varepsilon,1}^2 - \mathbf{v}'\mathbf{M}\mathbf{v}.\end{aligned} \tag{D.3}$$

Now consider the regression of the new explanatory variables on the existing ones:

$$\mathbf{x}_2 = \mathbf{b}'\mathbf{x}_1 + \boldsymbol{\eta}.$$

The residuals are

$$\boldsymbol{\eta} = \mathbf{x}_2 - \boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\mathbf{x}_1,$$

so their covariance matrix is

$$\begin{aligned}\mathbb{V}[\boldsymbol{\eta}] &= \boldsymbol{\Sigma}_2 - \boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\boldsymbol{\Sigma}_{12}\\ &= \mathbf{M}^{-1},\end{aligned}$$

and the correlation matrix is

$$\begin{aligned}\mathbb{C}\mathrm{orr}[\boldsymbol{\eta}] &= \sqrt{\mathrm{diag}\mathbb{V}[\boldsymbol{\eta}]}^{-1}\mathbb{V}[\boldsymbol{\eta}]\sqrt{\mathrm{diag}\mathbb{V}[\boldsymbol{\eta}]}^{-1}\\ &= \sqrt{\mathrm{diag}\mathbb{V}[\boldsymbol{\eta}]}^{-1}\mathbf{M}^{-1}\sqrt{\mathrm{diag}\mathbb{V}[\boldsymbol{\eta}]}^{-1}\end{aligned} \tag{D.4}$$

The covariance vector between the residuals and y is

$$\begin{aligned}\mathbb{C}\mathrm{ov}[y, \boldsymbol{\eta}] &= \mathbf{C}_2 - \boldsymbol{\Sigma}_{12}'\boldsymbol{\Sigma}_1^{-1}\mathbf{C}_1\\ &= \mathbf{v},\end{aligned}$$

so the correlation vector is

$$\mathbb{C}\mathrm{orr}[y, \boldsymbol{\eta}] = \frac{1}{\sigma_y}\sqrt{\mathrm{diag}\mathbb{V}[\boldsymbol{\eta}]}^{-1}\mathbf{v}, \tag{D.5}$$

where σ_y is the volatility of y.

With Equations (D.4) and (D.5), we obtain

$$\mathbf{v}'\mathbf{M}\mathbf{v} = \sigma_y^2\mathbb{C}\mathrm{orr}[y, \boldsymbol{\eta}]'\mathbb{C}\mathrm{orr}[\boldsymbol{\eta}]^{-1}\mathbb{C}\mathrm{orr}[y, \boldsymbol{\eta}], \tag{D.6}$$

so, by Equation (D.3), the variance of the residuals is

$$\sigma_{\varepsilon,2}^2 = \sigma_{\varepsilon,1}^2 - \sigma_y^2\mathbb{C}\mathrm{orr}[y, \boldsymbol{\eta}]'\mathbb{C}\mathrm{orr}[\boldsymbol{\eta}]^{-1}\mathbb{C}\mathrm{orr}[y, \boldsymbol{\eta}].$$

The R-square of regression (D.2) is $1 - \sigma_{\varepsilon,2}^2/\sigma_y^2$, so the increase in R-square from adding the regressors $\mathbf{x}_2$ is

$$\Delta R^2 = \mathbb{C}\mathrm{orr}[y, \boldsymbol{\eta}]'\mathbb{C}\mathrm{orr}[\boldsymbol{\eta}]^{-1}\mathbb{C}\mathrm{orr}[y, \boldsymbol{\eta}]. \tag{D.7}$$

When there is a single new regressor, the correlation "matrix" is the scalar 1, so the change in R-square is $\mathbb{C}\mathrm{orr}[y, \eta]^2$. Numerical examples of this formula are given in Appendix 3.2.2.

Alternative Formula for Gain in R-Square

The gain in R-square from adding regressors can also be expressed as the R-square of a regression. To see this, consider the regression of y on the residuals $\boldsymbol{\eta}$, which themselves come from the regression of new explanatory variables on the existing ones:

$$y = \mathbf{c}'\boldsymbol{\eta} + u. \tag{D.8}$$

The variance of the residuals is

$$\mathbb{V}[u] = \sigma_y^2 + \mathbf{c}'\mathbb{V}[\boldsymbol{\eta}]\mathbf{c} - 2\mathbf{c}'\mathbb{C}\mathrm{ov}[y, \boldsymbol{\eta}],$$

and the above calculations have shown that $\mathbb{V}[\boldsymbol{\eta}] = \mathbf{M}^{-1}$ and $\mathbb{C}\mathrm{ov}[y, \boldsymbol{\eta}] = \mathbf{v}$, such that

$$\mathbb{V}[u] = \sigma_y^2 + \mathbf{c}'\mathbf{M}^{-1}\mathbf{c} - 2\mathbf{c}'\mathbf{v}.$$

Moreover, the vector of slopes in Equation (D.8) is

$$\begin{aligned}\mathbf{c} &= \mathbb{V}[\boldsymbol{\eta}]^{-1}\mathbb{C}\mathrm{ov}[y, \boldsymbol{\eta}]\\ &= \mathbf{M}\mathbf{v},\end{aligned}$$

so

$$\begin{aligned}\mathbb{V}[u] &= \sigma_y^2 + \mathbf{v}'\mathbf{M}\mathbf{v} - 2\mathbf{v}'\mathbf{M}\mathbf{v}\\ &= \sigma_y^2 - \mathbf{v}'\mathbf{M}\mathbf{v}.\end{aligned}$$

Therefore, the R-square in regression (D.8) is

$$R^2_{y/\boldsymbol{\eta}} = 1 - \frac{\mathbb{V}[u]}{\sigma_y^2} = \frac{\mathbf{v}'\mathbf{M}\mathbf{v}}{\sigma_y^2}.$$

From Equations (D.6) and (D.7), it follows that

$$\Delta R^2 = R^2_{y/\boldsymbol{\eta}}.$$

This expression allows the gain in R-square to be interpreted as the marginal explanatory power of the new regressors when the influence of the existing regressors has been eliminated.

E. Decomposing the Returns on Constant-Maturity Bonds

In this paper, to hold a "constant-maturity" bond means to roll over pure discount bonds with identical maturities. The two parameters that describe the strategy are the constant maturity, denoted by u, and the time step between two consecutive bond purchases, denoted by h. Let t and $t + h$ be two consecutive roll-over dates. The arithmetic return on the strategy between these dates is

$$r_{t,t+h} = \frac{b_{t+h,u-h}}{b_{t,u}} - 1,$$

where $b_{t,u}$ is the price at date t of a pure discount bond with residual maturity u. The numerator in the right-hand side of this formula is $b_{t+h,u-h}$ because at date $t + h$, the maturity of the bond held in the portfolio is shorter by h.

The log return is

$$\begin{aligned}\log[1 + r_{t,t+h}] &= \log \frac{\exp(-[u-h]y_{t+h,u-h})}{\exp(-uy_{t,u})} \\ &= uy_{t,u} - [u-h]y_{t+h,u-h} \\ &= hy_{t,u} + [u-h][y_{t,u} - y_{t+h,u-h}] \\ &= hy_{t,u} + [u-h][y_{t+h,u} - y_{t+h,u-h}] + [u-h][y_{t,u} - y_{t+h,u}]. \qquad \text{(E.1)}\end{aligned}$$

The first term is an "income" term, which is always positive – as long as the rate of maturity u remains positive. The second term is a "roll-down return" that is positive if the yield curve at time $t + h$ is locally increasing near the maturity u: the rate of maturity u must be greater than the rate of maturity $u - h$. The last term is a "capital gain" term, which is positive if the rate of maturity u is lower at date $t + h$ than at date t, and negative otherwise.

Because the decomposition of the single-period log return is additive, a similar decomposition is obtained for the average log return over a period. It is plotted in Figure 4 for the sample period from January 1986 to January 2019.

F. Factor Exposures of Bonds

Assume that zero-coupon rates are exactly generated by K factors, so that the rate of maturity u prevailing at date t is of the form

$$y_{t,u} = c(u) + \mathbf{b}(u)'\mathbf{X}_t,$$

where $\mathbf{X}_t$ is a $K \times 1$ vector of factor values and $c(u)$ and $\mathbf{b}(u)$ are two functions of the maturity.

