April 4, 2012

Evaluating the Trade-Off between Hedge-Fund Returns & Peak-to-Valley

Back-Test Results on **Stress-Var** and ex-ante **Volatility**

Executive Summary

This Riskdata Research note examines the relationship between Hedge-Fund Returns and the Peak-To-Valley (*P2V.*)The study was implemented using the Portfolio Designer - an optimization algorithm developed by Riskdata – and applied to 100 random allocations on 359 Hedge funds reporting to *Hedge Fund Research* (*www.hedgefundresearch.com*). The simulations followed a rigorous “out-of-sample” protocol spanning the last 10 years. The goal was to reduce the Peak-to-Valley by 2 through two risk mitigation strategies: one based on ex-ante Volatility (*eaVol*), and the other using the Stress VaR (*SVaR*). The results were then adjusted to account for mortality bias.

We arrived at the following statistically-significant conclusions:

- Both eaVol and SVaR are effective tools for monitoring the P2V. However, SVaR is considerably more reliable than eaVol.
- The net impact of risk mitigation on performances is second order when using the SVaR, equivalent to a deleveraging policy if using the eaVol
- We estimate that over the past 10 years, a risk budgeting policy, that targeted a pre-defined level of P2V, produced an average post-fees excess return of 1.68% / year if based on the SVaR strategy, vs. no benefit based on the eaVol Strategy.
- This difference can be explained by the fact that returns have a high sensitivity to volatility – the likely consequence of 30 years of Sharpe ratio paradigm – while they still have low sensitivity to Peak-to-Valley – as markets currently mis-estimate the trade-off between returns and extreme risk, offering important arbitrage opportunities.

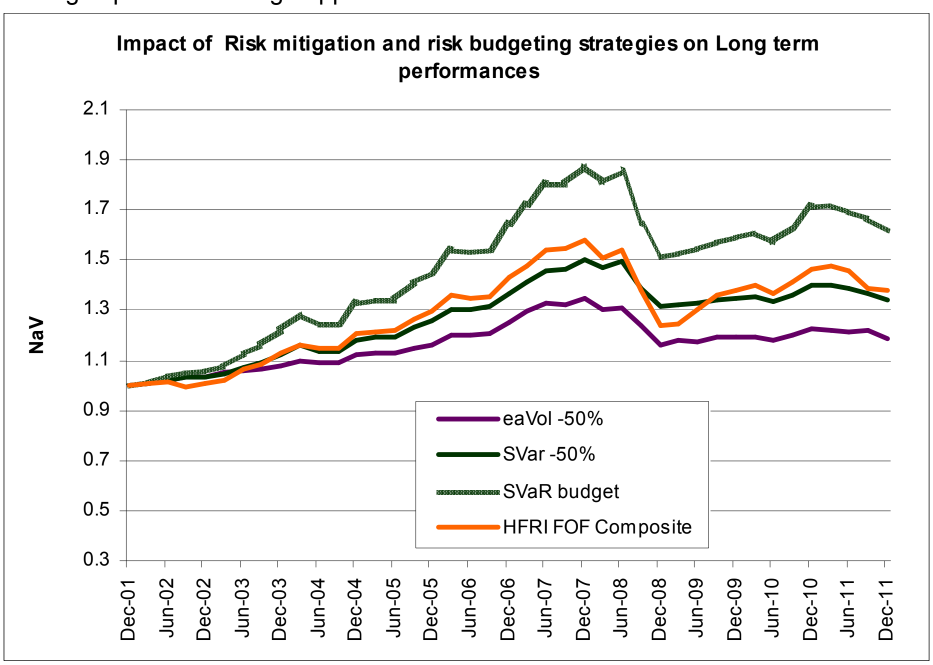

1. Introduction

A frequent statement by Portfolio managers is that returns go along with risk. The Modern Portfolio Theory modeled the relationship between returns and volatility. We examine the relationship between returns and extreme risk. Postfact, the representation of the extreme risk is the Portfolio Peak-to-Valley Ratio(P2V.) Thanks to Hedge Fund databases such as *Hedge Fund Research,* we have access to a large fund data sample to evaluate this relationship on all asset classes and strategies, over the past decade. Specifically, this study aims to answer the following questions::

- Is there a significant relationship between Excess Returns and P2V? If so, what is its nature?
- Can we monitor ex-ante P2V by other means than using a canonical method of adding cash?
- What is the cost of such strategies for both Business-as-usual returns (“Blue Sky” returns) and long-term returns?
- What are the benefits of a risk budgeting policy, and does the choice of Risk indicator matter?

In section 2 we describe the protocol used in the Back-Test procedure to ensure that results are truly out-of–sample and representative on the real conditions under which Hedge fund investors operate. In Section 3, we describe the Hedge funds and Fund allocations sample used in the Study, and estimate, based on this sample, the relationship between Blue Sky returns and P2V. In Section 4, we assess the efficiency of risk mitigation policy to monitor ex-ante P2V, using both ex-ante volatility (*eaVol*) and Stress VaR (*SVaR*) strategies. In Section 5, we quantify the cost of risk mitigation for excess returns and the potential benefits of risk budgeting policy. In Section 6, we examine the impact of Mortality bias on these results, and other adjustments and assumptions that we applied. In Section 7, we analyze the results and discuss implications.

2. The Back Test Protocol

The only way to estimate the real trade-off between excess returns and P2V is to run an out-of-sample back-test on the P2V reduction strategies. The general test protocol consisted of first selecting a set of allocations and a risk indicator *R* which was assumed to be a predictor of future potential P2V of the Portfolio, then iteratively performing the following sequence:

1. Create a portfolio model (M_t in Figure 1 below) at time t, following a systematic algorithm, which takes as input the only information available at the time t, both for the estimation of R and for the underlying asset valuation.
2. Then let portfolio model M_t drift over the liquidity period, keeping constant holdings, and turn it at the end of the period into an actual portfolio P_t, with the same holding structure as M_t, rescaled to match the market value inherited from P_{t-1}.. The liquidity period before implementation is here to account for liquidity constraints (such as lock up clauses)
3. Finally let the actual portfolio P_t drift over the holding period, keeping holdings constant. Then reiterate from step 1.

