Mean/Variance Analysis of Currency Overlays

Philippe Jorion

Global investors are now paying more attention to the management of the currency risks of their portfolios. Some have started to delegate currency management to "overlay" managers. These managers use currency futures and forwards to minimize the risks or maximize the returns of the underlying asset portfolios.

The overlay structure is inherently suboptimal because it ignores interactions between the assets in the underlying portfolio and exchange rates. Based on historical data, the efficiency loss appears to be on the order of 40 basis points for equity portfolios. This loss, of course, must be balanced against any excess returns that may be generated by specialized overlay managers.

Simulated results over the 1978–91 period indicate that overlay management can add value to global equity and bond portfolios. It cannot enhance performance by as much as an integrated approach to currency management, however.

With the accelerating trend toward international investing, U.S. investors are paying more attention to the impact of currency risk on their portfolio returns. Believing that many international equity managers do not have sufficient expertise in exchange rates, some institutional investors have turned to specialized "overlay" managers. They delegate selection of the "core" portfolio to a primary equity manager, either active or passive, but hire an expert to manage the currency risk of the portfolio separately. Out of a total of about $200 billion of U.S. pension funds invested abroad, about $50 billion are now actively managed as overlay portfolios.[1]

These developments raise several questions. Should currencies play a role in global portfolios? Is the delegation of currency management seriously suboptimal? Is there convincing evidence of predictability that would support the use of tactical currency allocation? This article reviews the theoretical and empirical arguments and analyzes conditions under which separate management of two asset classes, such as equities and currencies, is desirable.[2]

Three approaches to currency management are considered in a mean/variance framework—(1) a *joint*, full-blown optimization over the underlying assets (stocks or bonds) and currencies, (2) a *partial* optimization over the currencies, given a predetermined position in the core portfolio, and (3) a *separate* optimization over currencies. We compare the portfolio allocations derived from these approaches and analyze conditions under which the second and third approaches are globally optimal.

Approach 1 assumes the manager has expertise in many asset classes and can structure a portfolio to account for correlations between assets and currencies. The performance of the portfolio optimized over both currencies and underlying assets can be compared with the performance of an optimal portfolio comprised of the underlying assets only. The difference measures the benefit (or cost) of managing currency risk within the global portfolio.

Approaches 2 and 3 assume currency management via an overlay program. Investors may turn to an overlay manager if they feel the core manager lacks expertise in currencies. For instance, equity managers may neglect exchange rates because currencies are less volatile than equities, hence likely to contribute less to value added. Also, many equity managers may be regional specialists, or may focus on micro rather than on macro factors in the stock-picking process. In partial optimization, Approach 2, currencies are managed separately from the core portfolio, but the manager still controls *total* portfolio risk. In separate optimization, Approach 3, currencies are managed completely independently of the rest of the portfolio, and their performance is measured against a separate benchmark—cash, for instance.

Philippe Jorion is Associate Professor of Finance at the Graduate School of Management of the University of California at Irvine.

As we show, the mean/variance approach provides a unified framework for evaluating different approaches to currency management. It also includes, as special cases, "unitary" and "universal" currency hedging. Unitary ("full") currency hedging has attracted considerable attention because it has been advocated as a "free lunch," offering reduced risk with no commensurate reduction in returns.[3] Proponents of full currency hedging argue that international portfolio performance should be measured against a benchmark that is always fully hedged against currency risks.[4] In contrast, "universal" currency hedging implies that foreign asset positions should be only partially hedged, with the optimal hedge ratio being common to all assets and investors.[5]

More generally, this article discusses conditions under which currency hedging adds value to global portfolios, using historical data on major stock, bond and currency markets over the 1978–91 period. Currency positions are kept constant throughout the sample period. This may understate the benefits from currency management if active managers systematically outperform their benchmarks. To account for this possibility, the article presents an active management strategy, using a foreign exchange anomaly in the context of a disciplined investment management program.

OPTIMIZING GLOBAL PORTFOLIOS

The portfolio decision problem for global investments involves a joint choice over the underlying assets and currencies. In situations where there are no restrictions whatsoever on positions, closed-form solutions can be derived for the optimal portfolio. The appendix gives these solutions. We discuss here the intuition behind the results.