Factor Exposures of Retirement Bond

Based on Equation (3.1), the retirement bond price can be written as

$$B_t = \sum_i [1+\pi]^{T+i} \exp\left(-[T+i-t]\left[c(T+i-t) + \mathbf{b}(T+i-t)'\mathbf{X}_t\right]\right), \tag{F.1}$$

where we have alleviated the notation with respect to the aforementioned equation by omitting the summation bounds. The retirement bond price is a function of time and the factor values, so its arithmetic return over a period $[t,\ t+h]$ can be approximated at the first order as

$$r^B_{t,t+h} \approx \frac{1}{B_t}\frac{\partial B_t}{\partial t}h + \frac{1}{B_t}\left[\frac{\partial B_t}{\partial \mathbf{X}_t}\right]' \Delta \mathbf{X}_t,$$

where $\partial B_t / \partial \mathbf{X}_t$ denotes the gradient vector with respect to the factors, which can be calculated with Equation (F.1):

$$\frac{1}{B_t}\frac{\partial B_t}{\partial \mathbf{X}_t} = -\frac{1}{B_t}\sum_i [1+\pi]^{T+i} \times [T+i-t] \times \exp\left(-[T+i-t]y_{t,T+i-t}\right)\mathbf{b}(T+i-t).$$

If the k^{th} factor, denoted by $X_{t,k}$, is a "pure level factor", i.e. such that $b_k(u)$ is independent from the maturity u, then the sensitivity of the bond price with respect to this factor can be written as

$$\frac{1}{B_t}\frac{\partial B_t}{\partial X_{t,k}} = -D^B_{t,k} b_k,$$

where $D^B_{t,k}$ is defined as

$$D^B_{t,k} = \frac{1}{B_t}\sum_i [1+\pi]^{T+i} \times [T+i-t] \times \exp\left(-[T+i-t]y_{t,T+i-t}\right).$$

$D^B_{t,k}$ is the "duration" of the bond. It is not the usual modified duration of the bond because it is not defined as the sensitivity of the bond price with respect to a change in the yield to maturity, but it has a similar expression. The difference with respect to the modified duration is that each cash flow is discounted at its own zero-coupon rate, while in the modified duration, they are all discounted at the same rate, which is the yield to maturity.

Factor Exposures of Constant-Maturity Bonds

Consider a roll over of zero-coupon bonds with maturity u, and the period between two consecutive roll-over dates, t and $t + h$. The log return on the bond is given by Equation (E.1):

$$\log r_{t,t+h} = hy_{t,u} + [u - h][y_{t+h,u} - y_{t+h,u-h}] + [u - h][y_{t,u} - y_{t+h,u}].$$

Assume that the roll-down return (the second term in the right-hand side) is small compared to the sum of the income and capital gain terms (the other two). The decomposition of average bond returns in Figure 4 provides some support for this assumption. The log return can then be approximated as

$$\begin{aligned}\log r_{t,t+h} &\approx hy_{t,u} + [u - h][y_{t,u} - y_{t+h,u}]\\ &= hc(u) + h\mathbf{b}(u)'\mathbf{X}_t - [u - h]\mathbf{b}(u)'\Delta\mathbf{X}_t.\end{aligned}$$

At the first order, the arithmetic return is approximately equal to the log return, so we have

$$r_{t,t+h} - 1 \approx h[c(u) + \mathbf{b}(u)'\mathbf{X}_t] - [u - h]\mathbf{b}(u)'\Delta\mathbf{X}_t.$$

Now consider a portfolio invested in N constant-maturity bonds, with maturities $u_1, ..., u_N$, and weights $w_1, ..., w_N$. Denoting the return of each constituent by a superscript i, we find that the portfolio return between dates t and $t + h$ is

$$r^p_{t,t+h} = \sum_{i=1}^{N} w_i[r^i_{t,t+h} - 1].$$

Thus, it can be approximated as

$$r^p_{t,t+h} \approx h\sum_{i=1}^{N} w_i[c(u_i) + \mathbf{b}(u_i)'\mathbf{X}_t] - \left[\sum_{i=1}^{N} w_i[u_i - h]\mathbf{b}(u_i)\right]'\Delta\mathbf{X}_t.$$

If the k^{th} is a pure level factor, $b_k(u)$ is a flat function of the maturity u, so the derivative of the portfolio return with respect to a change in this factor is

$$-\left[\sum_{i=1}^{N} w_i[u_i - h]\right] b_k.$$

G. Maximum Annual Withdrawal Rate

In this appendix, we derive the expression for the maximum percentage of pre-retirement savings that can be withdrawn every year in decumulation from the pension pot.

The retirement date is denoted by T, and the number of replacement cash flows by τ. Withdrawals take place on dates $T + h$, where h ranges from 0 to $\tau - 1$ with a step of 1. During the entire decumulation phase, savings are invested in a fund whose return between dates T and $T + h$ is denoted by $r_{T,T+h}$. Consider an individual endowed with pre-retirement savings W_{T-} and making withdrawals that rise at rate π. The withdrawal at date $T + h$ is therefore

$$CF_{T+h} = CF_T[1+\pi]^h.$$

Wealth at date $T + 1$ is

$$W_{T+1} = [W_{T-} - CF_T]\, r_{T,T+1} - CF_{T+1},$$

and wealth at time $T + 2$ is

$$\begin{aligned} W_{T+2} &= W_{T+1} r_{T+1,T+2} - CF_{T+2} \\ &= [W_{T-} - CF_T]\, r_{T,T+1} r_{T+1,T+2} - CF_{T+1} r_{T+1,T+2} - CF_{T+2} \\ &= [W_{T-} - CF_T]\, r_{T,T+2} - CF_{T+1} r_{T+1,T+2} - CF_{T+2}. \end{aligned}$$

More generally, it is shown by induction on h that for each $h = 0, \ldots, \tau - 1$, we have

$$W_{T+h} = W_{T-} \times r_{T,T+h} - \sum_{i=0}^{h} CF_{T+i} r_{T+i,T+h}.$$

In particular, wealth just after the last withdrawal is

$$W_{T+\tau-1} = W_{T-} \times r_{T,T+\tau-1} - \sum_{h=0}^{\tau-1} CF_{T+h} r_{T+h,T+\tau-1}.$$

It is non-negative if, and only if,

$$\sum_{h=0}^{\tau-1} CF_{T+h} r_{T+h,T+\tau-1} \leq W_{T-} \times r_{T,T+\tau-1},$$

or, equivalently,

$$CF_T \sum_{h=0}^{\tau-1} [1+\pi]^h r_{T+h,T+\tau-1} \leq W_{T-} \times r_{T,T+\tau-1}.$$

So the maximum annual withdrawal rate, i.e. the maximum possible value for CF_T/W_{T-}, is

$$\delta = \frac{r_{T,T+\tau-1}}{\sum_{h=0}^{\tau-1}[1+\pi]^h r_{T+h,T+\tau-1}} = \frac{1}{\sum_{h=0}^{\tau-1}\frac{[1+\pi]^h}{r_{T,T+h}}}. \tag{G.1}$$

Special Case: The Retirement Bond

Let

$$\rho_{T+h} = [1+\pi]^h$$

be the cash flow of the retirement bond at date $T+h$. The total return index, $\tilde{\beta}$, is defined by the following conditions

$$\tilde{\beta}_T = \beta_{T-},$$

$$\frac{\tilde{\beta}_t}{\tilde{\beta}_T} = \frac{\beta_t}{\beta_{T-}} \prod_{0\leq h\leq t-T} \frac{\beta_{[T+h]-}}{\beta_{T+h}}.$$

It is the value of a self-financed portfolio that contains one share of the retirement bond just before date T, and subsequently contains a growing number of shares as cash flows are reinvested in the bond.