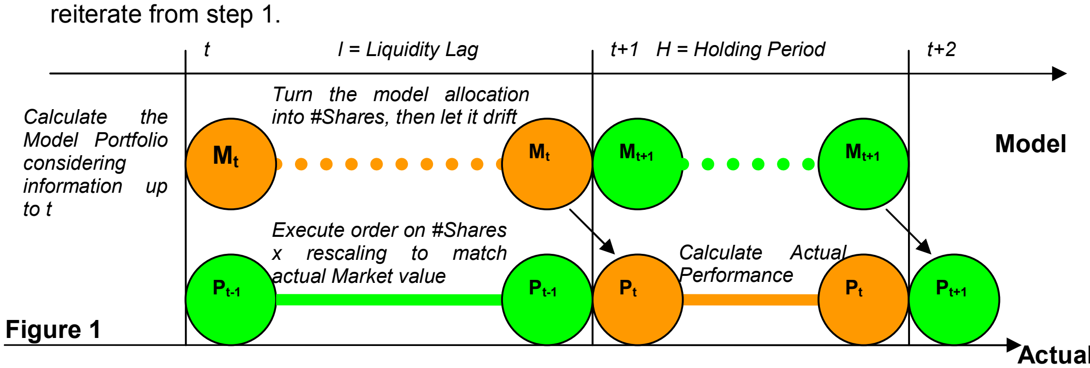

Figure 1

Our back test has been performed with a starting date of December 31 2001 running through December 31 2011, performing the above sequence on quarterly basis – meaning 40 iterations. This means, for instance, that the returns produced by an actual portfolio for the period from July 2008 to the end of September 2008 result from an allocation decision taken in April 2008, based on the information available at the end of March 2008.

These 40 iterations have been applied to 300 systematic rules based on 100 randomly-generated benchmark allocations, with 3 different algorithms tested:

- *Benchmark strategy*: track the allocations with no constraints other than the liquidity ones described above
- *eaVol Strategy*: based on the assumption that future Peak-to-Valley is proportional to the eaVol (as measured at time t). Explicit leverage adjustments are not allowed in this strategy
- *SVaR* based on the assumption that the future Peak-to-Valley is proportional to the Stress VaR as measured at t. Explicit leverage adjustments are not allowed in this strategy

For both eaVol and SVaR strategies, the optimization parameters applied were to stay as close as possible to the benchmark allocation (based on a quadratic distance) while dividing the level of risk by 2 (eaVol or SVaR) vs the benchmark strategy, without using cash or short positions, nor having a concentration higher than 20% on any fund. This resulted in 8000 optimization rounds, run using Riskdata Portfolio Designer optimization algorithm (figure 2), based on a generic Bayesian routine.

