Currency Allocation As Dual Benchmark Optimization

Dr. Bernd Scherer, Dr. Andreas Gintschel

Currency allocation for central banks is formulated as a dual benchmark optimization problem in which central banks may attach different weights to a nominal wealth preservation benchmark (local cash) and a liquidity benchmark (short term debt and imports). Currency returns are modeled as drawings from two regimes, allowing for shifts in correlation and volatility as well as for non-normality. To construct optimal portfolios we use a novel approach to dual benchmark optimization that allows for multiple benchmarks, multiple risk regimes, and non-normality

---

1. Currency Allocation and Central Banks

Currency reserve management centers around a set of related questions. Central banks need to determine the optimal level of reserves, the optimal split between highly liquid assets (which can easily be converted into cash in times of market stress) and less liquid assets (also called alternatives), as well as the optimal choice between currencies (currency allocation). This paper will concentrate on the later problem.

Allocating among currencies is complicated by the fact that they represent assets and the numeraire at the same time. For example: the volatility of short term US bonds (assets) depends on whether we view them from a European or a Japanese perspective (numeraire). However, so far the literature on currency benchmarking problems has narrowly focused on either one of the following optimization problems.[1]

> **Wealth preservation approach.** Wealth put into foreign assets is volatile in terms of the numeraire (home) currency. Hence positions are often diversified across many currencies to enjoy a reduction of risk.
>
> **Liquidity preservation approach.** To ensure international solvency (at least in the short run) central banks need enough reserves to cover imports, or to service foreign debt payments. The closer the currency composition to this allocation the less risky it is even though it might be very volatile in terms of the home currnency.

Rather than simply following one of these approaches in isolation the currency allocation is presented as a multi-objective optimization problem in which central banks can not only attach different weights to each subproblems, but are also allowed to consider alternative risk regimes.

2. Currency Returns: Rival Risk Regimes and Non-Normality

Portfolio construction requires inputs. There is a vast literature on currency returns in particular, and on input estimation for portfolio optimization in general. We do not provide a

[1] Currency return optimization relative to various benchmarks has been addressed in Boorman/Ingves (2001), Ramaswamy (1999) and Scobie/Cagliesi (2000). Dual Benchmark Optimization is covered in Scherer (2002).

comprehensive review or argue for a preferred methodology. Instead we focus on two aspects of currency returns that ared used in the following sections. The currencies we focus on are US Dollar, UK Pound, Japanese Yen, Australian Dollar, Swiss Franks (all against the Euro). We use weekly return data (currency plus local cash return) from J.P. Morgan covering the period from January 1986 to December 2002.

It is well known that correlations break down in times of market meltdowns, i.e. when portfolio managers need them most. We will not attempt to forecast the change in input parameters. However we look for a tool to evaluate the diversifying properties of currencies in rivalling risk regimes. As supervisory boards become more and more concerned about short term performance, investors often do not have the luxury to bet on average correlation or average volatility. To come up with correlation and volatility estimates for normal and hectic times have to define first what exactly do we mean with unusual times.[2] We define unusual times according to their statistical distance from the mean vector as given in

$$(1.1) \qquad \left(\mathbf{R}_m - \bar{\boldsymbol{\mu}}\right)^T \bar{\boldsymbol{\Omega}}^{-1} \left(\mathbf{R}_m - \bar{\boldsymbol{\mu}}\right) = \mathbf{d}_m^T \bar{\boldsymbol{\Omega}}^{-1} \mathbf{d}_m$$

where $\mathbf{d}_m$ reflects the distance vector at time m, $\mathbf{R}_m$ is a vector of currency return observations for N currencies at time m, $\bar{\boldsymbol{\mu}}$ denotes a vector of average currency returns and $\hat{\boldsymbol{\Omega}}$ is the unconditional covariance matrix (over all $m = 1, \cdots, M$ observations). For each cross section of stock returns we calculate (1.1) and compare it to the critical value of a $\chi^2_{0.95}(N)$. If we define an unusual observation as the outer 5% of a distribution (alternatively one might call it outlier) for 5 return series, the cut off distance is 11.07. In (1.1) the return distance is weighted by the inverse of the covariance matrix. Thus, we take into account currency volatilities (The deviation from mean might be significant for low volatility series, but not necessarily for high volatility series) and correlations (return differences of opposite sign for two highly correlated series might be more unusual than for series with negative correlation). Hence outliers are, in theory, not necessarily associated with down markets (although they are often in practice). The correlation matrices for normal and hectic times are given below.