Remember that

$$\beta_{T+h} = \beta_{[T+h]-} - \rho_{T+h},$$

so we can derive an alternative expression for the total return on the retirement bond between dates T and t when $t > T$

$$\frac{\tilde{\beta}_t}{\tilde{\beta}_T} = \frac{\beta_t}{\beta_{T-}} \times \frac{\beta_{t-}}{\beta_t} \times \frac{\beta_{T-}}{\beta_T} \times \prod_{0<h<t-T}\left[1+\frac{\rho_{T+h}}{\beta_{T+h}}\right] = \frac{\beta_{t-}}{\beta_T}\prod_{0<h<t-T}\left[1+\frac{\rho_{T+h}}{\beta_{T+h}}\right]. \tag{G.2}$$

We now want to show that δ defined in Equation (G.1) equals $1/\beta_{T-}$ when savings are invested in the retirement bond. To this end, we consider the following induction hypothesis, $\mathcal{H}_k$, for $k = 1, ..., \tau - 1$:

$$\sum_{h=\tau-k}^{\tau-1}\frac{\tilde{\beta}_T}{\tilde{\beta}_{T+h}}\rho_{T+h} = \beta_T \prod_{0<h<\tau-k}\frac{\beta_{T+h}}{\beta_{[T+h]-}}.$$

$\mathcal{H}_1$ holds true. Indeed,

$$\sum_{h=\tau-1}^{\tau-1}\frac{\tilde{\beta}_T}{\tilde{\beta}_{T+h}}\rho_{T+h} = \frac{\tilde{\beta}_T}{\tilde{\beta}_{T+\tau-1}}\rho_{T+\tau-1} = \frac{\beta_T}{\beta_{[T+\tau-1]-}}\rho_{T+\tau-1}\prod_{0<h<\tau-1}\frac{\beta_{T+h}}{\beta_{[T+h]-}},$$

where the second equality uses Equation (G.2). Moreover, $\rho_{T+\tau-1} = \beta_{[T+\tau-1]-}$ because date $T + \tau - 1$ is the last cash flow date and the bond price is zero just after this payment. So

$$\sum_{h=\tau-1}^{\tau-1} \frac{\tilde{\beta}_T}{\tilde{\beta}_{T+h}} \rho_{T+h} = \beta_T \prod_{0<h<\tau-1} \frac{\beta_{T+h}}{\beta_{[T+h]-}}.$$

Next, $\mathcal{H}_k$ implies $\mathcal{H}_{k+1}$ when $k < \tau - 1$. Indeed, assume that $\mathcal{H}_k$ holds true. Then

$$\begin{aligned}
\sum_{h=\tau-k-1}^{\tau-1} \frac{\tilde{\beta}_T}{\tilde{\beta}_{T+h}} \rho_{T+h} &= \beta_T \prod_{0<h<\tau-k} \frac{\beta_{T+h}}{\beta_{[T+h]-}} + \frac{\tilde{\beta}_T}{\tilde{\beta}_{T+\tau-k-1}} \rho_{T+\tau-k-1} \\
&= \beta_T \prod_{0<h<\tau-k} \frac{\beta_{T+h}}{\beta_{[T+h]-}} + \rho_{T+\tau-k-1} \frac{\beta_T}{\beta_{[T+\tau-k-1]-}} \prod_{0<h<\tau-k-1} \frac{\beta_{T+h}}{\beta_{[T+h]-}} \\
&= \beta_T \left[\prod_{0<h<\tau-k-1} \frac{\beta_{T+h}}{\beta_{[T+h]-}} \right] \times \left[\frac{\beta_{T+\tau-k-1}}{\beta_{[T+\tau-k-1]-}} + \frac{\rho_{T+\tau-k-1}}{\beta_{[T+\tau-k-1]-}} \right] \\
&= \beta_T \prod_{0<h<\tau-k-1} \frac{\beta_{T+h}}{\beta_{[T+h]-}}.
\end{aligned}$$

To skip from the second-last to the last equality, we have used the fact that $\beta_{[T+\tau-k-1]-} = \beta_{T+\tau-k-1} + \rho_{T+\tau-k-1}$.

Hence, $\mathcal{H}_{\tau-1}$ holds true, such that

$$\begin{aligned}
\sum_{h=1}^{\tau-1} \frac{\tilde{\beta}_T}{\tilde{\beta}_{T+h}} \rho_{T+h} &= \beta_T \prod_{0<h<1} \frac{\beta_{T+h}}{\beta_{[T+h]-}} \\
&= \beta_T.
\end{aligned} \tag{G.3}$$

(An empty product is equal to 1.) Hence, when savings are fully invested in the retirement bond, the reciprocal of the maximum withdrawal rate is

$$\begin{aligned}
\frac{1}{\delta} &= \sum_{h=0}^{\tau-1} \rho_{T+h} \frac{\tilde{\beta}_T}{\tilde{\beta}_{T+h}} \\
&= \rho_T + \sum_{h=1}^{\tau-1} \rho_{T+h} \frac{\tilde{\beta}_T}{\tilde{\beta}_{T+h}} \\
&= \rho_T + \beta_T \\
&= \beta_{T-}.
\end{aligned}$$

The third equality follows from the second because Equation (G.3) holds.

H. Excess Returns on Zero-Coupon Bonds

Let $b_{t,T}$ denote the price at date t of a pure discount bond with unit face value maturing at date T: thus, the bond maturity is $T - t$. Now let $y_{t,u}$ denote the annualized and continuously compounded zero-coupon rate of maturity u that prevails at date t. The log excess return of the bond between dates t and $t + h$ is

$$\begin{aligned}
\log \frac{b_{t+h,T}}{b_{t,T}} - hy_{t,h} =& [T-t]y_{t,T-t} - [T-t-h]y_{t+h,T-t-h} - hy_{t,h} \\
=& [T-t][y_{t,T-t} - y_{t,h}] - [T-t-h][y_{t+h,T-t-h} - y_{t,h}] \\
=& [T-t][y_{t,T-t} - y_{t,h}] - [T-t-h][y_{t+h,T-t-h} - y_{t+h,h}] - [T-t-h][y_{t+h,h} - y_{t,h}] \\
=& [T-t][y_{t,T-t} - y_{t,h}] - [T-t-h][y_{t+h,T-t-h} - y_{t+h,h}] - [T-t-h][y_{t+h,h} - y_{t,h}] \\
=& [T-t][y_{t,T-t} - y_{t,h}] - [T-t-h][y_{t+h,T-t} - y_{t+h,h}] \\
& - [T-t-h][y_{t+h,h} - y_{t,h}] - [T-t-h][y_{t+h,T-t-h} - y_{t+h,T-t}].
\end{aligned}$$

Let the term spread $y_{t,T-t} - y_{t,h}$ be the slope measure at date t, denoted by s_t, and take the short-term rate, $y_{t,h}$, as the level measure, l_t. Then

$$\begin{aligned}
\log \frac{b_{t+h,T}}{b_{t,T}} - hy_{t,h} =& [T-t]s_t - [T-t-h]s_{t+h} - [T-t-h][l_{t+h} - l_t] \\
& - [T-t-h][y_{t+h,T-t-h} - y_{t+h,T-t}] \\
=& -[T-t-h][s_{t+h} - s_t] - [T-t-h][l_{t+h} - l_t] \\
& + hs_t + [T-t-h][y_{t+h,T-t} - y_{t+h,T-t-h}].
\end{aligned}$$

The first two terms in the right-hand side of the second equation are respectively proportional to the negatives of changes in slope and in level. The third term is known as of date t and is proportional to the current slope, and the fourth is a roll-down return, proportional to the spread between the rates of maturities $T - t$ and $T - t - h$ at date $t + h$. This equation shows that the sensitivities of the excess return to the changes in the slope and in the level are both equal to $-[T - t - h]$.

I. Defining Contributions of Factors to Risk

Consider the decomposition of the variance into a systematic term and a specific term (for the sake of generality, we consider a K-factor model):

$$\sigma^2 = \sum_{i,j=1}^{K} \beta_i \beta_j \mathbb{C}\mathrm{ov}[F_i, F_j] + \sigma_\varepsilon^2.$$

Grouping terms in the sum in the right-hand side, we rewrite the variance as

$$\sigma^2 = \sum_{i=1}^{K} \beta_i \sum_{j=1}^{K} \beta_j \mathbb{C}\mathrm{ov}[F_i, F_j] + \sigma_\varepsilon^2.$$

Define the contribution of the i^{th} factor as

$$q_i = \frac{1}{\sigma} \beta_i \sum_{j=1}^{K} \beta_j \mathbb{C}\mathrm{ov}[F_i, F_j],$$

and the contribution of idiosyncratic risk as

$$q_\varepsilon = \frac{\sigma_\varepsilon^2}{\sigma}.$$

This gives us

$$\sum_{i=1}^{K} q_i + q_\varepsilon = \sigma.$$

By definition of the R-square, we have

$$R^2 = 1 - \frac{\sigma_\varepsilon^2}{\sigma^2},$$

so

$$q_\varepsilon = \sigma[1 - R^2].$$

The contribution of factor i can be rewritten as

$$\begin{aligned} q_i &= \frac{\beta_i}{\sigma} \mathbb{C}\mathrm{ov}\left[F_i, \sum_{j=1}^{K} \beta_j F_j\right] \\ &= \frac{\beta_i}{\sigma} \mathbb{C}\mathrm{ov}\left[F_i, \sum_{j=1}^{K} \beta_j F_j + \varepsilon\right] \\ &= \frac{\beta_i}{\sigma} \mathbb{C}\mathrm{ov}[F_i, r]. \end{aligned}$$

J. Derivation of Tracking Error-Matching Allocation

In this section, we calculate the tracking error-matching allocation to an alternative equity benchmark. It is the percentage allocation to this equity index such that a portfolio invested in the index and in a perfect LHP has the same ex-ante tracking error as a portfolio invested in a reference equity index and the LHP.