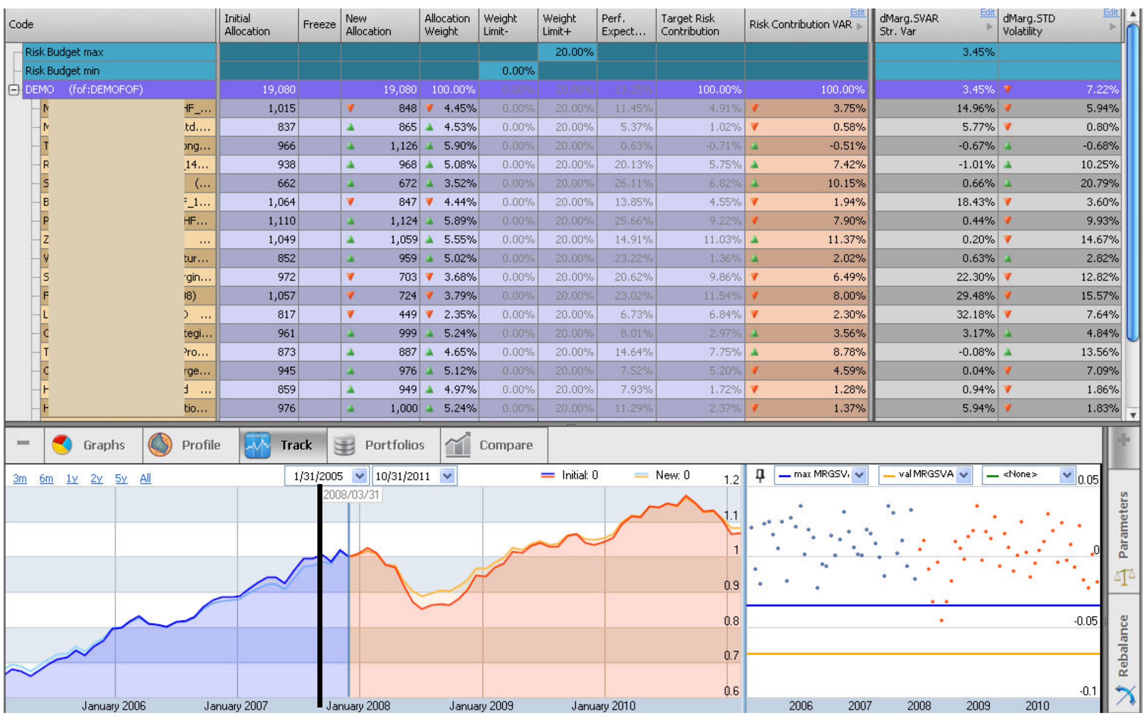

Code	Initial Allocation	Freeze	New Allocation	Allocation Weight	Weight Limit-	Weight Limit+	Perf. Expect…	Target Risk Contribution	Risk Contribution VAR	dMarg.SVAR Str. Var	dMarg.STD Volatility
Risk Budget max						20.00%				3.45%	
Risk Budget min					0.00%						
DEMO (fof:DEMOFOF)	19,080		19,080	100.00%	0.00%	20.00%	13.25%	100.00%	100.00%	3.45%	7.22%
M HF_…	1,015		848	4.45%	0.00%	20.00%	11.45%	4.91%	3.75%	14.96%	5.94%
M td.…	837		865	4.53%	0.00%	20.00%	5.37%	1.02%	0.58%	5.77%	0.80%
T ong…	966		1,126	5.90%	0.00%	20.00%	0.63%	-0.71%	-0.51%	-0.67%	-0.68%
R 14…	938		968	5.08%	0.00%	20.00%	20.13%	5.75%	7.42%	-1.01%	10.25%
S (…	662		672	3.52%	0.00%	20.00%	26.11%	6.82%	10.15%	0.66%	20.79%
B _1…	1,064		847	4.44%	0.00%	20.00%	13.85%	4.55%	1.94%	18.43%	3.60%
P HF…	1,110		1,124	5.89%	0.00%	20.00%	25.66%	9.22%	7.90%	0.44%	9.93%
Z …	1,049		1,059	5.55%	0.00%	20.00%	14.91%	11.03%	11.37%	0.20%	14.67%
V tur…	852		959	5.02%	0.00%	20.00%	23.22%	1.36%	2.02%	0.63%	2.82%
S gin…	972		703	3.68%	0.00%	20.00%	20.62%	9.86%	6.49%	22.30%	12.82%
F 8)	1,057		724	3.79%	0.00%	20.00%	23.02%	11.54%	8.00%	29.48%	15.57%
L) …	817		449	2.35%	0.00%	20.00%	6.73%	6.84%	2.30%	32.18%	7.64%
C tegi…	961		999	5.24%	0.00%	20.00%	8.01%	2.97%	3.56%	3.17%	4.84%
T Pro…	873		887	4.65%	0.00%	20.00%	14.64%	7.75%	8.78%	-0.08%	13.56%
C rge…	945		976	5.12%	0.00%	20.00%	7.52%	5.20%	4.59%	0.04%	7.09%
H d …	859		949	4.97%	0.00%	20.00%	7.93%	1.72%	1.28%	0.94%	1.86%
H tio…	976		1,000	5.24%	0.00%	20.00%	11.29%	2.37%	1.37%	5.94%	1.83%

This protocol was designed to strictly comply with "out-of-sample" procedure while mimicking real-life constraints:

- Views on future performances of assets were expressed through random allocations – therefore with no possible ex-post bias of knowing ex-post how the assets have performed
- Liquidity constraints were incorporated, a critical point for the realism of the back test: it is definitely not the same thing to be bound by an allocation decision taken 4 months before

Lehman bankrupcy – as in our back test – vs. an erroneous assumption that that a portfolio can be reshaped 1 hour before.

- Risk indicators are calculated considering only market and fund time series truncated at the date of simulation.

There was however one bias in the protocol: the selection / survival bias. The only available time series today are the ones on funds who did not stop reporting. This obviously eliminated funds who disappeared from the radar screen, generally because of catastrophic performances. In section 6, we propose a method to correct our results in order to account for this bias.

3. The selected sample of Hedge Funds and the benchmark allocations

The 100 random allocation process was performed on all 359 USD-denominated funds who continuously reported to the HRF database at least from 1 January '99 up to 31 December 2011. None of the random portfolios contained Cash or cash-assimilated products. The allocations and concentrations of for the average portfolio are presented below:

	Average	Std	Min	Max
#Funds	20.4	4.4	9	32
Equity Hedge	47%	11%	18%	73%
Relative Value	12%	7%	0%	28%
Event-Driven	16%	9%	0%	37%
Macro	24%	9%	5%	46%

For each of the 100 random allocations, we first follow the protocol defined in section 1 without applying any constraints. We are then able to produce a realistic track record for each allocation accounting for liquidity constraints.

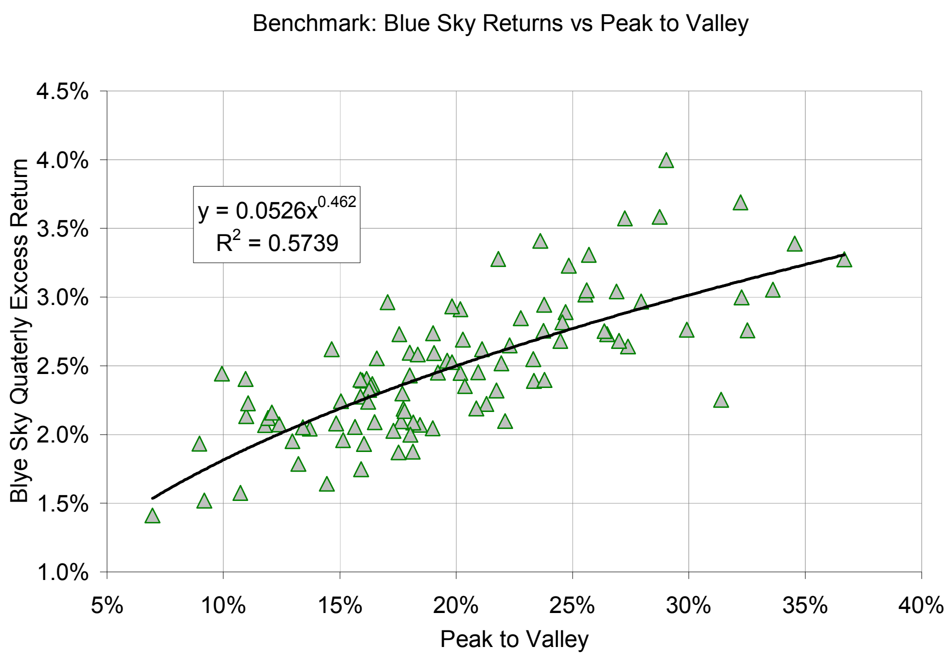

Figure 3

We can observe ex-post that there is a strong relationship between P2V and "blue-sky excess returns" (Figure 3). It means that someone with a crystal ball could monitor their Peak-to-Valley by selecting the appropriate portfolio.

The sensitivity is however much lower that the one of Blue Sky Returns vs Ex Post volatility:

Benchmark: Blue Sky Returns vs Ex Post Volatility

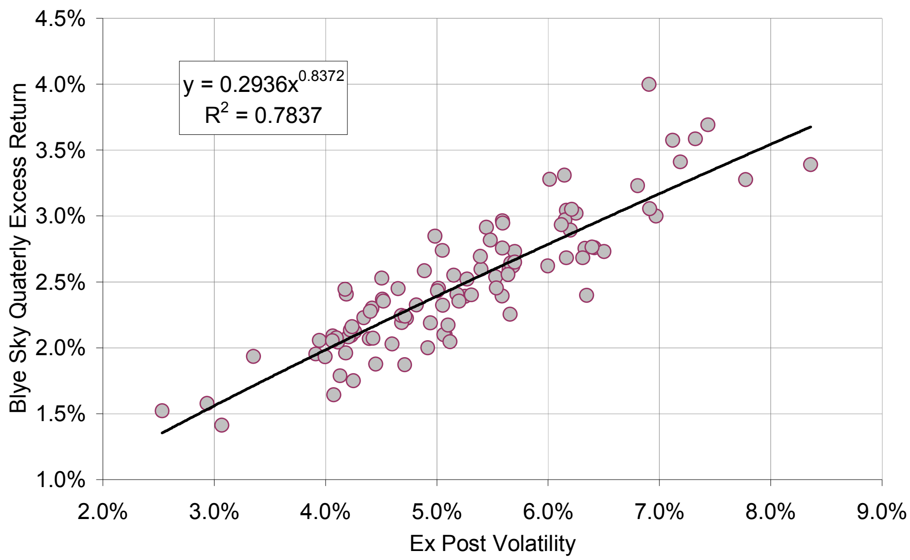

Figure 4

Ex post sensitivity of Blue Sky return vs. P2V is 46% vs. 83% against the Volatility.