[2] See Chow et al (1999).

$$\boldsymbol{\rho}_{normal} = \begin{array}{c} \begin{array}{cccccc} \text{USD} & \text{UKP} & \text{YEN} & \text{ASD} & \text{FRK} & \text{EUR} \end{array} \\ \begin{pmatrix} 1,0 & 0,5 & 0,5 & 0,7 & -0,1 & -0,1 \\ 0,5 & 1,0 & 0,3 & 0,4 & 0,1 & 0,0 \\ 0,5 & 0,3 & 1,0 & 0,4 & 0,1 & 0,0 \\ 0,7 & 0,4 & 0,4 & 1,0 & -0,1 & -0,1 \\ -0,1 & 0,1 & 0,1 & -0,1 & 1,0 & 0,0 \\ -0,1 & 0,0 & 0,0 & -0,1 & 0,0 & 1,0 \end{pmatrix} \end{array} \boldsymbol{\rho}_{hectic} = \begin{array}{c} \begin{array}{cccccc} \text{USD} & \text{UKP} & \text{YEN} & \text{ASD} & \text{FRK} & \text{EUR} \end{array} \\ \begin{pmatrix} 1,0 & 0,4 & 0,5 & 0,8 & 0,0 & 0,1 \\ 0,4 & 1,0 & 0,1 & 0,3 & 0,0 & 0,0 \\ 0,5 & 0,1 & 1,0 & 0,4 & 0,2 & 0,0 \\ 0,8 & 0,3 & 0,4 & 1,0 & 0,0 & 0,1 \\ 0,0 & 0,0 & 0,2 & 0,0 & 1,0 & 0,1 \\ 0,1 & 0,0 & 0,0 & 0,1 & 0,1 & 1,0 \end{pmatrix} \end{array}$$

Table 1. Correlation in normal and hetic times

Interestingly to see that correlations between currency returns remain virtually unchanged in crisis times. The only exception seems to be the Australian Dollar, which exhibits a slight increase in correlation in hectic periods. With diversification properties unchanged, international currency allocation keeps its attractivness in times of market crisis. Apart from correlations we can also calculate the corresponding volatilities.

$$\boldsymbol{\sigma}_{normal} = \begin{pmatrix} 9,3 & & & & & \\ & 6,6 & & & 0 & \\ & & 9,8 & & & \\ & & & 11,9 & & \\ & 0 & & & 3,6 & \\ & & & & & 0,4 \end{pmatrix} \boldsymbol{\sigma}_{hectic} = \begin{pmatrix} 20,6 & & & & & \\ & 14,6 & & & 0 & \\ & & 20,1 & & & \\ & & & 26,4 & & \\ & 0 & & & 7,8 & \\ & & & & & 0,4 \end{pmatrix}$$

Table 2. Volatility in normal and hectic times

The dramatic rise in risk (volatility more than doubles) as we move from normal to hectic times is apparent. We can now glue together the above results to arrive at the covariance matrices

$$\boldsymbol{\Omega}_s = \boldsymbol{\sigma}_s \boldsymbol{\rho}_s \boldsymbol{\sigma}_s$$

with $s = normal, hectic$ for later use. Effectivly, we split the coavriance matrix into a high and a normal volatility regime.

So far we have assumed normality in currency returns. How well does this approximization hold? Studying Figure 1and Figure 2 shows that returns are fat tailed (deviation from the straight line in the QQ-plot at both ends) and sometimes skewed (deviation from the straight line in the QQ-plot at both ends) as in the case of the Japanese Yen. Section 5 will investigate the impact of non-normality on optimal portfolio choice.