Let x_0 denote the weight of equities in the original portfolio and x their weight in the new portfolio, and let $r_{PSP,0}$ and r_{PSP} denote the respective returns of the broad equity index and the new index over the next rebalancing period.

Now let r_{LHP} be the return of the LHP over the same period. The respective returns of the LDI strategies that respectively use the original PSP and the alternative one are

$$r_{LDI,0} = x_0 r_{PSP,0} + [1 - x_0] r_{LHP},$$
$$r_{LDI} = x r_{PSP} + [1 - x] r_{LHP}.$$

The tracking errors are the standard deviations of the excess returns $r_{LDI,0} - r_L$ and $r_{LDI} - r_L$, where r_L is the relative change in the value of liabilities. But the LHP perfectly replicates liabilities, so r_{LHP} equals r_L, and the two excess returns are given by

$$r_{LDI,0} - r_L = x_0 [r_{PSP,0} - r_L],$$
$$r_{LDI} - r_L = x [r_{PSP} - r_L],$$

so their variances are

$$\mathbb{V}[r_{LDI,0}] = x_0^2 \mathbb{V}[r_{PSP,0} - r_L],$$
$$\mathbb{V}[r_{LDI}] = x^2 \mathbb{V}[r_{PSP} - r_L].$$

The tracking error-matching allocation is such that the two variances are equal. Remember that x_0 is a fixed positive number (taken to be 40% in the empirical illustration), and we are interested in long-only strategies, so x must be positive too. As a result, the tracking error-matching weight is

$$x = x_0 \times \sqrt{\frac{\mathbb{V}[r_{LDI,0} - r_L]}{\mathbb{V}[r_{LDI} - r_L]}}.$$

The right-hand side involves the respective tracking errors of the two equity indices, $TE_{PSP,0}$ and TE_{PSP}, so the tracking error-matching allocation can be rewritten as

$$x = x_0 \times \frac{TE_{PSP,0}}{TE_{PSP}}.$$

References

References

• Alquist, R., R. Israel and T. Moskowitz. 2018. Fact, Fiction, and the Size Effect. *The Journal of Portfolio Management* 45(1): 34–61.

• Amenc, N., R. Deguest, F. Goltz, A. Lodh and E. Shirbini. 2014a. Risk Allocation, Factor Investing and Smart Beta: Reconciling Innovations in Equity Portfolio Construction. EDHEC-Risk Institute Publication.

• Amenc, N., F. Ducoulombier, M. Esakia, F. Goltz and S. Sivasubramanian. 2017. Accounting for Cross-Factor Interactions in Multifactor Portfolios without Sacrificing Diversification and Risk Control. *The Journal of Portfolio Management* 43(5): 99–114.

• Amenc, N., F. Ducoulombier, F. Goltz, A. Lodh and S. Sivasubramanian. 2016. Diversified or Concentrated Factor Tilts? *Journal of Portfolio Management* 42(2): 64–76.

• Amenc, N. and F. Goltz. 2013. Smart Beta 2.0. *Journal of Index Investing* 4(3): 15–23.

• Amenc, N., F. Goltz and A. Lodh. 2012. Choose Your Betas: Benchmarking Alternative Equity Index Strategies. *The Journal of Portfolio Management* 39(1): 88–111.

• Amenc, N., F. Goltz, A. Lodh and L. Martellini. 2014b. Towards Smart Equity Factor Indices: Harvesting Risk Premia Without Taking Unrewarded Risks. *The Journal of Portfolio Management* 40(4): 106–122.

• Amenc, N., F. Goltz, L. Martellini and P. Retkowsky. 2011. Efficient Indexation: An Alternative to Cap-Weighted Indices. *Journal of Investment Management* 9(4): 1–23.

• Ang, A., R. J. Hodrick, Y. Xing and X. Zhang. 2006. The Cross-Section of Volatility and Expected Returns. *Journal of Finance* 61(1): 259–299.

• —. 2009. High Idiosyncratic Volatility and Low Returns: International and Further US Evidence. *Journal of Financial Economics* 91(1): 1–23.

• Asness, C. S., A. Frazzini and L. H. Pedersen. 2019. Quality Minus Junk. *Review of Accounting Studies* 24(1): 34–112.

• Asness, C. S., T. J. Moskowitz and L. H. Pedersen. 2013. Value and Momentum Everywhere. *Journal of Finance* 68(3): 929–985.

• Baker, M. P., B. Bradley and J. Wurgler. 2011. Benchmarks as Limits to Arbitrage: Understanding the Low-Volatility Anomaly. *Financial Analysts Journal* 67(1): 40–54.

• Baker, N. L. and R. A. Haugen. 2012. Low Risk Stocks Outperform Within All Observable Markets of the World. Working Paper. .

• Bali, T. and N. Cakici. 2008. Idiosyncratic Volatility and the Cross-Section of Expected Returns. *Journal of Financial and Quantitative Analysis* 43(1): 29–58.

• Bansal, R., R. F. Dittmar and C. T. Lundblad. 2005. Consumption, Dividends, and the Cross Section of Equity Returns. *Journal of Finance* 60(4): 1639–1672.

• Banz, R. W. 1981. The Relationship Between Return and Market Value of Common Stocks. *Journal of Financial Economics* 9(1): 3–18.

• Barberis, N. and M. Huang. 2001. Mental Accounting, Loss Aversion, and Individual Stock Returns. *Journal of Finance* 56(4): 1247–1292.

• —. 2008. Stocks as Lotteries: The Implications of Probability Weighting for Security Prices. *The American Economic Review* 98(5): 2066–2100.

References

• Basu, S. 1977. Investment Performance of Common Stocks in Relation to Their Price-Earnings Ratios: A Test of the Efficient Market Hypothesis. *Journal of Finance* 32(3): 663–682.

• –. 1983. The Relationship Between Earnings' Yield, Market Value and Return for NYSE Common Stocks: Further Evidence. *Journal of Financial Economics* 12(1): 129–156.

• Benartzi, S., A. Previtero and R. H. Thaler. 2011. Annuitization Puzzles. *Journal of Economic Perspectives* 25(4): 143–64.

• Black, F. 1972. Capital Market Equilibrium with Restricted Borrowing. *Journal of Business* 45(3): 444–455.

• Black, F., M. Jensen and M. Scholes. 1972. The Capital Asset Pricing Model: Some Empirical Tests. In Studies in the Theory of Capital Markets. New York: Praeger.

• Blitz, D. and P. Van Vliet. 2007. The Volatility Effect: Lower Risk without Lower Return. *Journal of Portfolio Management* 34(1): 102–113.

• Bollerslev, T. 1986. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics* 31(3): 307–327.

• Breeden, D. 1979. An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities. *Journal of Financial Economics* 7(3): 265–296.

• Campbell, J. 1993. Intertemporal Asset Pricing Without Consumption Data. *The American Economic Review* 83(3): 487–512.

• –. 1996. Understanding Risk and Return. *Journal of Political Economy* 104(2): 298–345.

• Campbell, J., Y. Chan and L. Viceira. 2003. A Multivariate Model of Strategic Asset Allocation. *Journal of Financial Economics* 67(1): 41–80.

• Campbell, J. and R. Shiller. 1988. Stock Prices, Earnings and Expected Dividends. *Journal of Finance* 43(3): 661–676.

• Campbell, J. and L. Viceira. 2001. Who Should Buy Long-Term Bonds? *American Economic Review* 91(1): 99–127.