In a mean/variance framework, investors choose investment weights, w, so as to maximize an objective function that is positively related to the portfolio mean, μ_p, and negatively related to the portfolio variance, σ_p^2. This tradeoff between return and risk is reflected in the function $U(\mu_p, \sigma_p^2)$. In what follows:

p = the total portfolio,
x = the underlying assets (stocks or bonds) and
f = the currency forward contracts.

We can now summarize the three implementations of currency management. In *a joint optimization*, Approach 1, positions in assets and currencies are determined simultaneously so as to optimize the tradeoff between risk and return for the portfolio as a whole. In this case, the problem is stated:

$$Max_{w_x, w_f} U(\mu_p, \sigma_p^2). \tag{1}$$

The optimal positions are $w_x^*(x, f)$ and $w_f^*(x, f)$ where the optimal hedge positions, $w_f^*(x, f)$, generally depend on the optimal asset positions, $w_x^*(x, f)$, which themselves are affected by the presence of currencies in the portfolio. Also, the hedge positions contain both a speculative return component, driven by expected excess returns on currency forwards, and a variance-reduction component.

A *partial optimization*, Approach 2, is conditioned on predetermined underlying asset positions. The asset weights, w_x, are first determined optimally without regard to the hedges. The currency weights are then optimally determined, given w_x. This two-step approach is described by:

$$\begin{pmatrix} Max_{w_x} U(\mu_x, \sigma_x^2) \\ Max_{w_f} U(\mu_p, \sigma_p^2|w_x) \end{pmatrix}. \tag{2}$$

The positions can be written as $w_x^*(x)$ and $w_f^*(f|x)$. As before, the optimal hedge positions depend on the positions in the core portfolio.

In *a separate optimization*, Approach 3, asset and hedge positions are determined independently. The weights are found by independently solving two optimization problems:

$$\begin{pmatrix} Max_{w_x} U(\mu_x, \sigma_x^2) \\ Max_{w_f} U(\mu_f, \sigma_f^2) \end{pmatrix}. \tag{3}$$

The optimal asset and hedge positions are $w_x^*(x)$ and $w_f^*(f)$. Neither the core position nor the currency position depends on the other asset class.

Currency Overlays are Suboptimal

Clearly, going from Approach 1 to Approach 3 is successively less optimal. Figure A illustrates the loss of efficiency. It displays the performances of portfolios built using these three approaches. The portfolios consist of five bonds and four forward contracts (data presented below) and assume no restrictions on the positions.

The curve on the left represents the efficient portfolios obtained under the joint optimization scenario. It can be obtained from maximizing the investor's utility function over nine parameters (the investment weights for the five bonds and four forward contracts) for different levels of risk aversion. By definition, this line dominates all other approaches.

The curves in the middle and on the right correspond to the partial and separate optimiza-

Figure A. International Bond Portfolios with Different Optimization Methods, 1978–91

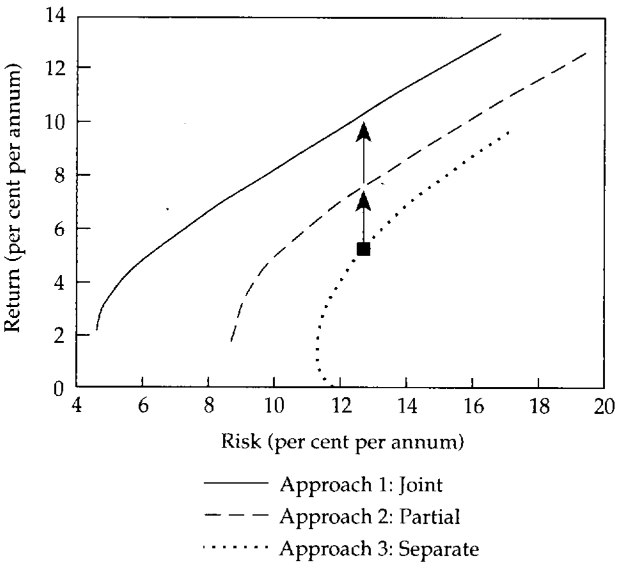

tions, respectively. In both approaches, the optimal bond portfolio, represented by the square, is first determined by maximizing the Sharpe ratio of the bond portfolio. With a partial optimization, the overlay manager then chooses the portfolio of currencies that provides the best portfolio risk/return combination. Because the bond position has already been determined, the manager cannot fully exploit interactions between bonds and currencies. The curve representing the second efficient set must therefore be inside the first efficient set, reflecting the suboptimality of the partial optimization approach.

