EDHEC RISK AND ASSET MANAGEMENT RESEARCH CENTRE
Edhec -1090 route des crêtes - 06560 Valbonne - Tel. +33 (0)4 92 96 89 50 - Fax. +33 (0)4 92 96 93 22
Email: research@edhec-risk.com – Web: www.edhec-risk.com

Revisiting Core-Satellite Investing - A Dynamic Model of Relative Risk Management[i]

Noël Amenc
Professor of Finance, EDHEC Graduate School of Business
Director, EDHEC Risk and Asset Management Research Centre

Philippe Malaise
Professor of Finance, EDHEC Graduate School of Business
Associate Researcher, EDHEC Risk and Asset Management Research Centre

Lionel Martellini
Professor of Finance, EDHEC Graduate School of Business
Scientific Director, EDHEC Risk and Asset Management Research Centre

EDHEC
BUSINESS SCHOOL
Lille - Nice

Abstract

Tracking error is not necessarily bad. Just like with good and bad cholesterol, there is “good” tracking error, which refers to out-performance of a portfolio with respect to the benchmark, and “bad” tracking error, which refers to underperformance with respect to the benchmark. By severely restricting the amounts invested in active strategies as a result of tight tracking error constraints, investors forgive an opportunity for significant out-performance, especially during market downturns. In this paper, we introduce a new methodology that allows investors to gain full access to good tracking error, while maintaining the level of bad tracking error below a given threshold, based on an optimal dynamic adjustment of the fractions invested in a passive core versus an active satellite portfolio. This method can be regarded as a natural extension of constant proportion portfolio insurance techniques, originally designed to ensure the respect of absolute performance, to a relative return context.

Recent difficulties have drawn attention to the risk management practices of institutional investors in general, and pension plans in particular. A perfect storm of adverse market conditions over the past three years has devastated many corporate defined benefit pension plans. Negative equity market returns have eroded plan assets at the same time as declining interest rates have increased benefit obligations. In extreme cases, this has left corporate pension plans with funding gaps as large as or larger than the market capitalization of the plan sponsor. For example, the FTSE 100 companies faced a cumulative deficit of 55 billion GBP in 2003 (Standard Life Investments (2003)), while the worldwide deficit reached an estimated 1,500 to 2,000 billion USD (Watson Wyatt (2003)). These events have spotlighted the weakness of current funding standards for corporate defined benefit pension plans. They have also emphasized the weakness of investment practices.

Driven by the desire to improve investment efficiency, a growing number of institutional investors have moved to a core-satellite approach to portfolio management over the past several years. Core-satellite investment breaks down a portfolio into a core mostly passive component that is managed by a mainstream manager and satellite active components that are allocated to less efficient markets requiring specialist skills.[ii] The purpose of the core component is to control manager risk and to improve the efficiency of the overall portfolio by limiting costs. The role of the active components, meanwhile, is to provide diversification and to generate out-performance.

The move towards core-satellite management has brought with it some key changes in the asset management industry, such as the increased demand for high alpha products such as hedge funds, pursuing absolute performance strategies in the absence of tight tracking error constraints. Whether it is an adequate answer to the challenges met by institutional investors remains to be seen. On the one hand, there is no doubt that the core-satellite approach to portfolio management has proven to be a cost-efficient way to control a portfolio *relative* risk (also known as *tracking error* risk), i.e., the risk of deviation from a given benchmark, either a commercial index or a liability-consistent customized benchmark. On the other hand, it does not address the challenge of managing a portfolio *absolute* risk, understood as the risk of a major decrease in the core portfolio value. With a dominant fraction of the portfolio invested in the core component, asset value will be dramatically impacted by severe market downturns such as the ones experienced over the last three years.

By severely restricting the amounts invested in active strategies as a result of tight tracking error constraints, investors actually forgive an opportunity for return enhancement and risk reduction, especially during market downturns when active absolute return strategies are expected to out-perform passive investments in equity. Tracking error is not necessarily bad. Just like with good and bad cholesterol, there is “good” tracking error, which refers to out-performance of a portfolio with respect to the benchmark, and “bad” tracking error, which refers to underperformance with respect to the benchmark.

In this paper, we introduce a new methodology that allows investors to gain full access to good tracking error, while maintaining the level of bad tracking error at a reasonable threshold, based on an optimal dynamic adjustment of the fractions invested in core versus satellite portfolio. This method can be regarded as a natural extension of constant proportion portfolio insurance techniques, originally designed to ensure the respect of absolute performance, to a relative return context. Our approach is also somewhat reminiscent of contingent immunization that was introduced by Leibowitz and Weinberger (1982, 1983) in the context of duration gap management. Our contribution is to show that the relevance of the method applies more generally than to a mere asset-liability management context.

The rest of the paper is organized as follows. We first review the standard, static, approach to core-satellite investing, before introducing the dynamic version, allowing investors to implement a dissymmetric control of relative risk. We then present an example of application of the method in the context of a bond portfolio investment strategy.

STANDARD STATIC CORE-SATELLITE APPROACH

Most active managers still have dominant passive exposure to their benchmark. Instead of paying high fees on the passively managed part of their portfolio, the core-satellite approach suggests passively investing in a low-fee index fund (or an enhanced index product) as a core portfolio and in a variety of satellite active managers with higher tracking error. Ultimately, investors may want to invest in market-neutral managers who provide only portable alpha benefits without passive exposure to the index, so that they only compensate active managers for their *abnormal* returns, not for their passive exposure to rewarded sources of risk.

The Arithmetic of Core-Satellite Portfolio Management

We first consider a core-satellite approach with a single satellite portfolio. We show how to derive the optimal proportion to invest in satellite versus core portfolio by setting the problem in a simple mean-variance analysis. We also demonstrate that, if the core portfolio perfectly replicates the benchmark, then the information ratio of the overall portfolio is independent of the proportion in core versus satellite and equal to the information ratio of the satellite portfolio.

We first consider a core-satellite approach with a single satellite portfolio. The mathematics of a core-satellite approach is then straightforward. The overall portfolio corresponds to: $P = wS + (1-w)C$, where w is the fraction invested in the satellite (S), and *1-w* is the fraction invested in the core (C). We now calculate the tracking error with respect to a benchmark B. We have: $P - B = wS + (1-w)C - B = w(S-B) + (1-w)(C-B)$. If we now assume for simplicity that the core portfolio is perfectly replicating the benchmark, we have *C=B*, then we have: $P - B = w(S-B)$. As a result, we obtain that $TE(P) = \sqrt{\mathrm{var}(P-B)} = w\sqrt{\mathrm{var}(S-B)} = wTE(S)$.

Let us consider the following example. We assume an investor has a target level of risk relative to a given benchmark, such as a 2.5% tracking error budget. Two options are possible. Either the investor hires one manager with a tracking error equal to 2.5% for the entire portfolio, or the investor forms a passive core portfolio and leaves 20% in an aggressively managed satellite with a $12.5\% = \frac{TE(P)}{w} = \frac{2.5\%}{20\%}$ tracking error.