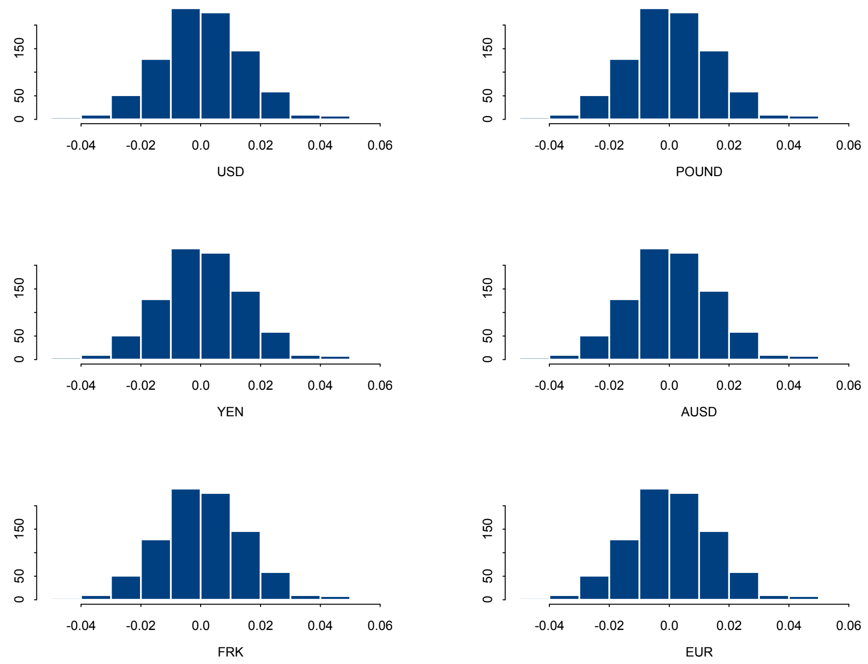

Figure 1. Histogram of currency returns

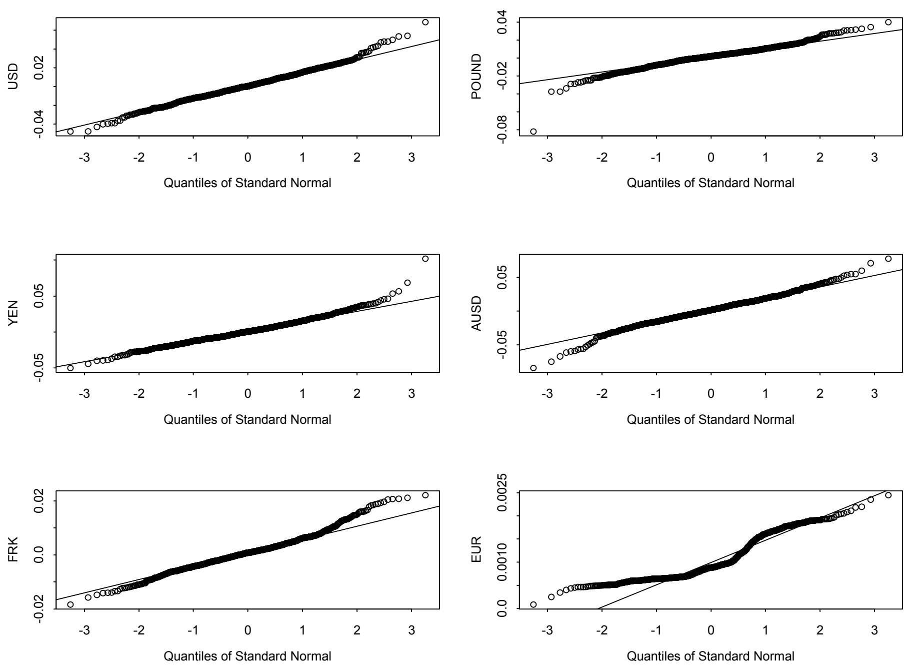

Figure 2. QQ-plot of currency returns

Finally, we need estimates of expected currency returns. We use the James-Stein estimator to reduce estimation error in means and, hence reduce the tendency to arrive at corner solutions (solutions where asset weights are limited by their constraints). The estimator rests on the notion that we can improve the efficiency of the historical mean estimate on an individual return series by pooling the information in all series.

Effectively the JS estimator shrinks the individual historical means towards the grand mean of all time series. In the case of complete shrinkage we arrive at means equal to the grand mean and will ultimately end up with the minimum variance portfolio. The JS estimator

$$\boldsymbol{\mu} = \phi \bar{\mu} \mathbf{I} + (1 - \phi) \bar{\boldsymbol{\mu}} \tag{1.2}$$

where $\bar{\mu}$ denotes the grand mean of all series $\bar{\mu} = \sum_{n=1}^{N} \mu_n$, $\mathbf{I}$ represents a vector of ones and $\bar{\boldsymbol{\mu}}$ reflects the vector of historical means. The shrinkage factor can be derived from

$$\phi = \min\left[1, \frac{(N-2)}{M\left(\bar{\boldsymbol{\mu}} - \bar{\mu}\mathbf{I}\right)^T \boldsymbol{\Omega}^{-1} \left(\bar{\boldsymbol{\mu}} - \bar{\mu}\mathbf{I}\right)}\right] \tag{1.3}$$

If $\left(\bar{\boldsymbol{\mu}} - \bar{\mu}\mathbf{I}\right)^T \boldsymbol{\Omega}^{-1} \left(\bar{\boldsymbol{\mu}} - \bar{\mu}\mathbf{I}\right)$ is small, either because the distance between the historical means and the grand mean is small, or because the precision of our estimates is small, the James-Stein estimator will shift our estimates towards the grand mean. The same is true if the number of assets is large, or the number of observations is low.

3. Multiple Benchmarks and Rival Risk Regimes

When predicting future currency returns, decision makers often fail to agree on the probability ordering of alternative currency risk regimes. Additionaly, they might differ in the importance they attach to a wealth preservation benchmark (local cash) versus a liquidity preservation benchmark (short term debt and imports). Suppose central bankers face $s = 1 \cdots S$ risk regimes, reflected in the associated currency covariance matrices $\boldsymbol{\Omega}_s$ and

$b=1\cdots B$ currency benchmark portfolios. The central bank is allowed to invest $n=1\cdots N$ assets summarized in the vector of currency holdings, $\mathbf{w}$. The optimization problem becomes