The curve on the right, representing the separate optimization case, is also suboptimal. Here also bond positions are predetermined, and the optimization focuses on the risk and return of the currency portfolio. Extremely risk-averse managers will choose to have no position in currencies; their overall portfolios will coincide with their bond portfolios. Of course, this choice will not minimize total portfolio risk. This illustrates the basic inconsistency between focusing on currencies as a stand-alone portfolio and incorporating them as a part of the overall portfolio.

In fact, optimizing separately is equivalent to managing in terms of *relative*, rather than absolute, returns. Roll shows that a manager who tries to minimize the variance of total tracking error, relative to a benchmark, will not produce an overall mean/variance efficient portfolio.[6] Implementing a separate optimization is the least efficient method of managing currencies and should be ruled out in favor of partial optimization.

The remaining question is whether partial optimization can provide a close substitute for the more general model. The appendix shows that positions that solve Objectives 1 and 2 will be identical if returns on underlying assets, measured in dollars, are uncorrelated with exchange rates. In this case, the currency positions are equal and there is no loss of efficiency from optimizing assets and currencies separately. Furthermore, if expected returns on currencies are zero, the currency positions are zero, and there is no reason to invest in currencies.

Currency returns in our context are payoffs from positions in forward contracts. For them to equal zero, the forward rate must be an unbiased forecast of future spot rates; that is, "uncovered interest rate parity" holds. Whether this assumption is true is subject to active debate, to which the following section is devoted.

The first assumption—that dollar returns on underlying assets are uncorrelated with exchange rates—is unlikely to be met in practice. If the underlying assets are foreign stocks or bonds, then their *dollar* returns are likely to be correlated with dollar exchange rates. A less unrealistic assumption is that local-market returns are uncorrelated with exchange rates. In this case, unitary hedging is optimal, or:

$$w_f^*(x, f) = -w_x^*(x, f),$$

and each foreign market should be fully hedged against currency risk.

Unitary hedging is therefore optimal if local-currency returns are unrelated to exchange rates and if currency returns are expected to be zero. Under these stringent assumptions, full currency hedging provides less risk at no cost. Note, however, that the partial-optimization approach is still inefficient unless core positions are determined on the basis of currency-hedged returns.

The Zero Risk Premium

It is important to emphasize that the assumption of a zero risk premium drives the "free lunch" argument first advanced by Pérold and Schulman.[7] One argument posits that there should be no long-term payoff from buying a forward contract on another currency, because any gains must be offset by losses to the counterparty. Because of this symmetry, and because forward contracts are in zero net supply, currencies should offer no risk premium.

This argument is incorrect, as an analogy with stock index futures shows. Stock index futures are also in zero net supply, but they are linked to the underlying stocks through a cost-of-carry relationship: When held to maturity, a long position in futures is equivalent to a long position in the underlying cash instruments. But it is generally accepted that stocks generate long-term excess returns of about 5% to 10% annually. This risk premium must also be embedded in long positions in stock index futures. The fact that a contract is in zero net supply thus does not necessarily imply a zero risk premium.

International asset pricing models, such as the IAPM developed by Solnik, actually show that, in equilibrium, risk premiums depend on investors' risk aversion and on whether countries are net investors or net borrowers.[8] To simplify, assume a world with two investors only—U.S. and British. If British investors as a whole are net investors in the U.S., they will generally seek to reduce risk by hedging to some extent against exchange rate changes. They will thus sell dollars even if it involves a slight loss or if the forward price of the dollar is lower than the expected future spot price. In this world, U.S. investors must be net borrowers. Therefore, U.S. investors will seek to hedge by buying the pound forward, even if the forward price of the pound is above the expected future spot price. Thus, in equilibrium, there will be a non-zero expected return to forward contracts, despite the fact that for every Briton selling dollars, an American will be buying the currency. Currencies can very well, in equilibrium, be characterized by non-zero expected returns.