The next step consists of deriving the optimal proportion w^* to invest in satellite versus core portfolio. We solve the problem in the context of a simple mean-variance analysis.

The optimization program reads: $U = E(P-B) - \lambda\sigma^2(P-B) = IR(P)\times TE(P) - \lambda TE^2(P)$, where *IR(P)* is the information ratio of the portfolio P with respect to the benchmark, i.e., $IR(P) = \frac{E(P-B)}{\sigma(P-B)} = \frac{E(P-B)}{TE(P)}$ (see for example Grinold and Kahn (2000)).

It should be noted that when the core portfolio perfectly replicates the benchmark, the information ratio of the overall portfolio *IR(P)* is actually independent of the proportion in core versus satellite and equal to the information ratio of the satellite portfolio *IR(S)* (as long as the proportion *w* is strictly positive). This can easily be seen from:

$$IR(P) = \frac{E(wS + (1-w)C - B)}{\sigma(wS + (1-w)C - B)} = \frac{wE(S-B)}{wTE(S)} = IR(S)$$

We may rewrite the optimization program as: $U(w) = IR \times w \times TE(S) - \lambda w^2 TE^2(S)$, and the first-order condition reads: $\frac{\partial U}{\partial w}(w^*) = 0 \Rightarrow w^* = \frac{IR}{2\lambda TE(S)}$.

For example, let us assume that the tracking error of the active fund is 5%, that the Information Ratio (*IR*) is 0.5, and that the coefficient of risk-aversion with respect to relative risk is $\lambda = 0.2$. Then, the optimal proportion invested in the active portfolio is:
$w^* = \frac{IR}{2\lambda TE(S)} = \frac{.5}{2\times .2\times 5\%} = 25\%$
The resulting tracking error is $TE(P) = 25\% \times 5\% = 1.25\%$.

Extending the analysis to the case of a satellite $S = \sum_{i=1}^{n} w_i S_i$ invested in a number *n* of active portfolio managers S_i according to the proportions w_i is straightforward. The excess return on the satellite portfolio is then $S - B = \sum_{i=1}^{n} w_i (S_i - B)$, and the tracking error of the satellite portfolio reads $TE(S) = \left(\sum_{i,j=1}^{N} w_i w_j \sigma_{ij} - 2\sum_{i=1}^{N} w_i \sigma_{iB} + \sigma_B^2\right)^{1/2}$, where σ_{ij} is the covariance between portfolio managers S_i and S_j, and σ_B is the volatility of the benchmark.

One can then find the optimal fraction invested in each active manager within the satellite portfolio so as to achieve the highest possible Information Ratio. One can show (see for example Scherer (2002)) that the optimal condition is that the ratio of return to risk contribution is the same for all managers, which reads:

$$\frac{w_k \alpha_k}{\left(w_k^2 \sigma_{\alpha_k}^2 + \sum_j w_k w_j \sigma_{kj}\right) \Big/ TE(S)} = \frac{w_l \alpha_l}{\left(w_l^2 \sigma_{\alpha_l}^2 + \sum_j w_l w_j \sigma_{kj}\right) \Big/ TE(S)}$$

As a conclusion, a satellite/core approach seems to be perfectly suited for investors who attempt to use hedge funds to add portable alpha benefits to their long-only portfolio without modifying their passive exposure to a reference index, as it allows for a separate control on the tracking error of the satellite and core portfolios, so as to ensure that the overall portfolio is consistent with a target level of deviation with respect to the chosen benchmark.

The Economics of Core-Satellite Portfolio Management

We emphasize that a core-satellite portfolio approach can be used as an effective strategy for institutions that want to diversify their portfolios without giving up the potential for higher returns generated by selected active management strategies.

Exhibit 1 illustrates how the core-satellite approach provides the framework for targeting and controlling those areas where investors are willing to take more risk in a cost-efficient manner.

Exhibit 1: The economics of core-satellite management.

	"Core"	"Satellite"	Global
Weight	75%	25%	**100%**
Tracking Error	0%	20%	**5%** (0%x0.75 + 20%x0.25)
Management Fees	16bps	40bps	**22bps** (16*0.75+40x0.25)

Let us assume that an investor has a relative risk tracking error budget equal to 5%. The first solution is to allocate 100% of the portfolio to an active manager who will commit to respecting the 5% tracking error constraint. We assume that the level of management fees charged by this active manager is, say, 40 basis points. The second solution consists of allocating 75% of the portfolio to a purely passive product, e.g., an Exchange Traded Fund (ETF), and 25% of the portfolio to a 20% tracking error manager. This latter solution, which is consistent with a core-satellite approach to active asset management, offers two benefits.

First, allowing the active manager to significantly deviate from the benchmark leads to a better use of the manager's skills. If the manager has reliable views on market trends and directions, a 5% tracking error constraint leaves him with too little room for implementing active decisions consistent with these views. Beating the market is a notoriously tough game to play. Playing it with one hand tied behind one's back does not seem to be a good starting point!

The other benefit of the core-satellite approach is obvious in terms of cost control. Assuming a realistic 16 basis points fee structure on the core portfolio passively invested in an Exchange Traded Fund, we obtain a total level of fees equal to 22 basis points. In an increasingly competitive environment, such a difference in fee structure is likely to prove decisive.

EXTENDING THE APPROACH TO DYNAMIC CORE-SATELLITE MANAGEMENT

In this section, we show how to extend the approach to a dynamic context, where we let the proportion invested in the active portfolio vary as a function of the current cumulative out-performance of the global portfolio with respect to the benchmark. To achieve this objective, we transport the traditional constant proportion portfolio insurance method (CPPI) to the context of core-satellite portfolio management, so as to allow for a more efficient relative risk control.

Standard CPPI Approach for Absolute Risk Management

Introduced by Black and Jones (1987) and Black and Perold (1992), the CPPI procedure allows the production of option-like positions through systematic trading rules. This procedure dynamically allocates total assets to a risky asset in proportion to a multiple of the *cushion*, i.e., the difference between current wealth and a desired protective floor. This produces an effect similar to owning a put option. Under such a strategy, the portfolio's exposure tends to zero as the cushion approaches zero; when the cushion is zero, the portfolio is completely invested in cash. Thus, in theory, the guarantee is perfect: the strategy of exposure ensures that the portfolio never descends below the floor; in the event that it touches the floor, the fund is "dead" - it can deliver no performance beyond the guarantee.[iii]

The method is perhaps best understood through an illustration (see Exhibit 2). Suppose that a portfolio manager seeks to guarantee an investor a floor equal to 90% of an initial investment normalized at 100, so that the cushion is equal to portfolio value-floor=10 at the initial date T0. Assume also for simplicity of exposure that the risk-free rate is zero. To achieve this goal, the manager uses a multiplier equal to 4. By definition, the fraction invested by the manager in the risky asset is 10×4=40, which leaves 100-40=60 in cash. At date T1, the value of the risky asset has increased by 10%, so that the cushion is now equal to 14. The fraction invested by the manager in the risky asset is now 14×4=56, which leaves 104-56=48 in cash. At date T2, the value of the risk asset has decreased by 9.1%, and goes from 56 to 50.91. The portfolio value is then equal to 50.91+48=98.91, and the cushion is now 98.91-90=8.91. The position in risky asset needs to be adjusted back to 8.91×4=35.64, which leaves 98.91-35.64=63.27 in cash.