(1.4) $$\max_{\mathbf{w},\mathbf{w}\geq 0}\left(\mathbf{w}^T\boldsymbol{\mu}-\lambda\max_{s,b}\left[\left(\mathbf{w}-\mathbf{w}_b\right)^T\boldsymbol{\Omega}_s\left(\mathbf{w}-\mathbf{w}_b\right)\right]\right)$$

where λ reflects the central bank's risk aversion. This approach allows us to find protection against the risk of adopting an investment strategy based on the wrong benchmark or/and the wrong risk regime. We can reformulate (1.4) in a way digestable to solvers for constrained quadratic programs.[3]

$$\begin{aligned}
&\max_{\mathbf{w},\sigma^2_{\max}}\left(\mu-\lambda\sigma^2_{\max}\right)\\
&\left(\mathbf{w}-\mathbf{w}_b\right)^T\boldsymbol{\Omega}_s\left(\mathbf{w}-\mathbf{w}_b\right)\leq\sigma^2_{\max}\\
&\mathbf{w}^T\mathbf{I}=1\\
&\mathbf{w}^T\boldsymbol{\mu}=\mu\\
&\mathbf{w}\geq 0
\end{aligned}$$

Note that this defaults to standard markowitz optimization for $s=1,b=1$. In the following, we assume two benchmarks (liquidity and wealth preservation) as

$$\begin{aligned}
w_{wealth}&=\begin{pmatrix}0 & 0 & 0 & 0 & 0 & 1\end{pmatrix}^T\\
w_{liquidity}&=\begin{pmatrix}0{,}5 & 0 & 0{,}5 & 0 & 0 & 1\end{pmatrix}^T
\end{aligned}$$

and two risk regimes. The wealth benchmark consists of 100% Euro cash, while the liquidity preservation benchmark is assumed to contain 50% US Dollar and 50% Yen. These figures hypothetical, but provide a numerical example. We can now investigate the effect of two benchmarks on portfolio choice. Suppose we start with a portfolio of 100% US Cash and gradually move allocations into Euro Cash. Which effect would this have on benchmark relative risk? Suppose a central bank with the above benchmark definitions and risk regimes chooses a 20% allocation in the US Dollar. The relative riskiness varies between 1.87% (wealth benchmark in normal times) and 13.99% (liquidity benchmark in hectic times) in Figure 3. Board members might argue about the riskiness of any given strategy depending on which weight they place on the respective benchmarks and risk regimes.

[3] All optimizations have been performed using NUOPT for S-Plus. See Scherer/Martin (2003) for an extensive treatment of portfolio optimization problems in S-Plus.

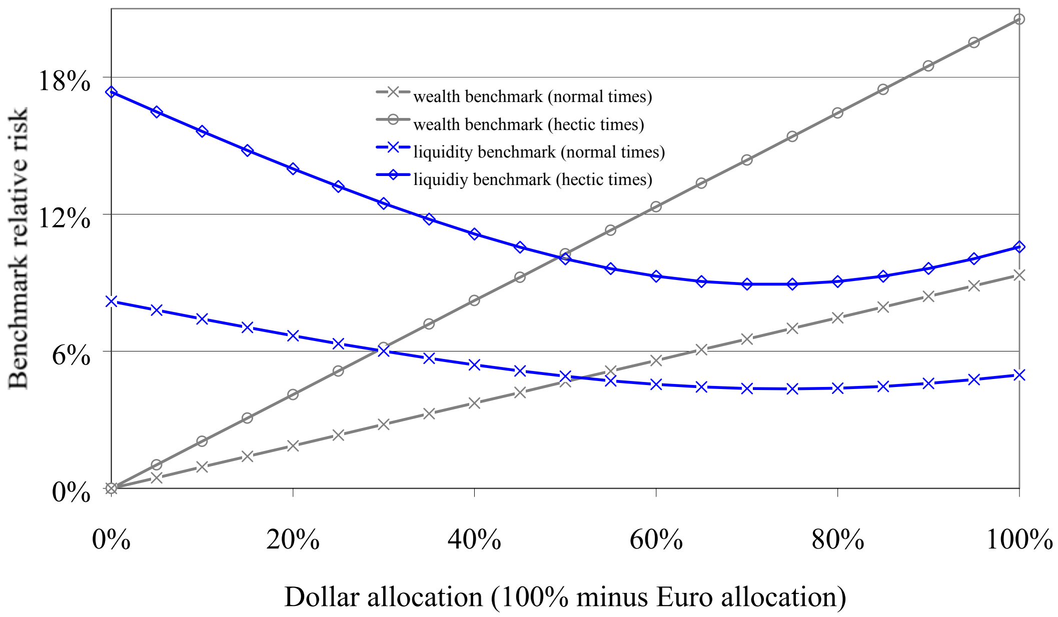

Figure 3. Realtive risks under alternative benchmarks and risk regimes

Hence the regret of having decided to manage against the wrong benchmark in the wrong risk regime is potentially large. However allocating 50% to both US Dollar and Euro will significantly narrow the range of potential outcomes anywhere between 4.67% and 10.27%. In contrast to the 20%/80% allocation he worst case riskiness falls to 10.275 (from 13.99%).