Black's "universal" hedge ratio has both speculative and risk-minimization motives and implicitly assumes non-zero risk premiums.[9] Black's results, as in the IAPM, derive from the aggregation of individual optimal portfolios across countries, but also from the assumptions that all investors have the same risk tolerance and that each national wealth is exactly equal to the value of each stock market. Under these assumptions, the hedge ratio reduces to the "universal" value $h = (1 - \lambda)$.

Admittedly, the risk premium on forward contracts can be positive or negative, unlike the risk premium on stocks, and may be difficult to identify. It can even change sign if asset supplies change over time. Furthermore, the empirical evidence on expected returns on forward contracts seems to indicate that, on average, returns are close to zero. Table 1 reports average returns on forward contracts, measured over 1978–91; because these positions involve no net investment, these are excess returns. Over this period, returns ranged from −0.16% (deutschemarks) to 3.01% (pounds sterling) per annum.

None of the returns is statistically different from zero. This, however, may have more to do with the power of the tests than with a truly zero risk premium. From an economic perspective, the numbers are significant because they are of the same magnitude as the value added we would expect from active management. An added value of a few percentage points per annum is more than satisfactory. Yet conventional hypothesis tests with 14 years of data are not powerful enough to detect statistical significance: Assuming a 12% annual volatility for the pound, a t-test yields $3.01/(12/\sqrt{14}) = 0.9$.[10] To attain a t-statistic above 2, which would correspond to the conventional 5% significance level, one would need 64 years of data. Over this period, however, a dollar invested at the excess return of 3% would have grown to \$6.6! Clearly, waiting for statistical significance entails a very significant opportunity cost.

Properly interpreted, the statistics tests indicate that the evidence on average returns is inconclusive, not that currency forwards have zero expected returns. In addition, realized returns on currencies appear to vary over time, from positive to negative values; measuring average returns might very well hide predictable temporal variations in return.

EMPIRICAL EVIDENCE

Instead of assuming a zero risk premium and a simplified covariance matrix structure, we use actual estimates of return and risk measured over the 1978–91 period. Five stock and bond markets are considered—the U.S., Japan, Germany, Britain and France. Tables 1 through 3 present estimates of average excess returns, risks and correlations for currencies, stocks and bonds, respectively. Corre-

Table 1. Annualized Excess Returns, Risks and Correlations of Forward Contracts, 1978–91

	Yen	Mark	Pound	Franc
Return	1.76%	−0.16%	3.01%	2.03%
Risk	13.40	13.07	12.67	12.47
Correlations				
JY	1.00			
DM	0.65	1.00		
BP	0.58	0.72	1.00	
FF	0.67	0.96	0.72	1.00

Table 2. Excess Annualized Returns and Risks of Stocks, 1978–91

	$	Yen	Mark	Pound	Franc
Return	7.63%	12.38%	8.02%	11.52%	13.47%
Risk	16.18	24.19	23.39	22.36	25.10
Correlations					
$	1.00				
Yen	0.24	1.00			
Mark	0.35	0.39	1.00		
Pound	0.57	0.41	0.48	1.00	
Franc	0.40	0.37	0.55	0.46	1.00
Unhedged					
Yen	−0.03	0.64	0.32	0.27	0.35
Mark	−0.03	0.32	0.53	0.29	0.44
Pound	−0.01	0.32	0.37	0.53	0.30
Franc	0.00	0.34	0.50	0.32	0.49
Hedged					
Yen	−0.03	0.12	−0.04	−0.07	0.02
Mark	−0.03	−0.05	−0.03	−0.13	−0.04
Pound	−0.01	−0.01	−0.04	−0.04	−0.07
Franc	0.00	−0.04	−0.04	−0.10	−0.01

lations between currencies and both unhedged and hedged returns are reported. All returns are translated into U.S. dollars and measured in excess of the U.S. risk-free rate (taken as the one-month Treasury bill rate).

Unhedged stock returns appear to be positively correlated with exchange rates, with average correlation coefficients ranging from 0.3 to 0.6. Therefore, optimizing currencies and underlying unhedged assets separately cannot be optimal. In contrast, hedged stock returns appear uncorrelated with currencies. Full currency hedging may be an acceptable strategy for stocks.