Exhibit 2: An example of CPPI

Multiplier = 4

	T0				T1				T2
Market index	**100.00**	+ 10%	110.00		**110.00**	-9.1%	100.00		**100.00**
Portfolio value	**100.00**		104.00		**104.00**		98.91		**98.91**
Risky asset	40.00	→	44.00	→	56.00	→	50.91	→	35.64
				Adjustments				Adjustments	
Cash	60.00		60.00	→	48.00		48.00	→	63.27
Floor	90.00				90.00				90.00
Cushion	10.00				14.00				8.91
Max Exposure	40.00				56.00				35.64

Introducing a Relative Risk Approach CPPI

In this section, we introduce a new methodology that allows investors to gain full access to good tracking error, while maintaining the level of bad tracking error at a reasonable threshold, based on an optimal dynamic adjustment of the fractions invested in core versus satellite portfolio. This method, which can be seen as a structured form of active management, is a natural extension of constant proportion portfolio insurance techniques, originally designed to ensure the respect of absolute performance, to a relative return context.
An intuitively interesting example of a similar dynamic strategy that is specific to pension funds depends upon the funded status of the plan. Even if the fiduciary's views on returns and interest rates do not change over time, it may still make sense to adjust the asset allocation and duration of the portfolio dynamically as the funded status of the plan changes. There may be strong disincentives for allowing the funded status of the plan to fall short of certain thresholds (e.g., fully funded). Balancing this shortfall risk against return would require the fiduciary to make adjustments depending upon funded status. For example, when a plan is very much over-funded, a high allocation in equities, or more generally in an asset class showing low correlation with the liability-driven benchmark, would not incur much shortfall risk. However, when the plan is in danger of falling below one of the funded status thresholds, it may be more prudent to shift more heavily into fixed income.

In the example above (Exhibit 2), we considered a standard CPPI approach, where a guarantee is offered on the absolute level of performance, with a maximum loss limited to 10% of the amount initially invested. We now argue that a similar approach can be followed to offer the investor a guarantee on the *relative* level of performance, with a maximum underperformance with respect to the benchmark. This underperformance will be limited say to 10% of the amount initially invested.

The techniques of traditional CPPI still apply, provided that the risky asset is re-interpreted as the satellite portfolio, which contains relative risk with respect to the benchmark, and the risk-free asset is re-interpreted as the core portfolio, which contains no relative risk with respect to the benchmark (see Exhibit 3).

Exhibit 3: Traditional CPPI versus relative approach CPPI.

Traditional CPPI	Relative Approach CPPI
Risky Asset	Satellite Portfolio
Risk-free Asset	Core Portfolio

Let us consider an example similar to the one in Exhibit 2, but cast in a relative risk framework. We assume that the benchmark is a passive investment, e.g., in a bond index. The guarantee, which was set in the example above at 90 = initial amount invested – 10, is now set at benchmark value – 10. We again assume that the multiplier is equal to 4.

At the initial date T0, portfolio value and benchmark value are normalized at 100, with a floor set at benchmark value –10 =100–10=90. The cushion is therefore equal to 100-90=10. The investment in satellite is then 10×4=40, which results in 100-40=60 in the core. At date T1, let us assume that the difference between the satellite and the benchmark is +10%, resulting for example from the following scenario: S=0%, C=-10%.[iv] In this case, the position invested in the core has decreased by 10% from 60 down to 54. Besides, the active portfolio value has remained stable at 40, while the benchmark has also decreased by 10%, going from 100 to 90. The difference between the fund value (94=54+40) and the benchmark value (90) is now equal to 4, and the cushion is then 14=94-90+10. The new optimal fraction to invest in the satellite portfolio is 14×4=56, which leaves 94-56=38 in the core portfolio. Date T1 resulting allocation is therefore 56/94=59.58% in the satellite and 38/94=40.42% in the core portfolio.

If we assume, on the other hand, that the difference between the satellite and the benchmark is -10%, resulting for example from the following scenario: S=0%, C=+10%. In this case, the position invested in the core has increased by 10% from 60 up to 66. Besides, the active portfolio value has remained stable at 40, while the benchmark has also increased by 10%, going from 100 to 110. The difference between the fund value (106=66+40) and the benchmark value (110) is now equal to -4, and the cushion has decreased to reach 6=106-110+10. The new optimal fraction to invest in the satellite portfolio is 6×4=24, which leaves 106-24=82 in the core portfolio. Date T1 resulting allocation is therefore 24/106=22.64% in the satellite and 82/106=77.36% in the core portfolio.

As can be seen from this example, the method leads to an increase in the fraction allocated to the satellite (from 40% to 59.57% in the example) when the satellite has outperformed the benchmark. Indeed such an accumulation of past out-performance has resulted in an increase in the cushion, and therefore in the potential for a more aggressive (and hence higher tracking error) strategy in the future. If on the other hand the satellite has under-performed with respect to the benchmark, the method leads to a tighter tracking error strategy (through a decrease of the fraction invested in the satellite portfolio) in an attempt to ensure the guarantee of the relative performance objective.

This approach allows for dissymmetric management of tracking error, ensuring that the underperformance of the portfolio with respect to the benchmark will be limited to a given level, while letting the investor gain fuller access to excess returns potentially generated by the active portfolio.

The benefit of this approach are illustrated in Exhibit 4, which shows that a dynamic version of core-satellite approach allows an investor to truncate the relative return distribution so as to allocate the probability weights away from severe relative under-performance to the profit of more potential for out-performance.

When the active portfolio (satellite) is an absolute return market-neutral type of fund with low levels of average risk (volatility) and extreme risks (VaR), the method can be implemented with relatively high multiplier values, allowing the investor to have high exposure to the benefits of positive tracking error, while limiting the exposure to negative tracking error.