So far we have only addressed risk issues for a simple two currency portfolio. We will now directly solve (1.4). The optimization results can be seen in Table 3 through Table 5.

Currency	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}
USD	0,00%	25,61%	19,07%	12,09%	5,13%	0,66%	0,02%	0,00%	0,00%	0,00%
Pound	0,00%	0,00%	6,95%	15,16%	23,35%	32,78%	44,05%	55,63%	67,22%	100,00%
Yen	0,00%	23,91%	24,65%	24,24%	23,83%	21,90%	17,67%	13,05%	8,43%	0,00%
AUSD	0,00%	0,00%	4,33%	9,43%	14,51%	18,35%	20,29%	21,91%	23,52%	0,00%
FRK	100,00%	10,65%	0,07%	0,09%	0,02%	0,05%	0,00%	0,00%	0,00%	0,00%
EUR	0,00%	39,84%	44,93%	38,99%	33,15%	26,26%	17,98%	9,41%	0,83%	0,00%
Risk	18,33%	8,71%	8,76%	9,07%	9,58%	10,28%	11,21%	12,36%	13,67%	19,22%
Return	-0,42%	-0,13%	0,17%	0,46%	0,76%	1,05%	1,35%	1,64%	1,94%	2,23%

Table 3. Optimal portfolios with rivalling risk regimes

Currency	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}
USD	0,00%	25,90%	19,87%	14,02%	8,08%	2,47%	0,12%	0,00%	0,01%	0,00%
Pound	0,00%	0,04%	8,29%	18,38%	28,35%	38,43%	49,83%	62,14%	74,49%	100,00%
Yen	0,00%	23,97%	24,29%	23,50%	22,65%	21,64%	18,26%	13,28%	8,20%	0,00%
AUSD	0,00%	0,06%	2,84%	5,78%	8,94%	11,89%	13,54%	14,28%	14,97%	0,00%
FRK	100,00%	10,95%	0,10%	0,01%	0,26%	0,00%	0,00%	0,00%	0,00%	0,00%
EUR	0,00%	39,08%	44,62%	38,31%	31,72%	25,57%	18,25%	10,31%	2,34%	0,00%
Risk	8,88%	4,11%	4,14%	4,28%	4,53%	4,85%	5,25%	5,75%	6,32%	7,85%
Return	-0,42%	-0,13%	0,17%	0,46%	0,76%	1,05%	1,35%	1,64%	1,94%	2,23%

Table 4. Optimal portfolios with dual benchmarks and single risk regime (normal times)

Currency	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}
USD	0,00%	25,61%	19,07%	12,10%	5,14%	0,03%	0,01%	0,00%	0,00%	0,00%
Pound	0,00%	0,00%	6,95%	15,15%	23,36%	32,46%	44,04%	55,63%	67,22%	100,00%
Yen	0,00%	23,91%	24,65%	24,25%	23,82%	22,29%	17,68%	13,05%	8,42%	0,00%
AUSD	0,00%	0,00%	4,34%	9,42%	14,50%	18,66%	20,29%	21,91%	23,53%	0,00%
FRK	100,00%	10,65%	0,10%	0,04%	0,03%	0,00%	0,00%	0,00%	0,00%	0,00%
EUR	0,00%	39,84%	44,90%	39,05%	33,14%	26,55%	17,98%	9,41%	0,83%	0,00%
Risk	18,33%	8,71%	8,76%	9,07%	9,58%	10,27%	11,21%	12,36%	13,67%	19,21%
Return	-0,42%	-0,13%	0,17%	0,46%	0,76%	1,05%	1,35%	1,64%	1,94%	2,23%

Table 5. Optimal portfolios with dul benchmarks and single risk regime (hectic times)

We see that our optimization results are (not surprisingly) dominated by the worst case risk scenario. Hence the solutions in which we used both regimes are virtually identical to the solutions in which we only used the hectic regimes. The associated risks are also much higher than in normal times. If we compare these results with allocations in which we only used the covariance matrix for normal times, we see that the differences are less pronounced than we might have expected. However as correlations hardly change and all volatilities increased by similar magnitude this makes perfect sense. The only exception is the Australian Dollar that attracts higher allocations because of the improved risk return tradeoff. Note that the minimum risk portfolio remains unchanged under all optimizations above, further supporting our observation that the relative riskiness is stable.

4. Multiple Benchmarks and Pareto Optimality

So far we did not allow for varying risk preferences which would allow us attach different weights to the various subproblems. However this will be introduced now.[4]

[4] See Shectman (2000) on pareto optimality and dual benchmark optimization.