In reality, investors do not have a crystal ball. It would be rather counter-intuitive to assume ex-ante that a portfolio with poor performances is in fact a better deal than a high performer, simply because one expects it to follow the statistical law observed ex-post.

4. Ex-Ante Monitoring of Peak-to-Valley

The true test of the relationship is to select ex-ante a reasonable predictor of the Peak-to-Valley, act on it and then check the consequences ex-post. This is exactly what we did with two strategies "eaVol" and "SVaR". We found that both indicators are reasonable tools to efficiently reduce the P2V.

April 4, 2012

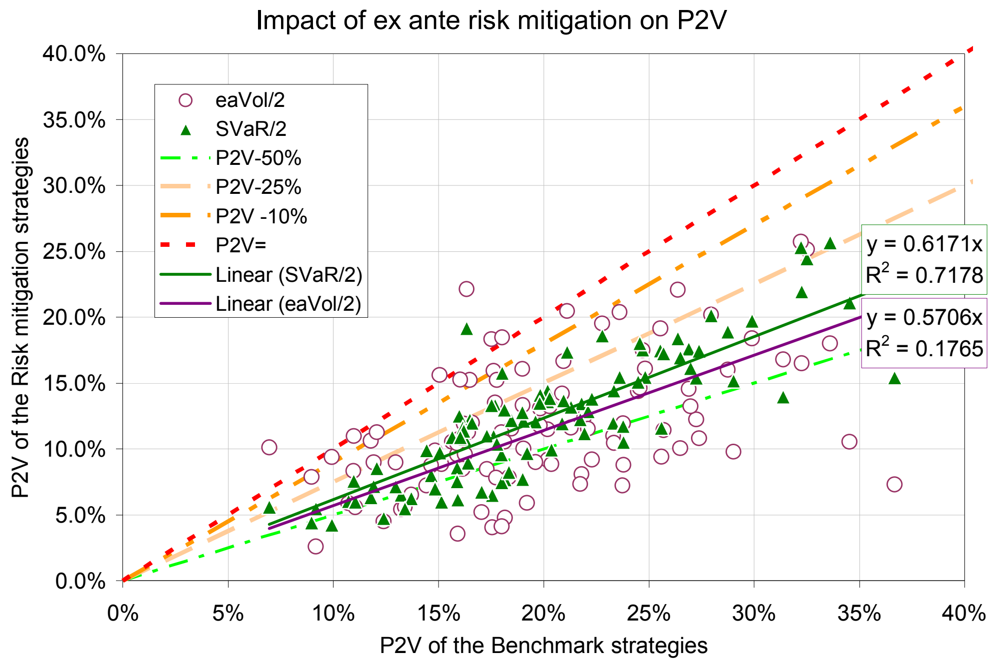

Figure 5

In figure 5, we plot the P2V for each of the 100 benchmark strategies on the horizontal axis, vs. eaVol and SVaR reduction strategies on the vertical axis. If the risk mitigation strategies were perfectly efficient, all points should fall be on the green axis. If the point is actually below the green line, the the risk is “over-mitigated”, ie that the P2V has been reduced by more than twice. If the point is over the line, we “under-mitigate” the risk, and if we are over the black line, we completely fail: the P2V of the risk mitigated strategy is *higher* than its benchmark!

In both cases, average P2V is significantly reduced. Not by a factor of 2 as we would have liked, but reasonably well: 38.3% (+/-1.3%) for the Stress VaR strategies, 43.0% (+/-2.4%) for the volatility.

For the skepticals, who suppose that the P2V reduction is obtained by chance (null hypothesis), we reply that it is impossible, statistically speaking, to be wrong on the average reduction, considering the uncertainty of the coefficients of reduction.

The eaVol strategy is likely to have an average higher impact on P2V than the SVaR (PValue of the null hypothesis being 5%). However, it is far less reliable than the SVaR strategy, considering the following empirical probability of success and failure:

#Cases (on a sample of 100)	eaVol/2	SVar/2
Failure: P2V reduced by less than 10%	13	1
Success: P2V reduced by more than 25%	75	89
Grey Zone: between 10 and 25%	12	10

If the positions of a risk manager were driven by the success of the risk mitigation strategy, with the eaVol approach, he would have a 75% chance of being promoted vs 13% risk of being fired, while with the SVaR strategy, he would have an 89% of chance of being promoted vs. 1% risk of fired.

These results are simply showing that the relationship between volatility and tail risk is loose, unlike that for the stress vaR. The reasons for this difference have been exposed in the SVaR paper in detail[1] . First of all, P2V usually occurs in a market regime different from the one used to estimate the volatility, meaning it is based on a different correlation regime. Secondly, P2V estimation needs to refer to a time window much longer than the actual track record of a fund. This can only be done through factor analysis. Finally, the distribution of tail risk is not really Gaussian.

5. The Cost of Peak-to-Valley Mitigation

The next question is the ex-post impact on the “blue-sky returns”, ie compounding returns out of the P2V period:

We can observe from the figure 5 that:

- Whatever the risk mitigation strategy, there is a cost in terms of “blue-sky” returns – there is no free lunch. However, this cost is always lower than that of the strategy consisting of dividing the leverage by 2, ie meaning dividing both the returns and the peak-to-valley.
- For both cases, we can estimate with a high level of confidence the average cost expressed in % of the Benchmark excess return to be: 16.6% (+/-0.8%) for the Stress Var strategy, and 30.5% (+/-0.8%) for the Volatility one.
- The cost of using volatility is significantly higher than that of using the Stress VaR: with a confidence of 99.95% at least 12% more!