Exhibit 4: Distribution of Portfolio Excess Return with Respect to the Benchmark: Static versus Dynamic Core-Satellite Approach

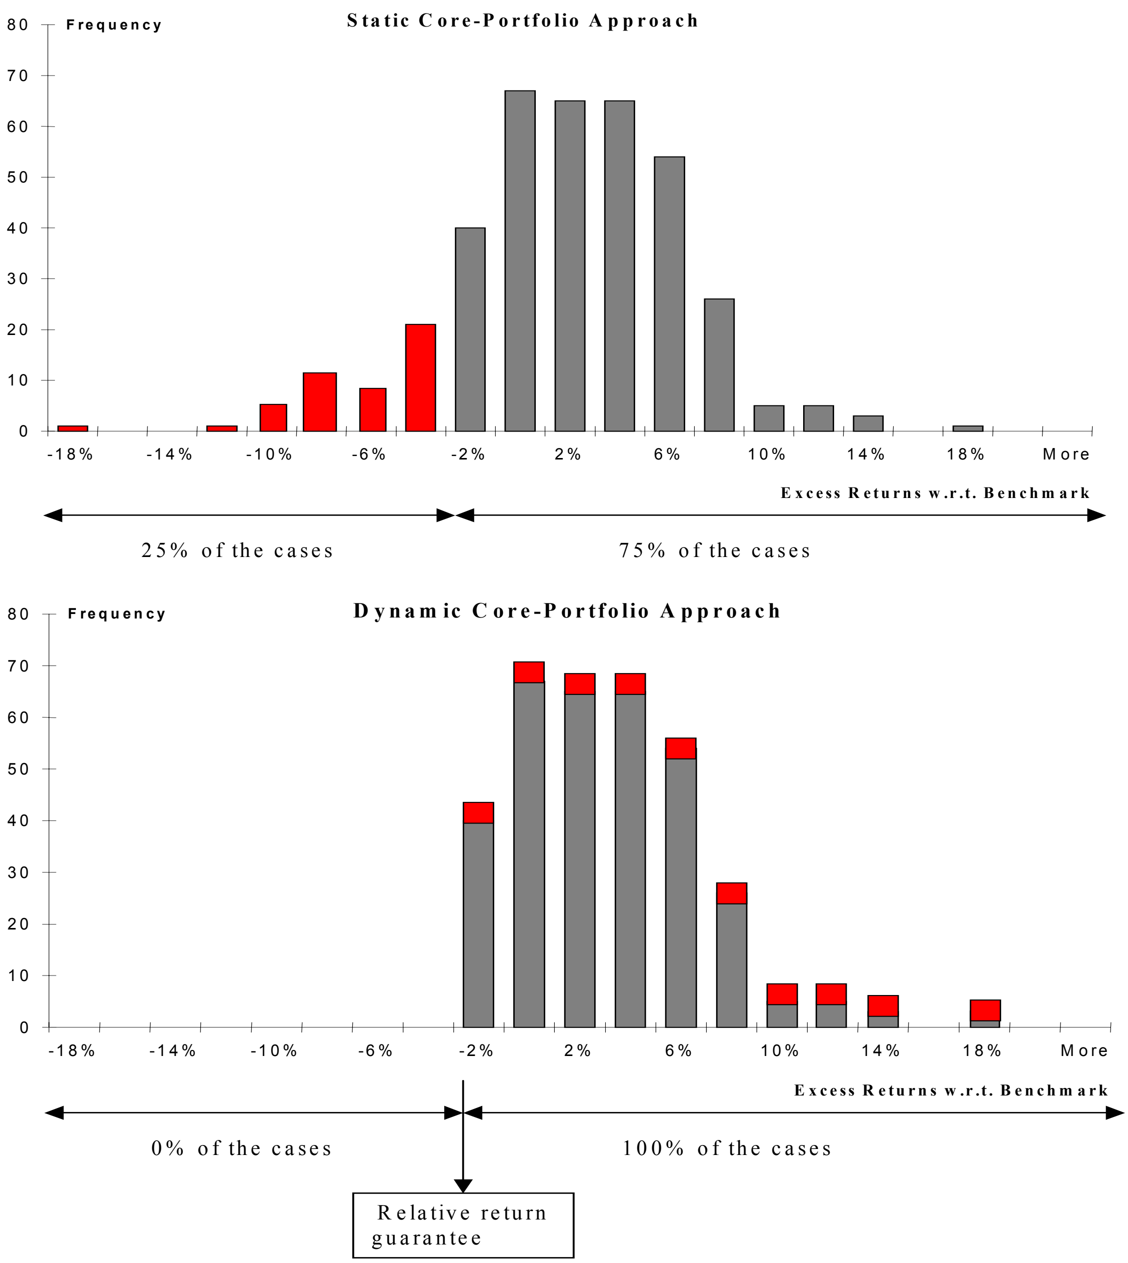

BACKTESTING THE METHOD

For the sake of illustration, we present in this section the results of a numerical experiment demonstrating the benefits of implementing a dynamic core-satellite approach in a context where a bond Exchange Traded Fund (ETF) is used as a core portfolio and a systematic maturity rotation strategy is used as an active satellite.

Introducing a Core and a Satellite

We have chosen to illustrate the method in the context of bond portfolio management.[v]

In an attempt to generate the performance of a possible candidate as a satellite portfolio, we consider the performance of a fictitious maturity style timer with various levels of forecasting ability. The motivation for maturity rotation strategies originates from the fact that different maturity indices strategies perform differently in different times and economic conditions, and that there is evidence of predictability in these patterns. For example Keim and Stambaugh (1986)) find that several ex-ante observable variables based on asset price levels predict, among other things, ex-post risk premiums on U.S. Government bonds of various maturities. Using multi-factor models for the return on bond indices, where the factors are chosen to measure the many dimensions of financial risks (market, volatility, credit and liquidity risks), one may be able to implement a strategy that generates abnormal return from timing between different segments of the yield curve (see for example Martellini, Priaulet and Priaulet (2003)).

In this example, we consider two European bond indices, the 3-5 year Euro-MTS Index and the 10-15 year Euro-MTS Index.[vi] We calculate the profit generated by investing 100% at the beginning of each month in the index with the highest return in the following month. This strategy can be implemented by dynamic trading in bond Exchange Traded Funds that are designed to passively replicate the performance of bond indices. More specifically, we show in Exhibit 5 the performance of a tactical style timer over the sample period ranging from January 1999 to December 2003. In this experiment, 100% of the portfolio is invested in the best performing index (3-5 or 10-15 years) with various degrees of predictive ability depicted by hit ratios ranging from 50% (no predictive ability) to 100% (perfect timer). The benchmark is an equally-weighted portfolio invested 50% in 3-5 year index and 50% in 10-15 year index, with rebalancing taking place at the beginning of each month to bring the allocation back to neutrality.

As can be seen from the results in Exhibit 5, we find that the performance of a style timer with perfect forecasting ability who invests 100% of a portfolio at the beginning of the year in the best performing style for the year generates an impressive information ratio equal to 4.8.

Exhibit 5: Performance of a tactical style timing strategy.

From Jan. 1999 to Dec. 2003	TSA Portfolio Risk & Return Analysis				Benchmark
	HR 100%	HR 75%	HR 65%	HR 50%	
Annualized Return	9.20%	6.48%	6.09%	4.69%	4.82%
Cumulative Return	57.52%	37.58%	34.91%	25.85%	26.72%
Cumulative Excess Return	30.80%	10.86%	8.20%	-0.86%	
Annualized Excess Return	4.38%	1.66%	1.27%	-0.13%	
Annualized Volatility	4.10%	4.04%	4.20%	4.06%	3.91%
Annualized Tracking Error	0.91%	1.49%	1.52%	1.57%	
Information Ratio	4.80	1.10	0.83	-0.08	

This Exhibit shows the performance of a timing portfolio on the period ranging from January 1999 to December 2003. In this experiment, 100% of the portfolio is invested in the best performing index (3-5 or 10-15 year) with various degrees of predictive ability depicted by hit ratios ranging from 50% (no predictive ability) to 100% (perfect timer). The benchmark is 50% invested in each index, with rebalancing taking place at the beginning of each month to bring the allocation back to neutrality.