(1.5) $$\max_{\mathbf{w},\mathbf{w}\geq 0}\left(\min_{s,b}\left[\mathbf{w}^T\boldsymbol{\mu}-\lambda_{s,b}\left(\mathbf{w}-\mathbf{w}_b\right)^T\boldsymbol{\Omega}_s\left(\mathbf{w}-\mathbf{w}_b\right)\right]\right)$$

Equation (1.5) poses the problem of maximizinig the minimum risk adjusted performance (utility) across both alternative benchmarks, risk regimes, and the associated risk aversion coefficients. This is equivalent to maximizing the minimum utility (assuming mean variance preferences). The resulting solution is pareto optimal in the sense that we can not increase utility any further without pushing utility from another subproblem below the minimum utility. Note that (1.5) differs from the conventional treatment of multiple benchmark problems[5]

(1.6) $$\max_{\mathbf{w},\mathbf{w}\geq 0}\left(\mathbf{w}^T\mathbf{R}-\lambda_1\left(\mathbf{w}-\mathbf{w}_1\right)^T\boldsymbol{\Omega}\left(\mathbf{w}-\mathbf{w}_1\right)-\lambda_2\left(\mathbf{w}-\mathbf{w}_2\right)^T\boldsymbol{\Omega}\left(\mathbf{w}-\mathbf{w}_2\right)-\cdots\right)$$

We will look at the same benchmarks and risk regimes as in the previous section. Solutions will vary according to the risk aversion parameters attached to both benchmarks. Results are summarized in Table 6. In the case of decision makers exhibit low risk aversions ($\lambda_s = 1$) we find the solutions to focus on assets with the highest returns. If the import coverage benchmark becomes more important (higher penalty term for relative risk) allocations tend to become closer to this benchmark (assumed to be 50% USD and 50% Yen). In the case of both risk aversion parameters equal 30 (which can be interpreted as extremly risk averse) we arrive at the intermediate solution that invests into the equal weighted benchmarks. If on the other extreme central bankers put an overwhelming emphasis on the avoidance of wealth benchmark relative risk, we naturally come close to this benchmark (81% weighting in local cash). Obviously this is a solution favoured only by a small number of central banks. We should note that the case of extreme risk aversion with regard to both benchmarks is relatively close to the minimum variance solution. We see that the optimal solution depends on the degree of risk aversion to the respective benchmark relative risks. Hence the above optimization framework offers a straightforward way to arrive at the optimal currency allocation given that risk aversions with respect to benchmark relative risks differ. To assess which importance has been given to the rivalling risk regimes we can rerun the optimization above, this time however with only one risk regime, namely normal tines. The results are outlined in Table 7. When risks are considerably lower (normal times), it is optimal to allocate to the maximum return currency (99% in UK Pound for the used data set, time period and currency perspective). However for high risk aversions the results are almost the same as in Table 6. If risk aversions are sufficiently high, the level of risk does not matter. In general, there is a clear tendency to allocate more to the high return asset classes. All other results remain qualitatively the same.

	$\lambda_2=1$	$\lambda_2=3$	$\lambda_2=10$	$\lambda_2=30$	
$\lambda_1=1$	0%	7%	25%	35%	USD
	53%	32%	17%	9%	Pound
	12%	32%	39%	43%	Yen
	21%	20%	10%	6%	AUSD
	0%	0%	0%	0%	FRK
	14%	9%	9%	7%	EUR
$\lambda_1=3$	0%	8%	23%	33%	USD
	29%	20%	11%	6%	Pound
	10%	24%	32%	38%	Yen
	13%	12%	7%	4%	AUSD
	0%	0%	0%	0%	FRK
	48%	37%	27%	19%	EUR
$\lambda_1=10$	0%	8%	20%	29%	USD
	15%	11%	6%	3%	Pound
	8%	17%	25%	31%	Yen
	8%	7%	4%	2%	AUSD
	0%	0%	0%	0%	FRK
	69%	58%	46%	34%	EUR
$\lambda_1=30$	0%	6%	15%	23%	USD
	9%	6%	3%	2%	Pound
	5%	11%	18%	25%	Yen
	5%	4%	2%	1%	AUSD
	0%	0%	0%	0%	FRK
	81%	72%	61%	48%	EUR

Table 6. Pareto Optimality with dual benchmarks and dual risk regimes

[5] See Wang (1999) for the use of standard portfolio optimizers in multiple benchmark optimization.

	λ_2=1	λ_2=3	λ_2=10	λ_2=30	
λ_1=1	0%	0%	20%	35%	USD
	99%	52%	27%	17%	Pound
	0%	20%	39%	43%	Yen
	0%	28%	14%	8%	AUSD
	0%	0%	0%	0%	FRK
	0%	0%	0%	0%	EUR
λ_1=3	0%	0%	8%	25%	USD
	82%	85%	44%	27%	Pound
	0%	0%	32%	39%	Yen
	2%	15%	16%	9%	AUSD
	0%	0%	0%	0%	FRK
	16%	0%	0%	1%	EUR
λ_1=10	0%	0%	5%	21%	USD
	28%	42%	32%	18%	Pound
	0%	2%	22%	31%	Yen
	1%	8%	10%	6%	AUSD
	0%	0%	0%	0%	FRK
	7%	48%	31%	24%	EUR
λ_1=30	0%	0%	6%	18%	USD
	16%	23%	18%	11%	Pound
	0%	4%	16%	24%	Yen
	1%	5%	6%	3%	AUSD
	0%	0%	0%	0%	FRK
	83%	68%	54%	44%	EUR