• Carhart, M. M. 1997. On Persistence in Mutual Fund Performance. *Journal of Finance* 52(1): 57–82.

• Carli, T., R. Deguest and L. Martellini. 2014. Improved Risk Reporting with Factor-Based Diversification Measures. EDHEC-Risk Institute Publication.

• Chan, K. and N.-F. Chen. 1991. Structural and Return Characteristics of Small and Large Firms. *The Journal of Finance* 46(4): 1467–1484.

• Chan, L. K., J. Karceski and J. Lakonishok. 1998. The Risk and Return From Factors. *Journal of Financial and Quantitative Analysis* 33(2): 159–188.

• Chen, N.-F., R. Roll and S. A. Ross. 1986. Economic Forces and the Stock Market. *Journal of Business* 59(3): 383–403.

• Chordia, T. and L. Shivakumar. 2002. Momentum, Business Cycle, and Time-Varying Expected Returns. *Journal of Finance* 57(2): 985–1019.

• Cieslak, A. and P. Povala. 2015. Expected Returns in Treasury Bonds. *The Review of Financial Studies* 28(10): 2859–2901.

• Cochrane, J. 2005. Asset Pricing Theory. Princeton University Press. Revised Edition.

References

• Cochrane, J. H. and M. Piazzesi. 2005. Bond Risk Premia. *The American Economic Review* 95(1): 138–160.

• Conrad, J., R. F. Dittmar and E. Ghysels. 2013. Ex Ante Skewness and Expected Stock Returns. *The Journal of Finance* 68(1): 85–124.

• Conrad, J. and G. Kaul. 1998. An Anatomy of Trading Strategies. *The Review of Financial Studies* 11(3): 489–519.

• Coqueret, G., L. Martellini and V. Milhau. 2017. Equity Benchmarks with Improved Liability-Hedging Properties. *Journal of Portfolio Management* 43(2): 37–49.

• Daniel, K. and S. Titman. 1997. Evidence on the Characteristics of Cross Sectional Variation in Stock Returns. *Journal of Finance* 52(1): 1–33.

• Daniel, L., D. Hirshleifer and A. Subrahmanyam. 1998. Investor Psychology and Security Market Under- and Overreactions. *Journal of Finance* 53(6): 1839–1885.

• DeBondt, W. F. and R. Thaler. 1985. Does the Stock Market Overreact? *Journal of Finance* 40(3): 793–805.

• Deguest, R., L. Martellini and V. Milhau. 2014. Hedging versus Insurance: Long-Horizon Investing with Short-Term Constraints. *Bankers, Markets and Investors* (Special Issue March-April 2014): 33–47.

• –. 2018a. A Reinterpretation of the Optimal Demand for Risky Assets in Fund Separation Theorems. *Management Science* 64(9): 4333–4347.

• –. 2018b. Bond Portfolio Optimization in the Presence of Duration Constraints. *Journal of Fixed Income* 28(1): 6–26.

• Elton, E. and M. Gruber. 1973. Estimating the Dependence Structure of Share Prices – Implications for Portfolio Selection. *Journal of Finance* 28(5): 1203–1232.

• Engle, R. 1982. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* 55(2): 251–276.

• Fama, E. and K. French. 1988. Dividend Yields and Expected Stock Returns. *Journal of Financial Economics* 22(1): 3–25.

• –. 1989. Business Conditions and Expected Returns on Stocks and Bonds. *Journal of Financial Economics* 25(1): 23–49.

• –. 2014. A Five-Factor Asset Pricing Model. *Journal of Financial Economics* 116(1): 1–22.

• Fama, E. F. 1986. Term Premiums and Default Premiums in Money Markets. *Journal of Financial Economics* 17(1): 175–196.

• –. 1991. Efficient Capital Markets: II. *Journal of Finance* 46(5): 1575–1617.

• Fama, E. F. and R. R. Bliss. 1987. The Information in Long-Maturity Forward Rates. *The American Economic Review* 77(4): 680–692.

• Fama, E. F. and K. R. French. 1992. The Cross-Section of Expected Stock Returns. *Journal of Finance* 47(2): 427–465.

• –. 1993. Common Risk Factors in the Returns on Stocks and Bonds. *Journal of Financial Economics* 33(1): 3–56.

• –. 1995. Size and Book-to-Market Factors in Earnings and Returns. *The Journal of Finance* 50(1): 131–155.

References

- —. 1996. Multifactor Explanations of Asset Pricing Anomalies. *The Journal of Finance* 51(1): 55–84.
- —. 2006. Profitability, Investment and Average Returns. *Journal of Financial Economics* 82(3): 491–518.
- —. 2008. Dissecting Anomalies. *The Journal of Finance* 63(4): 1653–1678.
- —. 2012. Size, Value, and Momentum in International Stock Returns. *Journal of Financial Economics* 105(3): 457–472.
- Fama, E. F. and J. D. MacBeth. 1973. Risk, Return, and Equilibrium: Empirical Tests. *Journal of Political Economy* 81(3): 607–636.
- Fama, E. F. and G. W. Schwert. 1977. Human Capital and Capital Market Equilibrium. *Journal of Financial Economics* 4(1): 95–125.
- Flannery, M. J. and A. A. Protopapadakis. 2002. Macroeconomic Factors Do Influence Aggregate Stock Returns. *Review of Financial Studies* 15(3): 751–782.
- Frazzini, A. and L. H. Pedersen. 2014. Betting Against Beta. *Journal of Financial Economics* 111(1): 1–25.
- Fu, F. 2009. Idiosyncratic Risk and the Cross-Section of Expected Stock Returns. *Journal of Financial Economics* 91(1): 24–37.
- Geczy, C. C. and M. Samonov. 2016. Two Centuries of Price-Return Momentum. *Financial Analysts Journal* 72(5): 32–56.
- Gibbons, M. R. 1982. Multivariate Tests of Financial Models: A New Approach. *Journal of Financial Economics* 10(1): 3–27.
- Gibbons, M. R., S. A. Ross and J. Shanken. 1989. A Test of the Efficiency of a Given Portfolio. *Econometrica* 57(5): 1121–1152.
- Gompers, P. A. and A. Metrick. 2001. Institutional Investors and Equity Prices. *Quarterly Journal of Economics* 116(1): 229–259.
- Gordon, M. J. and E. Shapiro. 1956. Capital Equipment Analysis: The Required Rate of Profit. *Management science* 3(1): 102–110.
- Goyal, A. and I. Welch. 2003. Predicting the Equity Premium with Dividend Ratios. *Management Science* 49(5): 639–654.
- Graham, B. and D. L. Dodd. 1934. Security Analysis: Principles and Technique. McGraw-Hill.
- Guo, H., H. Kassa and M. F. Ferguson. 2014. On the Relation Between EGARCH Idiosyncratic Volatility and Expected Stock Returns. *Journal of Financial and Quantitative Analysis* 49(1): 271–296.
- Gürkaynak, R., B. Sack and J. Wright. 2007. The US Treasury Yield Curve: 1961 to the Present. J*ournal of Monetary Economics* 54(8): 2291–2304.
- Harvey, C. R., Y. Liu and H. Zhu. 2016. ... and the Cross-Section of Expected Returns. *Review of Financial Studies* 29(1): 5–68.
- Haugen, R. A. and N. L. Baker. 1996. Commonality in the Determinants of Expected Stock Returns. *Journal of Financial Economics* 41(3): 401–439.
- Heston, S. L., K. G. Rouwenhorst and R. E. Wessels. 1995. The Structure of International Stock Returns and the Integration of Capital Markets. *Journal of Empirical Finance* 2(3): 173–197.

References

• Hong, H., T. Lim and J. C. Stein. 2000. Bad News Travels Slowly: Size, Analyst Coverage, and the Profitability of Momentum Strategies. *Journal of Finance* 55(1): 265–295.

• Hong, H. and J. C. Stein. 1999. A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. *Journal of Finance* 54(6): 2143–2184.

• Horowitz, J. L., T. Loughran and N. Savin. 2000. The Disappearing Size Effect. Research in Economics 54(1): 83–100.

• Hou, K. and M. A. Van Dijk. 2019. Resurrecting the Size Effect: Firm Size, Profitability Shocks, and Expected Stock Returns. *Review of Financial Studies* 37(2): 2850–2889.

• Hou, K., C. Xue and L. Zhang. 2015. Digesting Anomalies: An Investment Approach. *Review of Financial Studies* 28(3): 650–705.