Obviously, the assumption of perfect timing ability is not realistic. It can be argued that a realistic performance for a successful style timer is consistent with a hit ratio of around 65% (see for example Amenc et al. (2003)). Such a level of hit ratio allows for a 1.27% excess return and a 1.52% annual tracking error with respect to the equally-weighted benchmark, which results in a comfortable information ratio equal to 0.83 (see Grinold and Kahn (2000) for an empirical distribution of information ratios among active managers).

Exhibit 6: Performance distribution of a style timer with a 65% hit ratio.

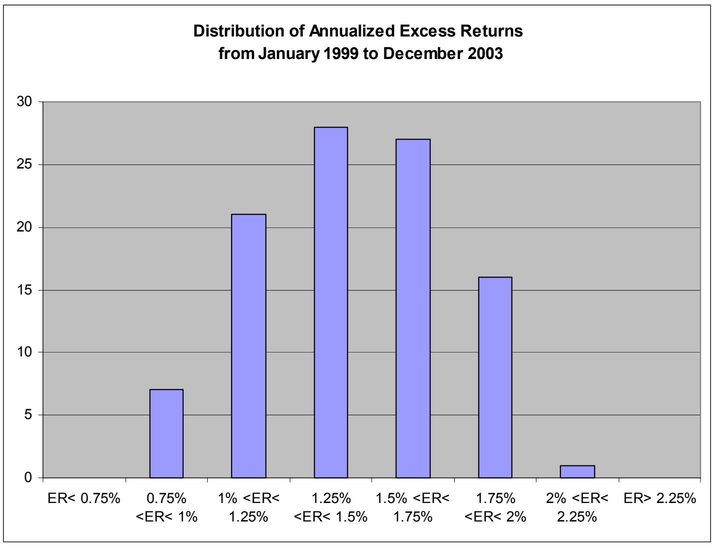

This Exhibit shows the distribution of performance obtained by a style timer with a 65% hit ratio as we let the months found to be successful vary across the 5x12 = 60 months in the entire sample.

To test the robustness of the results, we repeat the experiment 100 times by drawing the successful months randomly, while maintaining a 65% hit ratio level. Exhibit 6 below shows the distribution of performance obtained by a style timer, as we let the successful 65% of the months vary across the 5x12 = 60 months in the sample. As can be seen from Exhibit 6, the performance of the style timer is not a mere artifact of a particular choice of the winning months in the sample. This shows the robustness of abnormal performance that can be generated by a realistic timing strategy.

Not only can maturity rotation strategies be implemented in a long-only context, but they can also be used to generate absolute return benefits. Exhibit 7 below shows the performance of a strategy for the style timer with a 65% hit ratio, and long and short positions implemented with a market neutral exposure allowing 100% of initial capital to be invested in the risk-free rate (Eonia), so as to satisfy a level of leverage equal to 2. The resulting allocation is –50% in the index that is perceived as likely to underperform, and +50% in the index that is perceived as likely to outperform.[vii]

Exhibit 7: Absolute return approach.

Long/Short Positions	TSA Portfolio	EONIA
Annualized Return	6.09%	3.55%
Cumulative Return	35.19%	19.38%
Annualized Volatility	3.02%	0.27%

In this experiment, we focus on a 65% hit ratio, with long and short positions implemented with a € neutral exposure allowing 100% of initial capital to be invested in the risk-free rate (Eonia)

Exhibit 7 suggests that the benefits of maturity rotation strategies can be implemented in an absolute return approach, which allows for the portability of the abnormal performance to a core portfolio invested in a broad-based index. This is what we turn to next.

Base Case Experiment

In this section, we perform a numerical experiment demonstrating the benefit of implementing a dynamic core-satellite approach based on the global EuroMTS index as a core portfolio, and the absolute return maturity rotation strategy with 65% hit ratio as an active satellite.

For the sake of comparison, we first consider the performance of a standard, static, core-satellite portfolio approach, where the core portfolio is passively invested in the global Euro-MTS index, designed to be representative of the European Treasury bond market, and the satellite portfolio is based upon the absolute return timing strategy which performance is reported in Exhibit 7.

We perform several experiments, with an allocation to the satellite portfolio ranging from 5% to 40%,.the results of which are presented in Exhibit 8. From the analysis of the numbers in Exhibit 8, it appears that the benefits of the absolute performance maturity rotation strategies can be successfully transported to a core portfolio meant to reflect the strategic asset allocation of the investor.[viii]

Exhibit 8: Static Core-Satellite Portfolio Management.

Satellite Allocation	From Jan. 1999 to Dec. 2003	Overall Portfolio Risk & Return Analysis
10%	Annualized Return	4.82%
	Cumulative Return	26.90%
	Annualized Excess Return	0.14%
	Cumulative Excess Return	0.95%
	Annualized Volatility	3.05%
	Annualized Tracking Error	0.45%
	Information Ratio	0.31
20%	Annualized Return	4.96%
	Cumulative Return	27.85%
	Annualized Excess Return	0.28%
	Cumulative Excess Return	1.90%
	Annualized Volatility	2.77%
	Annualized Tracking Error	0.90%
	Information Ratio	0.31
30%	Annualized Return	5.10%
	Cumulative Return	28.79%
	Annualized Excess Return	0.42%
	Cumulative Excess Return	2.84%
	Annualized Volatility	2.54%
	Annualized Tracking Error	1.35%
	Information Ratio	0.31
40%	Annualized Return	5.24%
	Cumulative Return	29.72%
	Annualized Excess Return	0.56%
	Cumulative Excess Return	3.77%
	Annualized Volatility	2.37%
	Annualized Tracking Error	1.79%
	Information Ratio	0.31

This table shows the performance of a global core+satellite portfolio, where the core is passively invested in the Euro-MTS index, while the satellite is an active portfolio implementing a long-short maturity rotation strategy.

While we have presented here an experiment where the fraction of the portfolio allocated to the active satellite was as high as 40%, most investors are actually reluctant to increase the fraction of the portfolio invested in the satellite because of a concern over the tracking error of the overall portfolio. As recalled in the introduction, investors however forgive an opportunity for significant out-performance by severely restricting the amounts invested in active strategies as a result of tight tracking error constraints.

In what follows, we extend the approach to a dynamic context, where we let the proportion invested in the active portfolio vary as a function of the current out-performance of the global portfolio with respect to the benchmark. We show how the guarantee on the relative performance of the portfolio with respect to the benchmark bond index is performed. In particular, we show how the proportion invested in the satellite decreases as the active portfolio under-performs the benchmark. Conversely, we show how investors can gain higher exposure to the benefits of active portfolio management when the satellite outperforms the benchmark portfolio.