Table 7. Pareto Optimality with dual benchmarks and single (lnormal) risk regime

5. Pareto Optimality and Non-Normality

Although we established non-normality as one of the stylized facts in currency markets we have not incorporated it yet. We assume investors exhibit constant relative risk aversion of the form $U_m^j = \frac{1}{1-\gamma_j} W_m^{1-\gamma_j}$ where $W_m = \sum_{n=1}^{N} w_n \left(1+R_{mn}\right)$ denotes wealth in return scenario $m = 1 \cdots M$, R_{mn} reflects the return of asset n in scenario m, and γ_j is the risk aversion parameter for utility function j. Problem (1.5) under general return assumptions becomes then

(1.7)
$$\max_{\mathbf{w}, \mathbf{w} \geq 0} \left(\min_j \left[\frac{1}{m} \sum_{m=1}^{M} U_m^j \right] \right)$$

which is equivalent to maximizing the minimum expected utility. Currencies that exhibit positive skewness (Yen) will be will be favoured relative to currencies that show negative skewness (UK Pound). The results are shown in Table 8 and Table 9. We see that the Yen (positively skewed) is favoured versus Pound and Australian Dollar (both negatively skewed.) relative to the solution assuming normality. For very high risk aversions both solutions more or less coincide.

	$\gamma_2=3$	$\gamma_2=10$	$\gamma_2=30$	
$\gamma_1=3$	0%	0%	27%	USD
	82%	54%	20%	Pound
	0%	22%	42%	Yen
	18%	24%	11%	AUSD
	0%	0%	0%	FRK
	0%	0%	0%	EUR
$\gamma_1=10$	0%	1%	22%	USD
	22%	45%	23%	Pound
	1%	19%	36%	Yen
	7%	15%	10%	AUSD
	0%	0%	0%	FRK
	61%	21%	9%	EUR
$\gamma_1=30$	1%	0%	14%	USD
	22%	22%	15%	Pound
	1%	10%	24%	Yen
	4%	7%	7%	AUSD
	1%	0%	0%	FRK
	72%	61%	40%	EUR

Table 8. Pareto Optimality under Normality

	$\gamma_2=3$	$\gamma_2=10$	$\gamma_2=30$	
$\gamma_1=3$	0%	5%	32%	USD
	18%	19%	7%	Pound
	71%	63%	55%	Yen
	11%	13%	5%	AUSD
	0%	0%	0%	FRK
	0%	0%	0%	EUR
$\gamma_1=10$	1%	3%	25%	USD
	26%	23%	12%	Pound
	31%	43%	47%	Yen
	5%	8%	5%	AUSD
	0%	0%	0%	FRK
	38%	23%	11%	EUR
$\gamma_1=30$	0%	2%	17%	USD
	10%	11%	8%	Pound
	12%	21%	32%	Yen
	1%	4%	3%	AUSD
	0%	0%	0%	FRK
	76%	62%	41%	EUR

Table 9. Pareto Optimality under Non Normality

6. Summary

The optimal currency allocation problem is descibed as a multiple benchmark optimization problem. In contrast to traditional solutions we showed how a maxmin approach can be successfully applied to the currency allocation problem, including alternative risk aversions, risk regimes, and benchmarks. We also extended the above analysis to incorporate nonormality in return data using scenario optimizaton as the most general form of porfolio optimization. The proposed methodology is equally applicable to central banks of developed as well as developing countries as it allows complete freedom in the specification of benchmarks, relative importance of benchmark relative risks as well as risk regimes and normality assumptions. It is exactly the definition of these parameters that will determine the optimal solution for any given central bank.