Table 3. Excess Annualized Returns and Risks of Bonds, 1978–91

	$	Yen	Mark	Pound	Franc
Return	2.38%	3.18%	1.96%	1.46%	0.97%
Risk	10.51	6.63	5.84	10.13	6.05
Correlations					
$	1.00				
Yen	0.42	1.00			
Mark	0.52	0.58	1.00		
Pound	0.40	0.43	0.40	1.00	
Franc	0.34	0.40	0.51	0.32	1.00
Unhedged					
Yen	0.19	0.94	0.61	0.49	0.65
Mark	0.22	0.61	0.94	0.55	0.88
Pound	0.18	0.55	0.67	0.84	0.64
Franc	0.22	0.64	0.90	0.52	0.91
Hedged					
Yen	0.19	0.40	0.23	0.16	0.15
Mark	0.22	0.28	0.34	0.08	0.10
Pound	0.18	0.25	0.24	0.26	0.04
Franc	0.22	0.29	0.33	0.04	0.09

For bonds, we observe high correlations for both unhedged and hedged returns. The positive correlations for hedged bonds suggest that bond yields are negatively associated with the dollar value of foreign currencies. This implies that currency forwards can to some extent cross-hedge changes in foreign yields. Partial optimization and full hedging must be suboptimal for international bonds.

The question of interest, however, is how much efficiency is lost by separately optimizing the core portfolio and currencies in a realistic framework. A loss of a few basis points would cause no concern and possibly be more than offset by the value added from specialized currency managers. A loss of a few percentage points, however, would raise serious doubts about the currency overlay approach.

We consider the following three situations

- a full optimization over N + 1 assets (stocks or bonds, including U.S. assets) plus N currencies, compared with a full optimization over N + 1 assets only;
- a partial optimization over N currencies, given an optimal position of N + 1 assets; and
- a separate optimization over N currencies, given an optimal position in underlying assets.

Tables 4 and 5 present the optimal positions under each of these three approaches. To be realistic, we allow no short selling for the underlying assets and allow positions in the dollar, yen and European currencies to vary only between zero and the total portfolio value. These constraints are typical of restrictions placed on overlay managers, who are generally allowed to cross-hedge European currencies.

The tables report portfolio excess return, volatility and Sharpe ratio. For stocks, adding currencies to an optimized portfolio of stocks and reoptimizing increases the Sharpe ratio by 0.116 (= 0.805 − 0.689). At a typical 15% volatility level, this translates into added value of 173 basis points. Adding currencies and performing a separate optimization, given fixed positions in stocks, increases performance by 92 (= 0.781 − 0.689), which translates into 139 basis points of value added. Optimizing separately on currencies leads to an increase of only 72 basis points.

The 173 basis points, however, should be

Table 4. Optimal Positions with Stocks and Currencies, 1978–91

	No Currency	With Currency Joint	Partial	Separate
Return	11.1%	12.9%	11.8%	12.2%
Risk	16.1	16.1	15.1	16.6
Sharpe Ratio	0.689	0.805	0.781	0.735
Value Added (at 15% risk)		1.73	1.38	0.72
Positions (%)				
Stocks:				
U.S.	29*	3*	29	29
Japan	29*	39*	29	29
Germany	0*	4*	0	0
U.K.	16*	18*	16	16
France	26*	36*	26	26
Currency:				
Yen	–	−39*	−29*	0*
Mark	–	−58*	−42*	−42*
Pound	–	30*	28*	8*
Franc	–	15*	13*	40*
Net Foreign Exchange Exposure	71%	45%	41%	77%

* Denotes an optimized position.

interpreted with caution. Because the same period is used for the optimization and the performance measurement, adding assets must by definition improve portfolio performance. The number is biased upward because, at worst, the optimizer could take no position in currencies.[11]

The estimation error inherent in these numbers can be recognized explicitly by test statistics of performance improvement.[12] By comparing the observed value of the test statistic with a distribution obtained under the null hypothesis that currencies do not improve performance, one can gauge the extent to which the observed number may be due to sampling variability in the data. In this case, the statistic was exceeded in 23% of the samples, leading us to conclude that the performance-improvement numbers are in line with what could be expected because of chance alone. There is no statistical evidence that currencies add value to global equity portfolios.