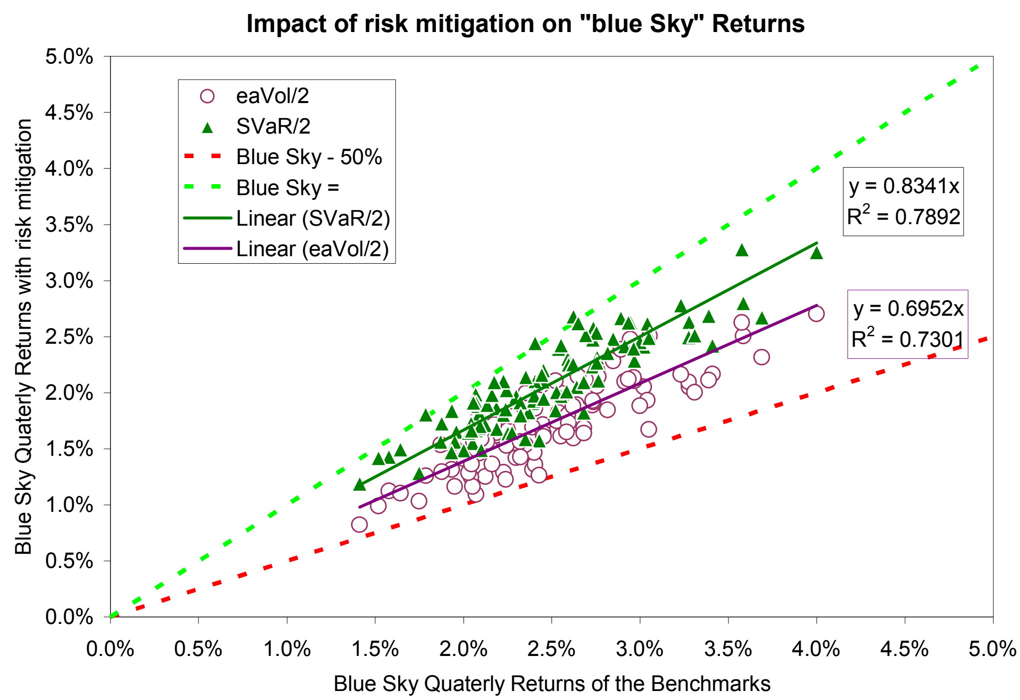

Figure 6

[1] Cyril Coste, Raphael Douady and Illija Zovko « The Stress VaR : A new concept for Extreme risk and Fund allocation », The Journal of Alternative Investment, Winter 2011, Vol.13.

We can then compound the long term the cost in term of Blue Sky excess return and the Peak to Valley mitigation, to deduce their net effect vs. the benchmark:

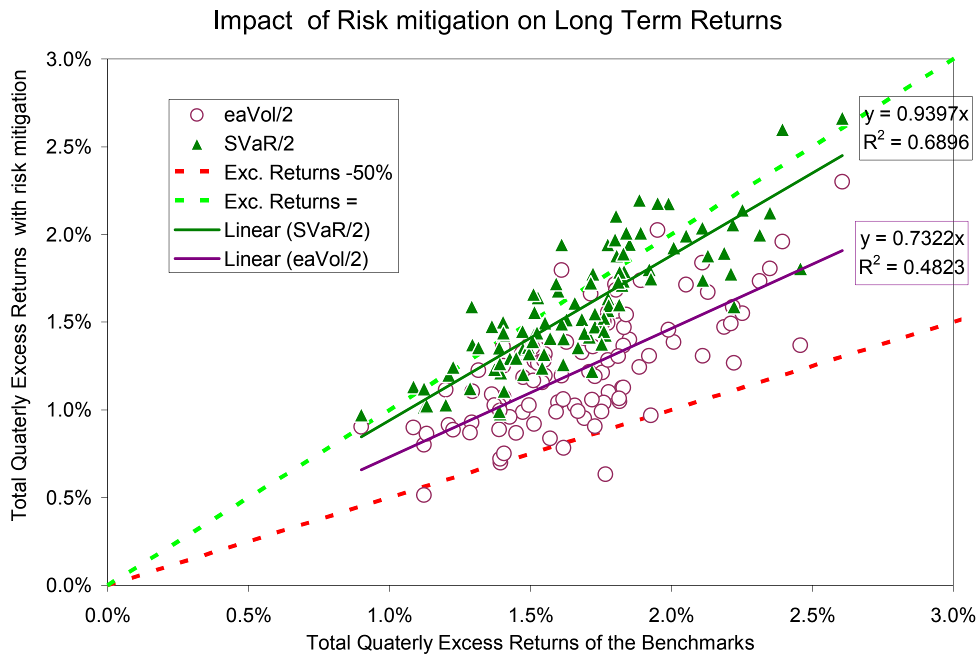

Figure 7

As we can see in figure 7, there are no miracles: on the average in the long run, risk mitigation has an impact on the costs of excess returns: 27% +/- 1.4% for the eaVol Strategy, 6% +/- 1.1% for the SVaR strategy.. The cost difference between the two approaches is significant: at least 14.6% with a 99.95% confidence!

Furthermore, to assess the benefit of risk mitigation, we "re- leverage" each of the risk mitigation strategies to the level of risk of the benchmark, assuming that the Benchmark risk as what is acceptable to the investor. In other words, we assess here the benefit of assigning a Risk budget vs. no Risk budget. We make a conservative assumption that the leverage provider has a crystal ball, that allows him to to predict the P2V reduction coefficient (Figure 5) – and hence to provide less than a factor-of-2 leverage, and less leverage for the SVaR strategy (1.62) than for the eaVol one (1.75).

We need of course to take into consideration that the higher leverage will involve an additional cost, which will reduce the benefit of the re-leveraging for the equity owner. In an efficient market, spread charged by the lender should match the cost of default, itself equal to (1-q)/q, q being the confidence that the fund will not default This confidence q is in itself driven by 2 main factors: a *market factor* a function of the leverage level– i.e. the risk to see the portfolio loss exceed the value of the equity – and the operational risk factor – a factor independent from the leverage level. This operational factor in our case discretionary - related to the credibility of the borrower: it can be viewed as a discretionary tax on the excess return. The market factor impact can be estimated in the following way. First we assume the lender has a crystal ball, and hence that the worst case scenario on 40 quarters (10 years) is the leveraged P2V, which becomes a proxy for the 97.5% level of risk. .The lender can then extrapolate the probability of default, considering that the tail distribution is the one of a leveraged SP500. It is interesting to use the historical distribution of SP500 quarterly returns since 1920 – incorporating fat tails. Using it we can observe that the cost of default is 1 bp if the leveraged P2V is 42%, 10 bp if it is 50%.