• Huang, W., Q. Liu, S. Rhee and L. Zhang. 2010. Return Reversals, Idiosyncratic Risk, and Expected Returns. *Review of Financial Studies* 23(1): 147–168.

• Israel, R. and T. J. Moskowitz. 2013. The Role of Shorting, Firm Size, and Time on Market Anomalies. *Journal of Financial Economics* 108(2): 275–301.

• Jagannathan, R. and T. Ma. 2003. Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps. *Journal of Finance* 58(4): 1651–1683.

• Jagannathan, R. and Z. Wang. 1996. The Conditional CAPM and the Cross-Section of Expected Returns. *Journal of Finance* 51(1): 3–53.

• Jegadeesh, N. 1990. Evidence of Predictable Behavior of Security Returns. *Journal of Finance* 45(3): 881–898.

• Jegadeesh, N. and S. Titman. 1993. Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *Journal of Finance* 48(1): 65–91.

• –. 2001. Profitability of Momentum Strategies: An Evaluation of Alternative Explanations. *The Journal of Finance* 56(2): 699–720.

• Jensen, M. C. 1969. Risk, the Pricing of Capital Assets, and the Evaluation of Investment Portfolios. *Journal of Business* 42(2): 167–247.

• Keim, D. B. 1983. Size-Related Anomalies and Stock Return Seasonality: Further Empirical Evidence. *Journal of Financial Economics* 12(1): 13–32.

• Kobor, A. and A. Muralidhar. 2018. How a New Bond Can Greatly Improve Retirement Security. Working paper.

• Kroon, E., A. Wouters and R. L. de Carvalho. 2017. Decomposing Funding-Ratio Risk: Providing Pension Funds with Key Insights into Their Liabilities Hedge Mismatch and Other Factor Exposures. *The Journal of Portfolio Management* 43(4): 71–86.

• Lakonishok, J., A. Shleifer and R. W. Vishny. 1994. Contrarian Investment, Extrapolation, and Risk. *Journal of Finance* 49(5): 1541–1578.

• Lettau, M. and S. Ludvigson. 2001a. Consumption, Aggregate Wealth, and Expected Stock Returns. *Journal of Finance* 56(3): 815–849.

• –. 2001b. Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying. *Journal of Political Economy* 109(6): 1238–1287.

References

- Lewellen, J. and S. Nagel. 2006. The Conditional CAPM Does Not Explain Asset-Pricing Anomalies. *Journal of Financial Economics* 82(2): 289–314.
- Liew, J. and M. Vassalou. 2000. Can Book-to-Market, Size and Momentum Be Risk Factors That Predict Economic Growth? *Journal of Financial Economics* 57(2): 221–245.
- Lintner, J. 1965. The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets. *The Review of Economics and Statistics* 47(1): 13–37.
- Liu, L. X. and L. Zhang. 2008. Momentum Profits, Factor Pricing, and Macroeconomic Risk. *Review of Financial Studies* 21(6): 2417–2448.
- Maeso, J., L. Martellini and R. Rebonato. 2019a. Factor Investing in Sovereign Bond Markets - A Time-Series Perspective EDHEC-Risk Institute Publication.
- —. 2019b. Factor investing in sovereign bond markets: Cross-sectional and time-series momentum in us sovereign bond market EDHEC-Risk Institute Working Paper.
- —. 2019c. Factor investing in sovereign bond markets: Defining and exploiting value in bonds EDHEC-Risk Institute Working Paper.
- —. 2019d. Factor Investing in US Sovereign Bond Market – A New Generation of Conditional Carry Strategies with Applications in Asset-Only and Asset-Liability Management Forthcoming in *The Journal of Portfolio Management*.
- Maillard, S., T. Roncalli and J. Teiletche. 2010. The Properties of Equally Weighted Risk Contribution Portfolios. *Journal of Portfolio Management* 36(4): 60–70.
- Mandelbrot, B. 1967. The Variation of Some Other Speculative Prices. *The Journal of Business* 40(4): 393–413.
- Markowitz, H. 1952. Portfolio Selection. *Journal of Finance* 7(1): 77–91.
- Martellini, L. 2008. Toward the Design of Better Equity Benchmarks: Rehabilitating the Tangency Portfolio from Modern Portfolio Theory. *The Journal of Portfolio Management* 34(4): 34–41.
- Martellini, L. and V. Milhau. 2012. Dynamic Allocation Decisions in the Presence of Funding Ratio Constraints. Journal of Pension *Economics and Finance* 11(4): 549–580.
- —. 2015. Factor Investing: A Welfare-Improving New Investment Paradigm or Yet Another Marketing Fad? EDHEC-Risk Institute Publication.
- —. 2018a. Proverbial Baskets Are Uncorrelated Risk Factors! A Factor-Based Framework for Measuring and Managing Diversification in Multi-Asset Investment Solutions. *Journal of Portfolio Management* Multi-Asset Special Issue 44(2): 8–22.
- —. 2018b. Proverbial Baskets Are Uncorrelated Risk Factors! A Factor-Based Framework for Measuring and Managing Diversification in Multi-Asset Investment Solutions. EDHEC-Risk Institute Publication.
- —. 2018c. Smart Beta and Beyond: Maximising the Benefits of Factor Investing. EDHEC-Risk Institute Publication.
- Martellini, L., V. Milhau and A. Tarelli. 2014a. Estimation Risk versus Optimality Risk : An Ex-Ante Efficiency Analysis of Alternative Equity Portfolio Diversification Strategies. Bankers, *Markets and Investors* 132: 26–42.

References

• —. 2014b. Hedging Inflation-Linked Liabilities without Inflation-Linked Bonds Through Long/Short Investments in Nominal Bonds. *Journal of Fixed Income* 24(3): 5–29.

• Martellini, L., P. Priaulet and S. Priaulet. 2003. Fixed-Income Securities: Valuation, Risk Management and Portfolio Strategies. John Wiley & Sons.

• Mayers, D. 1972. Nonmarketable Assets and Capital Market Equilibrium Under Uncertainty. In Jensen, M. C., ed., Studies in The Theory of Capital Markets. Praeger New York, 223–248.

• Menzly, L., T. Santos and P. Veronesi. 2004. Understanding Predictability. *Journal of Political Economy* 112(1): 1–47.

• Merton, R. 1973. An Intertemporal Capital Asset Pricing Model. *Econometrica* 41(5): 867–887.

• —. 1974. On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance* 29(2): 449–470.

• —. 1980. On Estimating the Expected Return on the Market: An Exploratory Investigation. *Journal of Financial Economics* 8(4): 323–361.

• —. 1992. An Intertemporal Capital Asset Pricing Model. In Merton, R., ed., Continuous-Time Finance. Blackwell Publishers, 475–523.

• —. 2014. The Crisis in Retirement Planning. Harvard Business Review.

• Merton, R. and A. Muralidhar. 2017. Time for Retirement 'SelFIES'? Working paper.

• Merton, R. C. 1987. A Simple Model of Capital Market Equilibrium with Incomplete Information. *Journal of Finance* 42(3): 483–510.

• Meucci, A., A. Santangelo and R. Deguest. 2015. Measuring Portfolio Diversification Based on Optimized Uncorrelated Factors. *Risk* 11(29): 70–75.

• Mitton, T. and K. Vorkink. 2007. Equilibrium Underdiversification and the Preference for Skewness. *Review of Financial Studies* 20(4): 1255–1288.

• MSCI. 2007. Barra Risk Model Handbook. MSCI.

• Muralidhar, A. 2015. New Bond Would Offer a Better Way to Secure DC Plans. *Pensions and Investments*.

• Muralidhar, A., K. Ohashi and S. H. Shin. 2016. The Most Basic Missing Instrument in Financial Markets: The Case for Forward Starting Bonds. *Journal of Investment Consulting* 47(2): 34–47.

• Nelson, C. R. and A. F. Siegel. 1987. Parsimonious Modeling of Yield Curves. *Journal of Business* 60(4): 473–489.

• Nielsen, L. and M. Vassalou. 2006. The Instantaneous Capital Market Line. *Economic Theory* 28(3): 651–664.

• Novy-Marx, R. 2013. The Other Side of Value: The Gross Profitability Premium. *Journal of Financial Economics* 108(1): 1–28.

• OECD. 2016. Pensions Outlook 2016.