The 173 basis points should thus be considered an upper bound on the value added from static management of currencies in equity portfolios. Because the same data and degrees of freedom are used in the partial and separate optimizations, one can conclude that partial optimization decreases returns by about 35 basis points, and that separate optimization decreases returns by a further 66 basis points. These are nonnegligible costs.

Including currencies in a fully optimized bond portfolio adds 273 basis points, given the 10% volatility typical of bond indexes. Currencies add 102 and 37 basis points, respectively, when managed using partial and separate optimizations. Because 273 is much higher than the 173 basis points reported above, it appears that fixed-income portfolios benefit much more than equity portfolios from the management of currencies. The reason for this can be traced to the high correlations between bond returns and currencies, which lead to substantial portfolio adjustments when currencies are included.

Table 5. Optimal Positions with Bonds and Currencies, 1978–91

	No Currency	With Currency Joint	Partial	Separate
Return	3.7%	3.6%	3.6%	4.2%
Risk	11.3	5.9	8.2	11.7
Sharpe Ratio	0.331	0.604	0.432	0.366
Value Added (at 10% risk)		2.73	1.02	0.37
Positions (%)				
Bonds:				
U.S.	44*	0*	44	44
Japan	39*	48*	39	39
Germany	0*	0*	0	0
U.K.	17*	3*	17	17
France	0*	49*	0	0
Currency:				
Yen	–	−48*	−39*	0*
Mark	–	−52*	−17*	−17*
Pound	–	18*	12*	7*
Franc	–	−9*	7*	16*
Net Foreign Exchange Exposure	56%	9%	19%	62%

* Denotes an optimized position.

In Table 5, for instance, the optimal bond portfolio is drastically altered when currencies are added. The initial composition of 44% U.S. bonds, 39% Japanese and 17% British becomes 48% Japanese, 3% British and 49% French, with positions in Japanese bonds fully hedged; a cross-hedge is also implemented, short the mark and long other European currencies.

These reallocations, much larger than for equity portfolios, translate into a higher value added, 273 versus 173 basis points. Furthermore, the 273 points are too large to be attributed to chance. The same test methodology as used above indicates that adding currencies to bond portfolios significantly improves performance.

Finally, the large drop from 273 to 102 basis points when the bond portfolio is predetermined demonstrates that it makes little sense to use currency overlays for bond portfolios. Given the high correlation between bonds and currencies, these two asset classes must be managed together. This is why, in practice, currency overlays are only applied to equity portfolios.

The last lines in Tables 4 and 5 report the net foreign currency positions. For instance, the net foreign exchange exposure of fully optimized stock portfolios is 71% (29% yen, 16% pound sterling, 26% franc). When adding currencies, the foreign exchange exposure of the portfolio decreases from 71% to 45% for the joint optimization.[13] Because hedging entailed an opportunity cost over this period, there was never full hedging into dollars. Hedging also appears to be far from universal. For instance, currency positions in the yen and mark are generally negative, while positions in the pound and franc are generally positive. This suggests that unitary hedging and universal hedging are both suboptimal.

Active Hedging

These results may be of little relevance to actual investment decision rules, as the optimal positions are derived from "ex post" data. The results suggest that, over this period, it would have been possible to enhance portfolio performance by using forward contracts; unfortunately, the optimal weights are revealed only after the fact. The problem is compounded by the apparent instability of the currency hedges. For instance, in the early '80s it would have been generally advantageous to hedge foreign currency exposures back into dollars, since the dollar appreciated, but the reverse was true later, when the dollar depreciated sharply.

Does the performance improvement survive an "ex ante" rule, where all investment decisions are based on prior information? Furthermore, is it possible to identify changes in expected returns and incorporate these changes into active management of currency allocations?

Previous research on the efficiency of the foreign exchange market has shown that the forward premium can help in predicting expected returns on forward contracts.[14] Because, by interest rate parity, the forward premium is also the interest rate differential, the evidence generally implies that investors should go long high-interest-rate currencies and short low-interest-rate currencies. The following discussion shows how to exploit this anomaly using a mean/variance optimizer.

The strategy is implemented as follows. Consider a passive benchmark actively hedged with four forward contracts. This is suboptimal but easy to implement, as bid–ask spreads in the foreign exchange market are very low, while stock and bond positions, more costly to alter, are fixed. The decision variables are the amounts to buy or sell in the four forward contracts.