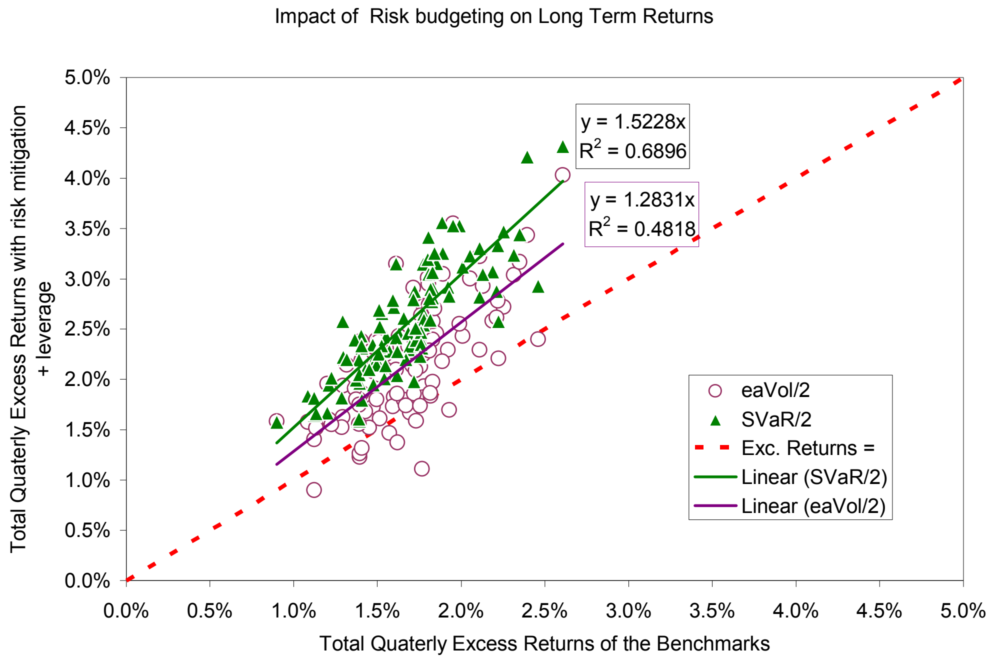

Figure 8

The benefit of managing the portfolio via risk constraints is quite striking, and is presented in figure 8:

- In both cases, the P-Value of the model – the fact that on average the excess returns are increased – is zero, meaning that it is impossible that these observations result from chance.
- For the eaVol strategy, the average long term performance is increased by 28%. The excess returns are improved in 88% of the cases if there is no discretionary spread.
- For the SVaR Strategy, the average long-term performance is increased by 52%, and the excess returns are improved in 100% of the cases with no spread.
- On the average, the SVaR strategy over-performs the eaVol one by 24%, with a 99.95% chance of being above 13%.

Based on these results, we arrive at the following, statistically significant conclusion: in the long run, a systematic policy of risk mitigation, targeting a controlled P2V, is profitable, if based on SVaR, unprofitable if based on the eaVol.

-

6. The selection bias impact

The selection bias can be estimated comparing our sample with an index which is known to have a low mortality bias. We choose HFRI Fund of Hedge Funds Composite index, because even if its underlying hedge funds stop reporting to HFR database, their performances are still captured through the fund of hedge funds performances. However, the gap between this index and our sample comes not only from the selection bias, but also from the management and performance fees charged by the fund of hedge funds. We take a hypothesis of a 2% management fees per year, and a 15% performance fees to apply on the yearly excess returns of the HFRI index. :

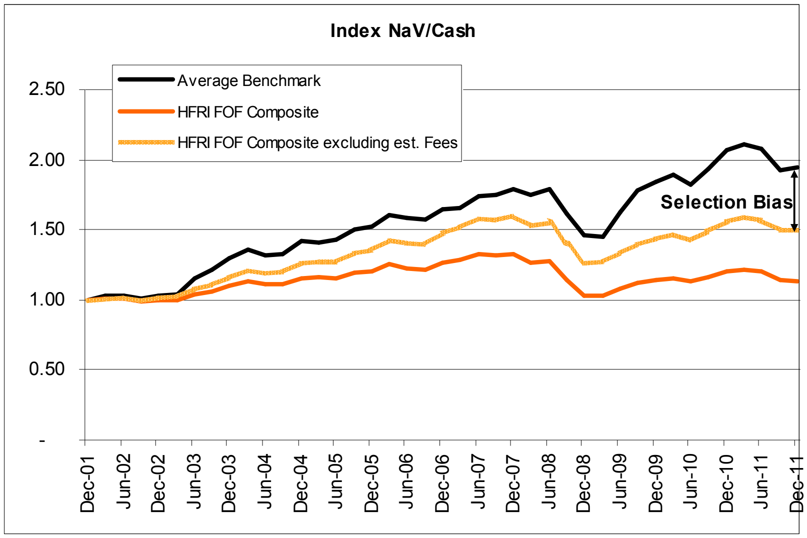

Figure 9

This leads us to an estimated "Mortality Tax" of 64 basis points per quarter over the period, with peak after crisis – typically when funds which have been the most impacted by redemptions stop reporting:

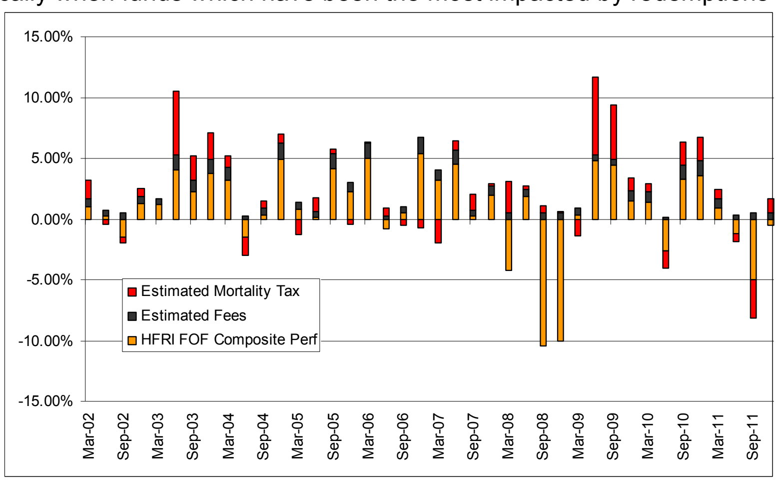

Figure 10

We could of course make the assumption that this bias is neutral vs. the relative performances reported in figure 8. This would be a rather optimistic assumption, it would imply that the Mortality Tax is reduced proportionally to P2V.We therefore make the following hypothesis, which we consider both realistic and highly conservative:

- The Mortality tax on the benchmarks is proportional to their long-term performance: the more aggressive the portfolio is, the higher the mortality bias would be. This leads us to a tax equal to 38.5% of the excess returns.
- For the risk diversification strategies, the Mortality tax in basis points is the same as that of the benchmark. For instance, if on the benchmark, the tax costs 60 b.p. per quarter, it will be the same for the corresponding risk diversification strategy. This means that we assume that the Mortality tax will not be reduced by the risk diversification, as if the hidden portion of the

portfolio based on our fund selection kept the same weights before and after risk diversification.