Exhibit 9 below contains the result of the experiment in the case of a guarantee equal to 95% of the benchmark, and an active portfolio invested in the long-short maturity rotation strategy with a hit ratio equal to 65% (see Exhibit 7). We assume that 100 million Euros are initially invested in the strategy, so that the guarantee is that the performance of the global portfolio shall never be more than 5 million Euros lower than that of the benchmark. We have tested 4 different values for the multiplier *m* (*m=2, 3, 4* or *5*).

Exhibit 9: Performance of relative return CPPI with a 95% guarantee.

Guarantee = 95% of Benchmark Performance	Initial cushion = 5 M€			
Multiplier	m = 2	m = 3	m = 4	m = 5
Satellite Initial Weighting	10%	15%	20%	25%
Cumulative Return	27.66%	28.58%	29.56%	30.56%
Annualized Return	4.95%	5.10%	5.25%	5.41%
Annualized Volatility	3.40%	3.44%	3.51%	3.60%
Annualized Excess Return	0.27%	0.42%	0.57%	0.73%
Annualized Excess Return (C+S static)	0.14%	0.21%	0.28%	0.35%
Annualized Tracking Error	0.35%	0.56%	0.79%	1.05%
Annualized Tracking Error (C+S static)	0.46%	0.67%	0.90%	1.12%
Information Ratio	0.77	0.75	0.72	0.70
Net Assets as from 12/31/2003	127,655,132	128,584,408	129,556,715	130,563,454
Guaranteed Value as of 12/31/2003	119,654,210	119,654,210	119,654,210	119,654,210
Difference	8,000,922	8,930,198	9,902,505	10,909,244

This Exhibit shows the result of the experiment in the case of a guarantee equal to 95% of the benchmark, and an active portfolio invested in the long-short maturity rotation strategy with a hit ratio equal to 65%. We assume that 100 million Euros are initially invested in the strategy, so that the guarantee is that the performance of the global portfolio shall never be more than 5 million Euros lower than that of the benchmark. We have tested 4 different values for the multiplier m (m=2, 3, 4 or 5). We present in line 7 the annual excess return with respect to the benchmark passively invested in the Euro-MTS index. In line 8 we show for comparison the excess return of a static core-portfolio alternative, with an investment in the satellite equal to the initial investment from the dynamic approach. The guaranteed value in line 13 is equal to the terminal value of the Euro-MTS index, from which we have subtracted 5 million Euros.

The performance of the method can perhaps be best understood by noting that the dynamic core-satellite approach always out-performs the comparable static counterpart (i.e., a static core-satellite portfolio with an investment in the satellite equal to the initial investment from the dynamic approach), with an excess return above the benchmark about twice as large in all cases. For example, we know (from Exhibit 8, also in line 7 of Exhibit 8) that the annualised excess return of a static core-portfolio strategy with 80% in the core and 20% in the satellite is 28 basis points. On the other hand, the comparable dynamic version of the strategy, with a 20% initial investment in the satellite corresponding to the case m=4, leads to a 57 basis points excess return. This is because the investor is allowed to increase the fraction allocated to the satellite when the performance of the active strategy allows him to do so, and therefore enjoys the benefits of a higher exposure to the satellite out-performance.

In Exhibit 10, we repeat a similar experiment in the case of a guarantee equal to 90% of the benchmark. As can be seen from Exhibit 10, even higher active benefits can be obtained if the guarantee is decreased. This is a standard trade-off between (relative) risk and return.

Exhibit 10: Performance of relative return CPPI with a 90% guarantee.

Guarantee = 90% of Benchmark Performance	Initial Cushion = 10 M€			
Multiplier	m = 2	m = 3	m = 4	m = 5
Satellite Initial Weighting	20%	30%	40%	50%
Cumulative Return	29.36%	31.22%	33.16%	35.18%
Annualized Return	5.22%	5.51%	5.81%	6.12%
Annualized Volatility	3.46%	3.59%	3.79%	4.07%
Annualized Excess Return	0.54%	0.83%	1.13%	1.44%
Annualized Excess Return (C+S static)	0.28%	0.42%	0.56%	0.70%
Annualized Tracking Error	0.69%	1.10%	1.56%	2.06%
Annualized Tracking Error (C+S static)	0.92%	1.35%	1.79%	2.24%
Information Ratio	0.78	0.75	0.73	0.70
Net Assets as from 12/31/2003	129,358,465	131,217,016	133,161,629	135,175,108
Guaranteed Value as of 12/31/2003	113,356,620	113,356,620	113,356,620	113,356,620
Difference	16,001,845	17,860,396	19,805,009	21,818,488

This Exhibit shows the result of the experiment in the case of a guarantee equal to 90% of the benchmark, and an active portfolio invested in the long-short maturity rotation strategy with a hit ratio equal to 65%. We assume that 100 million Euros are initially invested in the strategy, so that the guarantee is that the performance of the global portfolio shall never be more than 10 million Euros lower than that of the benchmark. We have tested 4 different values for the multiplier m (m=2, 3, 4 or 5). We present in line 7 the annual excess return with respect to the benchmark passively invested in the Euro-MTS index. In line 8 we show for comparison the excess return of a static core-portfolio alternative, with an investment in the satellite equal to the initial investment from the dynamic approach. The guaranteed value in line 13 is equal to the terminal value of the Euro-MTS index, from which we have subtracted 10 million Euros.

We now focus on one of these experiments (guarantee = 95% and m=5 from Exhibit 9), and we show in Exhibit 11 the relationship between the satellite performance and the fraction invested in it.

As can be seen from Exhibit 11, the method leads to an increase in the fraction allocated to the satellite (from 25% to about 45% in the example) following the fact that the satellite has outperformed the benchmark, as explained in the previous section.

Exhibit 11: Satellite performance and allocation.

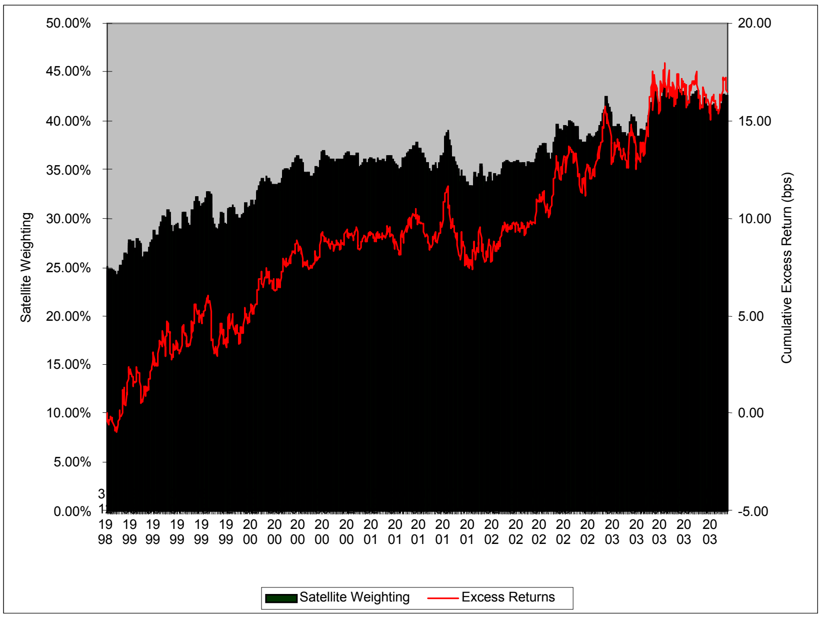

This Exhibit shows the relationship between the satellite performance and the fraction invested in it in the context of a guarantee equal to 95% and a multiplier equal to 5.