• Pashchenko, S. 2013. Accounting for Non-Annuitization. *Journal of Public Economics* 98: 53–67.

• Pastor, L. and R. F. Stambaugh. 2003. Liquidity Risk and Expected Stock Returns. *Journal of Political Economy* 111(3): 642–685.

References

• Petkova, R. 2006. Do the Fama-French Factors Proxy for Innovations in Predictive Variables? *Journal of Finance* 61(2): 581–612.

• Petkova, R. and L. Zhang. 2005. Is Value Riskier than Growth? *Journal of Financial Economics* 78(1): 187–202.

• Qian, E. 2005. Risk Parity Portfolios: Efficient Portfolios Through True Diversification. PanAgora.

• Ransenberg, D., P. Hodges and A. Hunt. 2012. LDI in a Risk Factor Framework. *Journal of Investing* 21(2): 105–116.

• Reinganum, M. R. 1981. Misspecification of Capital Asset Pricing: Empirical Anomalies Based on Earnings' Yields and Market Values. *Journal of Financial Economics* 9(1): 19–46.

• —. 1982. A Direct Test of Roll's Conjecture on The Firm Size Effect. *Journal of Finance* 37(1): 27–35.

• Roll, R. 1969. Bias in Fitting the Sharpe Model to Time Series Data. *Journal of Financial and Quantitative Analysis* 4(3): 271–289.

• —. 1977. A Critique of the Asset Pricing Theory's Tests. Part I: On Past and Potential Testability of the Theory. *Journal of Financial Economics* 4(2): 129–176.

• —. 1981. A Possible Explanation of the Small Firm Effect. *Journal of Finance* 36(4): 879–888.

• Rosenberg, B., K. Reid and R. Lanstein. 1985. Persuasive Evidence of Market Inefficiency. *Journal of Portfolio Management* 11(3): 9–16.

• Ross, S. 1976. The Arbitrage Theory of Capital Asset Pricing. *Journal of Economic Theory* 13(3): 341–360.

• Rouwenhorst, K. G. 1998. International Momentum Strategies. *Journal of Finance* 53(1): 267–284.
Scientific Beta. 2018. Scientific Beta Universe Construction Rules. Scientific Beta. Available at https://www.scientificbeta.com/#/documentation/ground-rules/eri-scientific-beta-universe-construction-rules.

• —. 2019. Scientific Beta Index Calculation Rules. Scientific Beta. Available at https://www.scientificbeta.com/#/documentation/ground-rules/eri-scientific-beta-strategy-calculation-rules.

• Scott, J. S., J. G. Watson and W.-Y. Hu. 2011. What Makes a Better Annuity? *Journal of Risk and Insurance* 78(1): 213–244.

• Scott, R. C. and P. A. Horvath. 1980. On the Direction of Preference for Moments of Higher Order Than the Variance. *Journal of Finance* 35(4): 915–919.

• Shanken, J. 1985. Multivariate Tests of the Zero-Beta CAPM. *Journal of Financial Economics* 14(3): 327–348.

• Sharpe, W. 1963. A Simplified Model for Portfolio Analysis. *Management Science* 9(2): 277–293.

• —. 1964. Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk. *Journal of Finance* 19(3): 425–442.

• Sharpe, W. F. 1965. Risk-Aversion in the Stock Market: Some Empirical Evidence. *Journal of Finance* 20(3): 416–422.

• Shumway, T. 1997. The Delisting Bias in CRSP Data. *The Journal of Finance* 52(1): 327–340.

• Soldofsky, R. M. and R. L. Miller. 1969. Risk Premium Curves for Different Classes of Long-Term Securities, 1950-1966. *Journal of Finance* 24(3): 429–445.

References

- Stattman, D. 1980. Book Values and Stock Returns. *The Chicago MBA: A Journal of Selected Papers* 4(1): 25–45.
- Stoll, H. R. and R. E. Whaley. 1983. Transaction Costs and the Small Firm Effect. *Journal of Financial Economics* 12(1): 57–79.
- Svensson, L. E. 1994. Estimating and Interpreting Forward Interest Rates: Sweden 1992-1994. National Bureau of Economic Research.
- Titman, S., K. J. Wei and F. Xie. 2004. Capital Investments and Stock Returns. *Journal of Financial and Quantitative Analysis* 39(4): 677–700.
- Van Dijk, M. A. 2011. Is Size Dead? A Review of the Size Effect in Equity Returns. *Journal of Banking & Finance* 35(12): 3263–3274.
- Vassalou, M. and Y. Xing. 2004. Default Risk in Equity Returns. *Journal of Finance* 59(2): 831–868.
- Welch, I. and A. Goyal. 2008. A Comprehensive Look at the Empirical Performance of Equity Premium Prediction. *Review of Financial Studies* 21(4): 1455–1508.
- Zhang, L. 2005. The Value Premium. *Journal of Finance* 60(1): 67–103.

About Amundi ETF, Indexing and Smart Beta

About Amundi ETF, Indexing and Smart Beta

With more than €133 billion[1] in assets under management, Amundi ETF, Indexing and Smart Beta is one of Amundi's strategic business areas and is a key growth driver for the Group.

Amundi ETF, Indexing and Smart Beta business line provides investors - whether institutionals or distributors - with robust, innovative, and cost-efficient solutions, leveraging Amundi Group's scale and large resources. The platform also offers investors fully customized solutions (ESG, Low Carbon, specific exclusions, risk constraints, etc.).

With over 30 years of benchmark construction and replication expertise, Amundi is a trusted name in ETF & Index management among the world's largest institutions. The team is also recognized for its ability to develop Smart Beta & Factor Investing solutions, with more than 10-year track-record.

1 - All figures and data are provided by Amundi ETF, Indexing & Smart Beta at 31/12/2019

About EDHEC-Risk Institute

About EDHEC-Risk Institute

Founded in 1906, EDHEC is one of the foremost international business schools. Operating from campuses in Lille, Nice, Paris, London and Singapore, EDHEC is one of the top 15 European business schools. Accredited by the three main international academic organisations, EQUIS, AACSB, and Association of MBAs, EDHEC has for a number of years been pursuing a strategy of international excellence that led it to set up EDHEC-Risk Institute in 2001. This Institute boasts a team of permanent professors, engineers and support staff, and counts a large number of affiliate professors and research associates from the financial industry among its ranks.

The Need for Investment Solutions and Risk Management

Investment management is justified as an industry only to the extent that it can demonstrate a capacity to add value through the design of dedicated and meaningful investor-centric investment solutions, as opposed to one-size-fits-all manager-centric investment products. After several decades of relative inertia, the much needed move towards investment solutions has been greatly facilitated by a true industrial revolution triggered by profound paradigm changes in terms of (1) mass production of cost- and risk-efficient smart factor indices; (2) mass customisation of liability-driven investing and goal-based investing strategies; and (3) mass distribution, with robo-advisor technologies. In parallel, the investment industry is strongly impacted by two other major external revolutions, namely the digital revolution and the environmental revolution.

In this fast-moving environment, EDHEC-Risk Institute positions itself as the leading academic think-tank in the area of investment solutions, which gives true significance to the investment management practice. Through our multi-faceted programme of research, outreach, education and industry partnership initiatives, our ambition is to support industry players, both asset owners and asset managers, in their efforts to transition towards a novel, welfare-improving, investment management paradigm.

EDHEC-Risk New Initiatives

In addition to the EDHEC Alternative Indexes, which are used as performance benchmarks for risk analysis by investors in hedge funds, and the EDHEC-IEIF Monthly Commercial Property index, which tracks the performance of the French commercial property market through SCPIs, EDHEC-Risk has recently launched a series of new initiatives.

• The EDHEC-Princeton Retirement Goal-Based Investing Index Series, launched in May 2018, which represent asset allocation benchmarks for innovative mass-customised target-date solutions for individuals preparing for retirement;

• The EDHEC Bond Risk Premium Monitor, the purpose of which is to offer to investment and academic communities a tool to quantify and analyse the risk premium associated with Government bonds;

• The EDHEC-Risk Investment Solutions (Serious) Game, which is meant to facilitate engagement with graduate students or investment professionals enrolled on one of EDHEC-Risk's various campus-based, blended or fully-digital educational programmes.