Impact of Risk diversification accounting for Mortality tax

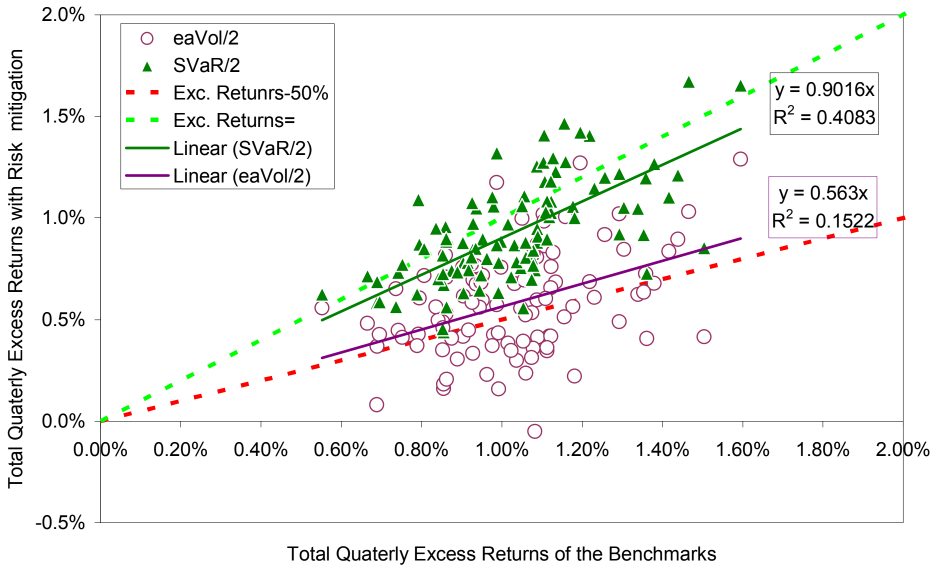

Figure 11

Impact of Risk budgeting on Long Term Returns, accounting for Mortality Tax

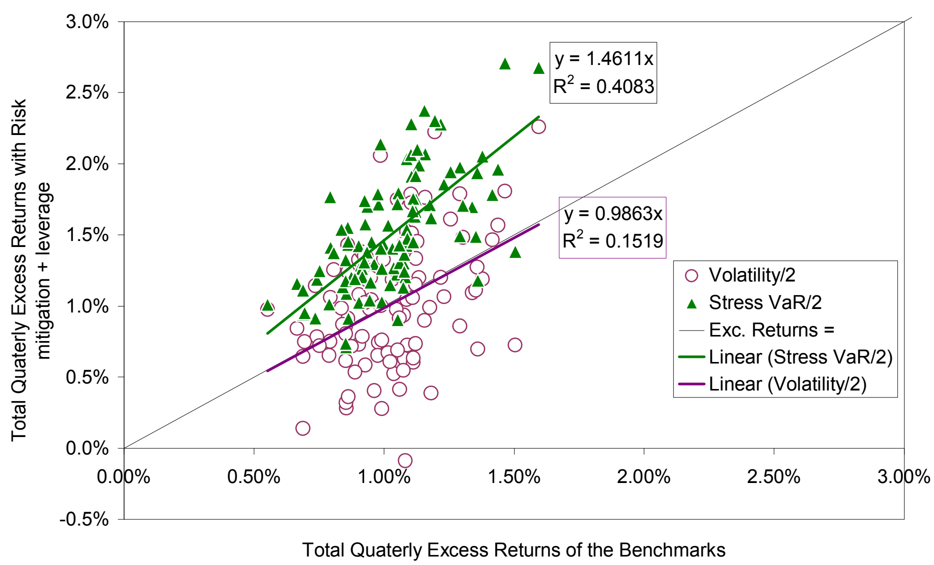

Figure 12

As we can see in figure 11, adding the Mortality tax mechanically amplifies the gap between the inexpensive and costly risk diversification strategies. This mechanical effect can be seen in the chart above: while Stress-VaR produces risk diversification at almost no cost, Volatility strategy rates quite badly: no significant improvement on the average compared to a simple deleveraging approach: in 41% of the cases it underperforms.

This mechanically translates into benefits for risk budgeting, assessed by re-leveraging the risk diversified portfolios in order to match, , the benchmark:

- A high and tangible benefit for SVaR strategy, with an average increase of 46% in the excess returns and only 5% of cases where it underperforms the benchmark
- it becomes unattractive for the eaVol strategy , with no significant gain on the excess returns, and with 50% of the cases where it underperforms the benchmark.

7. Conclusion

The results produce a straightforward conclusion about the trade-off between Peak-to -Valley and excess returns:

- Volatility Diversification strategy is both unreliable and costly, on the average, as much as a canonical deleveraging strategy`
- Stress VaR is reliable and cost almost nothing in terms of returns.

We should not infer from this that the returns and reliability costs are interconnected. In a perfectly efficient market, all risk mitigation strategies should be neutral in average – meaning impact equally excess returns, and should only differ by they dispersion around the mean – high for bad predictors, low for good ones.

What we observe here can be analyzed as a consequence of the following mechanism:

- Ex-ante volatility acts primarily on ex-post volatility. And returns are highly sensitive to ex-post volatility, thanks to 30 years of markets seeking to maximize the Sharpe Ratio ranging from 0.8 to 0.9, as we can see in figure 13. When using volatility to reduce P2V, we reduce returns too through this direct relationship.