Introducing Transaction Costs

It is likely that investors will face constraints on the fraction of their portfolio invested in the active satellite that might prevent them from implementing dynamic strategies involving frequent changes in allocation to active managers. In particular, whether or not an institutional investor may terminate, increase or decrease the size of positions in a given manager depends on the specific situation and the bargaining power of the parties involved.

One other important ingredient that we have assumed away, and to which we turn now, is the presence of transaction costs. While dynamic changes in allocation can easily be performed when both the core and the satellite(s) are invested in liquid passive investment vehicles such as Exchange Traded Funds (ETFs), the presence of transaction costs is likely to affect the performance of a dynamic core-satellite portfolio process.

Under the relative version of a CPPI strategy introduced in this paper, the exposure to the satellite portfolio tends to zero as the cushion approaches zero; when the cushion is zero, the portfolio is completely invested in the index. Thus, in theory, the (relative) guarantee is perfect: the strategy of exposure ensures that the portfolio never descends below the floor. In the event that it touches the floor, the fund is locked in a fully passive investment - it can deliver no performance beyond the benchmark. In practice, since the rebalancing of the portfolio is done at discrete intervals (here every day) rather than continuously, there is a small risk of the portfolio crashing through the floor in between two consecutive rebalancements. This is what happened with some insured portfolios during the 1987 crash. In such a case, it is impossible even to meet the absolute or relative guarantee. Therefore, one objective of management might be to minimize this possibility. In the presence of transaction costs, there is a trade-off between risk management and cost of trading: frequent trading leads to a low likelihood of not being able to satisfy the relative performance guarantee, but it is costly. On the other hand, infrequent trading is cost-efficient, but introduces the risk of a large change in relative performance of the satellite with respect to the core portfolio, which may result in failure to meet the guaranteed relative return target.

The attractiveness of the method lies in the leverage effect of the multiplier. In theory, an optimal value for the multiplier can be endogenously derived (Grossman and Vila (1992)) as the solution to a specific utility maximization problem, as shown by Merton (1971). In practice, however, the multiplier is set according to the level of risk tolerance of the investor based upon the following principle. The higher the multiplier, the higher the exposure to the active portfolio, and the higher the gains in the case that the satellite outperforms the benchmark; conversely, the more significant the loss in the case of a poor performing active satellite. The portfolio therefore has a return that increases with the multiplier, but so does the risk. More specifically, a lower multiplier is recommended in case the tracking error between the satellite and the benchmark is prone to sudden and large jumps in value, so as to ensure that the cushion is preserved. An important quantity is the ratio, denoted as λ, between the cushion and the fraction invested in the satellite. In theory, that value is set to $1/m$, where m denotes the level of the multiplier. In practice, however, continuous trading can only maintain the parameter λ at a fixed level.

Given that the strategy under consideration is a dynamic portfolio strategy, the presence of transaction costs is likely to be a main concern to portfolio managers. We focus here on trading costs and consider the problem of optimal trading strategies in the presence of transaction costs that are paid each time a transaction is made. In the previous section, we assumed away the problem by considering a stylised situation with no transaction costs and continuous (i.e., daily here) rebalancing. We now account for the presence of transaction costs. We set the proportional transaction cost paid at each trade (brokerage fees and spread) to 10 basis points.[ix] Exhibit 12 shows the results of a relative return CPPI strategy in the presence of transaction costs.

Exhibit 12: Performance of relative return CPPI in the presence of transaction costs.

Multiplier	m = 2	m = 3	m = 4	m = 5
Satellite Initial Weighting	10%	15%	20%	25%
Cumulative Return	27.60 %	28.40 %	29.17 %	29.87 %
Annualized Return	4.94%	5.07%	5.19%	5.30%
Annualized Volatility	3.40%	3.44%	3.51%	3.62%
Annualized Excess Return	0.26%	0.39%	0.51%	0.62%
Annualized Tracking Error	0.35%	0.56%	0.79%	1.05%
Information Ratio	0.76	0.70	0.64	0.59

This Exhibit shows the result of the experiment in the case of a guarantee equal to 95% of the benchmark, and an active portfolio invested in the long-short maturity rotation strategy with a hit ratio equal to 65% (see Exhibit 7). We assume that 100 million Euros are initially invested in the strategy, so that the guarantee is that the performance of the global portfolio shall never be more than 5 million Euros lower than that of the benchmark. We have tested 4 different values for the multiplier m (m=2, 3, 4 or 5). We present in line 6 the annual excess return with respect to the benchmark passively invested in the Euro-MTS index.

As can be seen from the numbers in Exhibit 12, and the comparison with Exhibit 9, the presence of transaction costs has a significant impact on the performance of the method. When we focus on a situation where trading occurs every day, the presence of transaction costs naturally leads to a decrease in the information ratio, for example in the case of an m=3 multiplier, from 0.75 down to 0.70 (from 0.72 down to 0.64 if m=4).

It should however be noted that the experiments reported here overestimate the impact of transaction costs on the benefits of dynamic core-satellite portfolio management. This is because it is not optimal to trade at fixed time intervals (here every day). Much attention has been devoted to the question of optimal dynamic trading strategies in the academic literature, and different types of strategies have been considered. Seminal contributions to optimal multi-period investment decisions in the presence of stylized proportional transaction costs have been made in particular by Constantinides (1979, 1984, 1986, 1993) and Dumas and Luciano (1991) (see also Dixit (1991), Dumas (1991), Dumas and Luciano (1991)).[x] Optimal control strategies require that no hedge transaction should occur inside a "no-trade" interval, with rebalancing to the nearest edge when deviation with respect to the target brings the portfolio outside this no-trading zone.