About EDHEC-Risk Institute

Academic Excellence and Industry Relevance

In an attempt to ensure that the research it carries out is truly applicable, EDHEC has implemented a dual validation system for the work of EDHEC-Risk. All research work must be part of a research programme, the relevance and goals of which have been validated from both an academic and a business viewpoint by the Institute's advisory board. This board is made up of internationally recognised researchers, the Institute's business partners, and representatives of major international institutional investors. Management of the research programmes respects a rigorous validation process, which guarantees the scientific quality and the operational usefulness of the programmes.

Seven research programmes have been conducted by the centre to date:
- Investment Solutions in Institutional and Individual Money Management;
- Equity Risk Premia in Investment Solutions;
- Fixed-Income Risk Premia in Investment Solutions;
- Alternative Risk Premia in Investment Solutions;
- Multi-Asset Multi-Factor Investment Solutions;
- Reporting and Regulation for Investment Solutions;
- Technology, Big Data and Artificial Intelligence for Investment Solutions.

EDHEC-Risk Institute's seven research programmes explore interrelated aspects of investment solutions to advance the frontiers of knowledge and foster industry innovation. They receive the support of a large number of financial companies. The results of the research programmes are disseminated through the EDHEC-Risk locations in the City of London (United Kingdom) and Nice, (France).

EDHEC-Risk has developed a close partnership with a small number of sponsors within the framework of research chairs or major research projects:
- **Financial Risk Management as a Source of Performance,**
in partnership with the *French Asset Management Association (Association Française de la Gestion financière – AFG)*;
- **ETF, Indexing and Smart Beta Investment Strategies,**
in partnership with *Amundi*;
- **Regulation and Institutional Investment,**
in partnership with *AXA Investment Managers*;
- **Optimising Bond Portfolios,**
in partnership with *BDF Gestion*;
- **Asset-Liability Management and Institutional Investment Management,**
in partnership with *BNP Paribas Investment Partners*;
- **New Frontiers in Risk Assessment and Performance Reporting,**
in partnership with *CACEIS*;
- **Exploring the Commodity Futures Risk Premium: Implications for Asset Allocation and Regulation,**
in partnership with *CME Group*;

About EDHEC-Risk Institute

• **Asset-Liability Management Techniques for Sovereign Wealth Fund Management,**
in partnership with *Deutsche Bank*;
• **The Benefits of Volatility Derivatives in Equity Portfolio Management,**
in partnership with *Eurex*;
• **Innovations and Regulations in Investment Banking,**
in partnership with the *French Banking Federation (FBF)*;
• **Dynamic Allocation Models and New Forms of Target-Date Funds for Private and Institutional Clients,**
in partnership with *La Française AM*;
• **Risk Allocation Solutions,**
in partnership with *Lyxor Asset Management*;
• **Infrastructure Equity Investment Management and Benchmarking,**
in partnership with *Meridiam and Campbell Lutyens*;
• **Risk Allocation Framework for Goal-Driven Investing Strategies,**
in partnership with *Merrill Lynch Wealth Management*;
• **Financial Engineering and Global Alternative Portfolios for Institutional Investors,**
in partnership with *Morgan Stanley Investment Management*;
• **Investment and Governance Characteristics of Infrastructure Debt Investments,**
in partnership with *Natixis*;
• **Advanced Investment Solutions for Liability Hedging for Inflation Risk,**
in partnership with *Ontario Teachers*' Pension Plan;
• **Cross-Sectional and Time-Series Estimates of Risk Premia in Bond Markets,**
in partnership with *PIMCO*;
• **Active Allocation to Smart Factor Indices,**
in partnership with *Rothschild & Cie*;
• **Solvency II,**
in partnership with *Russell Investments*;
• **Advanced Modelling for Alternative Investments,**
in partnership with *Société Générale Prime Services (Newedge)*;
• **Structured Equity Investment Strategies for Long-Term Asian Investors,**
in partnership with *Société Générale Corporate & Investment Banking*.

The philosophy of the Institute is to validate its work by publication in international academic journals, as well as to make it available to the sector through its position papers, published studies and global conferences.

To ensure the distribution of its research to the industry, EDHEC-Risk also provides professionals with access to its website, https://risk.edhec.edu, which is devoted to international risk and investment management research for the industry. The website is aimed at professionals who wish to benefit from EDHEC-Risk's analysis and expertise in the area of investment solutions. Its quarterly newsletter is distributed to more than 150,000 readers.

About EDHEC-Risk Institute

Research for Business

EDHEC-Risk Institute also has highly significant executive education activities for professionals, in partnership with prestigious academic partners. EDHEC-Risk's executive education programmes help investment professionals upgrade their skills with advanced asset allocation and risk management training across traditional and alternative classes.

In 2012, EDHEC-Risk Institute signed two strategic partnership agreements. The first was with the Operations Research and Financial Engineering department of Princeton University to set up a joint research programme in the area of investment solutions for institutions and individuals. The second was with Yale School of Management to set up joint certified executive training courses in North America and Europe in the area of risk and investment management.

As part of its policy of transferring know-how to the industry, in 2013 EDHEC-Risk Institute also set up Scientific Beta, which is an original initiative that aims to favour the adoption of the latest advances in smart beta design and implementation by the whole investment industry. Its academic origin provides the foundation for its strategy: offer, in the best economic conditions possible, the smart beta solutions that are most proven scientifically with full transparency in both the methods and the associated risks.

EDHEC-Risk Institute also contributed to the 2016 launch of EDHEC Infrastructure Institute (EDHEC*infra*), a spin-off dedicated to benchmarking private infrastructure investments. EDHEC*infra* was created to address the profound knowledge gap faced by infrastructure investors by collecting and standardising private investment and cash flow data and running state-of-the-art asset pricing and risk models to create the performance benchmarks that are needed for asset allocation, prudential regulation and the design of infrastructure investment solutions.

About EDHEC-Risk Institute

EDHEC-Risk Institute Publications and Position Papers (2017-2020)

EDHEC-Risk Institute Publications and Position Papers (2017-2020)

2020

- Martellini, L. and V. Milhau. Factor Investing in Liability-Driven and Goal-Based Investment Solutions (March).

2019

- Le Sourd, V. and L. Martellini. The EDHEC European ETF and Smart Beta and Factor Investing Survey 2019 (August).
- Maeso, J.M., Martellini, L. and R. Rebonato. Cross-Sectional and Time-Series Momentum in Sovereign Bond Markets (May).
- Maeso, J.M., Martellini, L. and R. Rebonato. Defining and Expoiting Value in Sovereign Bond Market (May).
- Maeso, J.M., Martellini, L. and R. Rebonato. Factor Investing in Sovereign Bond Markets - Time-Series Perspective (May).

2018

- Goltz, F. and V. Le Sourd. The EDHEC European ETF and Smart Beta and Factor Investing Survey 2018 (August).
- Mantilla-Garcia, D. Maximising the Volatility Return: A Risk-Based Strategy for Homogeneous Groups of Assets (June).
- Giron, K., L. Martellini, V. Milhau, J. Mulvey and A. Suri. Applying Goal-Based Investing Principles to the Retirement Problem (May).
- Martellini, L. and V. Milhau. Smart Beta and Beyond: Maximising the Benefits of Factor Investing (February).

2017

- Amenc, N., F. Goltz, V. Le Sourd. EDHEC Survey on Equity Factor Investing (November).
- Amenc, N., F. Goltz, V. Le Sourd. The EDHEC European ETF and Smart Beta Survey 2016 (May).
- Maeso, J.M., Martellini, L. Maximising an Equity Portfolio Excess Growth Rate: A New Form of Smart Beta Strategy? (November).
- Martellini, L. and V. Milhau. Mass Customisation versus Mass Production in Retirement Investment Management. Addressing a “Tough Engineering Problem” (May).
- Esakia, M., F. Goltz, S. Sivasubramanian and J. Ulahel. Smart Beta Replication Costs (February).
- Maeso, J.M., Martellini, L. Measuring Volatility Pumping Benefits in Equity Markets (February).

Notes

Notes

For more information, please contact:
Maud Gauchon on +33 (0)4 93 18 78 87
or by e-mail to: maud.gauchon@edhec-risk.com

EDHEC-Risk Institute
393 promenade des Anglais
BP 3116 - 06202 Nice Cedex 3
France
Tel. +33 (0)4 93 18 78 87

EDHEC Risk Institute—Europe
10 Fleet Place, Ludgate
London EC4M 7RB
United Kingdom
Tel: + 44 207 332 5600

risk.edhec.edu