Appendix: Selected S-Plus Code

Rivalling Risk Regimes

```
#  function to calculate hectic versus normal correlation matrices and
   standard deviations
#  methodology by Chow/Kritzman (1999)
#  Dr. Scherer, November 2001

hectic.vs.normal<-function(datamatrix, percentage){
# inputs
# expects rectangular datamatrix (no missing values)
# first row contains series names
# dimension T x k
# no first date column
# perecentage defines percentages of bad times
# dimension scalar 0.9 means 90%
# outputs
# calculates kxk correlation and covariance matrices
# and kx1 volatility vectors

series.names<-names(datamatrix)
# calculate distribution from datamatrix
covar<-var(datamatrix)
mean<-as.matrix(apply(datamatrix, 2, mean),ncol=1)
# set up distance vector and calculate distances
distance<-matrix(0,ncol=1, nrow=nrow(datamatrix))
for(i in 1:nrow(datamatrix)){distance[i]<-(datamatrix[i,]-
   mean)%*%solve(covar)%*%(t(datamatrix[i,])-mean)}
# calculate correlation matrices
normal<-matrix(datamatrix[distance<=qchisq(percentage, ncol(datamatrix))],
   ncol=ncol(datamatrix))
stdev.normal<-(apply(normal,2,stdev))
names(stdev.normal)<-series.names
hectic<-matrix(datamatrix[distance>qchisq(percentage, ncol(datamatrix))],
   ncol=ncol(datamatrix))
stdev.hectic<-(apply(hectic,2,stdev))
names(stdev.hectic)<-series.names
cor.normal<-cor(normal)
cor.hectic<-cor(hectic)
cov.normal<-var(normal)
cov.hectic<-var(hectic)
dimnames(cor.normal)<-list(series.names, series.names)
dimnames(cor.hectic)<-list(series.names, series.names)
dimnames(cov.normal)<-list(series.names, series.names)
dimnames(cov.hectic)<-list(series.names, series.names)

# create list of outputs
list("normal.correlation" = cor.normal , "hectic.correlation" = cor.hectic,
   "stdev.normal" = stdev.normal, "stdev.hectic"=stdev.hectic, "normal
   covariance" = cov.normal, "hectic covariance"=cov.hectic)
}
```

Pareto Optimality and Dual Benchmark Optimisation

```
model.pareto<-function(mean.return, cov.1, cov.2, bench.1, bench.2,
   lambda.1, lambda.2){
dimnames(mean.return)<-NULL
dimnames(cov.1)<-NULL
dimnames(cov.2)<-NULL
dimnames(bench.1)<-NULL
dimnames(bench.2)<-NULL
# index
I <- Set()
J <- Set()
i <- Element(set=I)
j <- Element(set=J)
# parameter
Q1 <- Parameter(cov.1, index = dprod(i, j))
Q2 <- Parameter(cov.2, index = dprod(i, j))
b1 <- Parameter(list(1:length(mean.return), bench.1), index=i)
b2 <- Parameter(list(1:length(mean.return), bench.2), index=i)
lambda.1<-Parameter(lambda.1)
lambda.2<-Parameter(lambda.2)
r.bar <- Parameter(list(1:length(mean.return), mean.return), index=i)
# variable
w <- Variable(index = i)
U.min<-Variable()
# expressions for tracking error (te)
sigma.1 <- Expression(index = i)
sigma.2 <- Expression(index = i)
sigma.3 <- Expression(index = i)
sigma.4 <- Expression(index = i)
sigma.1[j] ~ Sum((w[i]-b1[i]) * Q1[i,j], i)
sigma.2[j] ~ Sum((w[i]-b1[i]) * Q2[i,j], i)
sigma.3[j] ~ Sum((w[i]-b2[i]) * Q1[i,j], i)
sigma.4[j] ~ Sum((w[i]-b2[i]) * Q2[i,j], i)
U.1 <- Expression()
U.2 <- Expression()
U.3 <- Expression()
U.4 <- Expression()
U.1 ~ Sum(r.bar[i]*(w[i]-b1[i]), i)-lambda.1*Sum((w[i]-b1[i])*sigma.1[i],i)
U.2 ~ Sum(r.bar[i]*(w[i]-b1[i]), i)-lambda.1*Sum((w[i]-b1[i])*sigma.2[i],i)
U.3 ~ Sum(r.bar[i]*(w[i]-b2[i]), i)-lambda.2*Sum((w[i]-b2[i])*sigma.3[i],i)
U.4 ~ Sum(r.bar[i]*(w[i]-b2[i]), i)-lambda.2*Sum((w[i]-b2[i])*sigma.4[i],i)
U.1 >= U.min
U.2 >= U.min
U.3 >= U.min
U.4 >= U.min
U <- Objective(maximize)
U ~ U.min
w[i] >= 0
Sum(w[i], i) == 1
}
system.pareto<-System(model.pareto, mean.return, cov.1, cov.2, bench.1, bench.2,
   lambda.1=0, lambda.2=30)
solve(system.pareto)
```

Bibliography

Boorman J and S Ingwes (2001), *Issues in Reserve Adequacy Management*, IMF working paper

Chow G., Jacquier E., Kritzman M. and K. Lowry (1999), "Optimal Portfolios in Good Times and Bad", *Financial Analysts Journal*, 55, pp.65-73

Ramaswamy S. (1999), *Reserve Currency Allocation: An Alternative Methodology*, BIS Working Paper No 72

Scherer B. (2002), *Risk Budgeting and Portfolio Construction* (Riskwaters: London)

Scobie H. and G. Cagliesi (2000), *Reserve Management* (Riskwaters: London)

Shectman P. (2000), "Multiple Benchmarks and Multiple Sources of Risk", Working Paper, Northfields

Wang M. (1999), "Multiple-Benchmark and Multiple Portfolio Optimization", *Financial Analysts Journal*, v55, pp63-72