In the present context, the prescription of optimal control strategies would be as follows. Instead of trading every day, an optimal strategy would consist of a scheme where rebalancing takes place if and only if the actual inverse of the multiplier goes outside a given no-trade region. Let us for example consider the case of a multiplier m=5. The inverse of the multiplier is 1/5=20%. This is the target value for λ, the ratio between the cushion and the fraction invested in the satellite. Consistent with theoretical prescriptions from the theory of optimal control, trading should take place so as to bring the inverse of the multiplier ratio back to the nearest edge of the no-trade interval, and not to the optimal level. If the size of the no-trade interval is say 2%, then no trading is performed if the actual ratio wanders anywhere between 18% and 22%. When the ratio goes outside this range, say from below, it is brought back to 18%. Note that there is a standard risk-return trade-off involved here: increasing the no-trade interval leads to a decrease in transaction costs on the one hand, but it also leads to a less efficient, more static, implementation of the method. Ultimately, if we let the size of the no-trade interval go to infinity, no transaction is ever performed and we get back to the static case. The size of the no-trade interval may also be set optimally so as to be consistent with the investor's risk tolerance.

CONCLUSION

In this paper, we have introduced a method that leads to an increase in the fraction allocated to the satellite when the satellite has outperformed the benchmark. The intuition is that an accumulation of past out-performance results in a safety margin in terms of tracking error control, and therefore in the potential for a more aggressive (and hence higher tracking error) strategy in the future. If on the other hand the satellite has under-performed with respect to the benchmark, the method leads to a tighter tracking error strategy (through a decrease of the fraction invested in the satellite portfolio) in an attempt to ensure the guarantee of the relative performance objective.

We have shows that this dynamic version of core-satellite approach allows an investor to truncate the relative return distribution so as to allocate the probability weights away from severe relative under-performance to the profit of more potential for out-performance. This approach allows for a dissymmetric management of tracking error, ensuring that the underperformance of the portfolio with respect to the benchmark will be limited to a given level, while letting the investor gain fuller access to excess returns potentially generated by the active portfolio

REFERENCES

Amenc, N., P. Malaise, L. Martellini and D. Sfeir, 2003, Portable alpha and portable beta strategies in the Euro zone – Implementing active asset allocation decisions using index futures and options, *Journal of Portfolio Management*, forthcoming.

Black, F and Jones, R., 1987, Simplifying portfolio insurance, *Journal of Portfolio Management*, Fall, p. 48-51.

Black, F., and A. Perold, 1992, Theory of constant proportion portfolio insurance, *Journal of Economic Dynamics and Control*, 16, 403-426.

Constantinides, G., M., 1979, Multiperiod consumption and investment behavior with convex transaction costs, *Management Science*, 25, 1127-1137.

Constantinides, G. M., 1984, Optimal stock trading with personal taxes, *Journal of Finance*, 15, 676-713.

Constantinides, G. M., 1986, Capital market equilibrium with transaction costs, *Journal of Political Economy*, 94, 842-862.

Constantinides, G. M., 1993, Option pricing bonds with transaction costs, working paper, University of Chicago.

Dixit, A., 1991, A simplified treatment of some results concerning regulated Brownian motion, *Journal of Economic Dynamics and Control*, 15, 657-674.

Dumas, B., 1991, Super contact and related optimality conditions, *Journal of Economic Dynamics and Control*, 15, 675-686.

Dumas, B., and E. Luciano, 1991, An exact solution to a dynamic portfolio choice problem under transaction costs, *Journal of Finance*, 46, 577-595.

Grinold, R., and R. Kahn, 2000, Active portfolio management: a quantitative approach to providing superior returns and controlling risk, Second Edition, McGraw Hill.

Keim, D., and R. Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics*, 17, 357-390.

Grinold, R., and R. Kahn, 2000, Active portfolio management: a quantitative approach to providing superior returns and controlling risk, Second Edition, McGraw Hill.

Grossman, S., and Vila, J-L., 1992, Optimal dynamic trading with leverage constraints, *Journal of Financial and Quantitative Analysis*, 27, 2, 151-168.

Hirt, G., and J. C. Singleton (eds.), 2004, The core-satellite approach to portfolio management, , McGraw Hill, forthcoming book

Leibowitz, M., and A., Weinberger, 1982, Contingent immunization—part I: risk control procedures, *Financial Analysts Journal* (November–December), 17–31.

Leibowitz, M., and A., Weinberger. 1982, Contingent immunization—part II: problem areas, *Financial Analysts Journal* (January–February), 35–50.

Leland, H. E., 1999, Optimal portfolio management with transaction costs and capital gains taxes, working paper, Haas School of Business, UC Berkeley.

Martellini, L., and P. Priaulet, and S. Priaulet, 2003, Fixed-income securities – Valuation, risk management and portfolio strategies, *John Wiley*, 2003.

Merton, R., 1971, Optimum consumption and portfolio rules in a continuous time model, *Journal of Economic Theory*, 3, 373-413.

Scherer, B., 2002, Portfolio Construction and Risk Budgeting, RiskBooks.

Standard Life Investments, 2003, Bridging the pensions gap, Global Bytes, http://uk.standardlifeinvestments.com/content/strategy/strategy_index.html.

Watson Wyatt, 2003, Global asset study (ongoing); as cited by "Finanz und Wirtschaft" (28/01/2004), http://www.finanzinfo.ch.

ENDNOTES

---

[i] This research has been sponsored by Euronext and by the EDHEC Risk and Asset Management Research Centre.

[ii] See Hirt and Singleton (2004) for more details.

[iii] In practice, since the rebalancing of the portfolio is done at discrete intervals rather than continuously, there is a small risk of the portfolio crashing through the floor between two rebalances, as happened with some insured portfolios during the 1987 crash. In such a case, it is even impossible to meet the guarantee. Therefore, one objective of management might be to minimize this possibility.

[iv] Because only relative performance matters, any other scenario consistent with a 10% outperformance of the satellite would lead to the same strategy.

[v] We have also conducted similar experiences in the context of equity portfolios, which we hold available to readers upon request.

[vi] The Euro MTS indices model the total-return of the Euro zone bond market and currently cover the sovereign sector. They are calculated each day both in real-time and with two daily

fixings using prices from the MTS system. The Euro-MTS Index was formerly the CNO Etrix. The Euro-MTS Index reproduces the performance of the Euro zone government bond market by modeling the performance of a limited portfolio of bonds chosen to represent the performance of the wider market. The contribution of each representative bond's price and coupon accrual to the final index is determined by its issuer's size relative to the entire Euro zone government bond market.

[vii] It should be noted that, while the satellite's gross leverage (i.e. sum of absolute values of long and short positions) is equal to 2, the net leverage is equal to 0. Mixing long and short positions on bond ETFs allows investors to neutralize their exposure to interest rate risk. More accurately, when an investor is implementing such a strategy, the investor stays exposed to changes in the slope, as opposed to changes in the level, of the term structure of interest rates.

[viii] The results reported in Exhibit 8 allow us to confirm that the information ratio is independent of the fraction invested in the core versus the satellite portfolio, as was shown in full generality in a previous section.

[ix] This working hypothesis is consistent with transaction fees currently applying to substantial trades. Note that we assume away, for simplicity, a fixed-cost component that can obviously prove harmful for small trades.

[x] More recently, Leland (1999) extends this work in a slightly different setting to account for the effects of capital gains taxes.