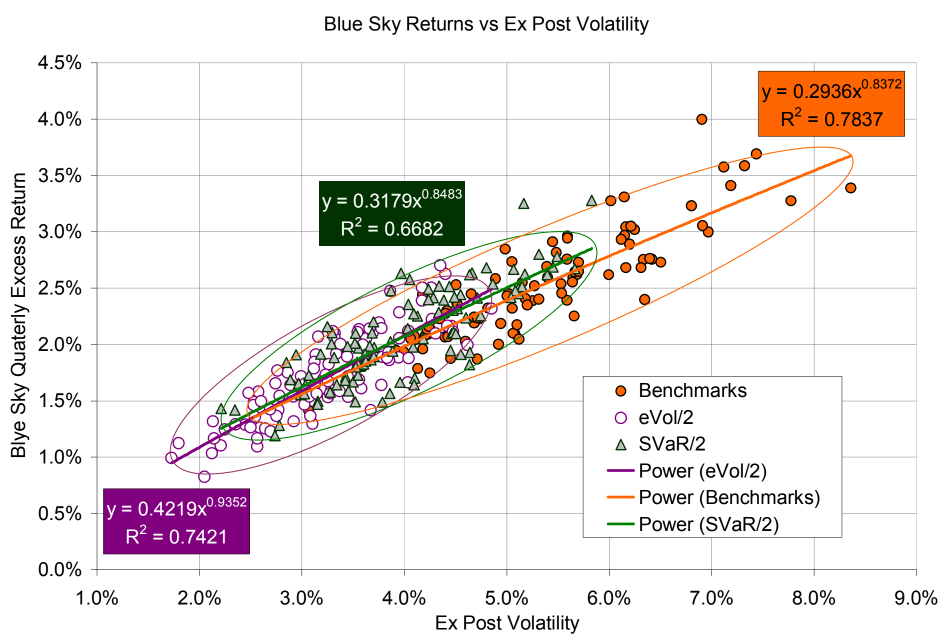

Figure 13

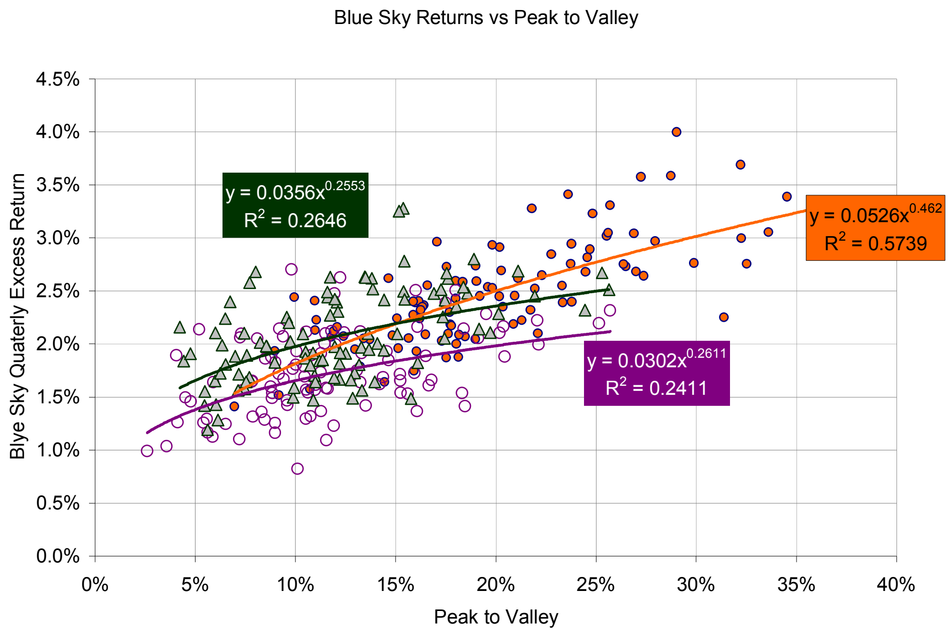

Figure 14

.SVaR strategy allows to reduce P2V without significantly impacting neither the-ex post volatility, nor the blue Sky returns: The direct sensitivity of “Blue Sky” Returns to the P2V is much lower vs. ex post volatility, ranging from 0.25 to 0.4, as we can see from the figure 14.

In reality what matters is the relationship between returns and tail risks (see Leveraged Portfolio Theory) and we simply observe here an arbitrage opportunity. At the end of the day, Stress-VaR strategy wins because it reduces P2V without proportionally impacting the volatilities, i.e. the actual source of returns in the market. As we can observe in the figure 15, the volatility of the Stress VaR portflios is systematically higher than that based on ex-ante volatility:

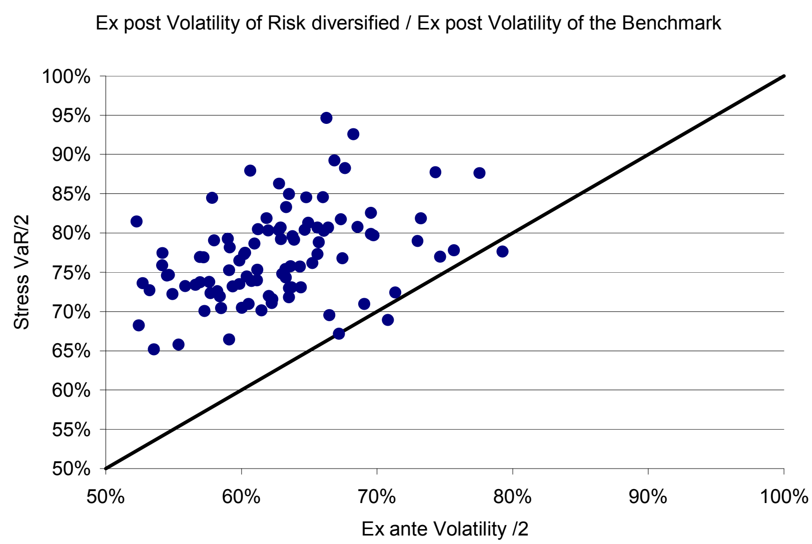

Figure 15

This explains why, over the entire period, in average, applying to the selection both mortality tax, management and performance fees to make the average selection comparable with HFRI FOF Composite index, we observe that an investor using the Stress-VaR strategy can significantly reduce his Peak-to-Valley without giving up an important fraction of performance . Generating a risk-budgeting policy based on this analysis, the investor would gain over the period, net of fees, an average 1.68% per year, which, compounded over the period, would represent a profit equal to 24% of his initial capital.