Each month, expected returns and risk measures are estimated from a four-year moving window. Returns are forecast from the estimated coefficients of a regression of past returns on the forward premium, combined with the most recent forward premium. The variance/covariance matrix is also estimated over the same time period. These parameters are fed into a portfolio optimizer that determines the forward positions of the portfolio with the highest excess return-to-risk ratio. The performance of the optimal portfolio is recorded for the following month, after which the process is repeated. The strategy is realistic in that investment decisions are based only on prior information, which is kept current, and only feasible positions are implemented.

Table 6 reports the results from this active hedging rule, as implemented over the 1978–91 period. The performance of the world stock index with actively managed currency exposure is compared with the performances of the unhedged and fully hedged passive indexes. Clearly, the currency rule adds substantial value. Assuming 15% volatility, excess returns are increased by 2.3% annually relative to the unhedged index and by 3.1% relative to the fully hedged index. The enhanced performance of active hedging more than offsets the transaction costs incurred by using forward contracts, which are very modest.[15] In addition, the results seem consistent across the

Table 6. Active Currency Hedging

	World Stock Index		
	Unhedged	Fully Hedged	Actively Hedged
1978–1991			
Average Annual Return	8.11%	6.81%	12.14%
Annual Volatility	15.21	14.11	17.54
Sharpe Ratio	0.533	0.483	0.692
1978–1984			
Average Annual Return	3.50	6.31	6.20
Annual Volatility	13.25	11.70	15.46
Sharpe Ratio	0.264	0.539	0.401
1985–1991			
Average Annual Return	12.77	7.59	18.22
Annual Volatility	16.72	16.08	19.11
Sharpe Ratio	0.764	0.472	0.953

two subperiods of dollar strength (1978–84) and dollar weakness (1985–91).

A caveat, however, is in order. For exchange rates measured against the dollar, risk can be reasonably well assessed from the volatility of historical series. Realignment risk in fixed-exchange-rate systems, however, may not be adequately represented by historical data. Because it underestimates devaluation risk, the forward-premium strategy applied to European cross-hedges tends to take large long positions in high-yield currencies offset by short positions in low-yield currencies, which can unravel with currency realignments. Over 1978–91, the strategy excess returns indicate that realignments have been more than offset by forward-premium gains; forward-premium gains, however, accumulate slowly over time, whereas realignment losses are large and occur suddenly. To address this asymmetry, the strategy should be augmented by realignment probabilities, and perhaps diversified with technical rules.[16] Nevertheless, one would hope that this simple model provides a lower bound on the value added by overlay managers.

CONCLUSIONS

Currency overlays are generally suboptimal. Overlays are implemented in two ways—either by optimizing on total portfolio return and risk, or by optimizing separately on currencies. In the former case, the performance of the overlay is measured in conjunction with that of the core portfolio. In the latter, the only objective of the overlay manager is to add value, measured in absolute returns. In both cases, returns suffer vis-à-vis a policy of managing underlying assets and currencies in an integrated fashion. Furthermore, underperformance increases, the higher the correlation between currencies and the underlying assets.

Whether the underperformance is significant must be judged against the benefits from using specialized overlay managers. There is some evidence that returns on currencies are predictable. If so, the value added by overlay managers could outweigh the inherent inefficiency of the set-up.

At the very least, performance should be measured in the context of the total portfolio. This means that the core portfolio manager must regularly communicate positions to the overlay manager, who would then evaluate risk and return for the combined portfolio. In this context, the partial-optimization approach may be an acceptable, albeit second-best, solution to the management of currencies in global portfolios. But the onus is on overlay specialists to prove their worth as active managers.[17]

APPENDIX

The portfolio decision problem for global investment involves a joint choice over the underlying assets and currencies. Assume for simplicity that investors maximize the following objective function, reflecting the tradeoff between expected portfolio return (in excess of the risk-free rate) $\mu_p = w'\mu$ and risk $\sigma_p^2 = w'\Sigma w$:

$$\mu_p - (1/2\lambda)\sigma_p^2, \tag{A1}$$

where w represents positions, μ represents expected excess returns on all assets (with covariance matrix Σ) and λ is the investor's risk tolerance. Returns on all assets are measured in U.S. dollars in excess of the risk-free rate.

The optimal positions are given by $w^* = \lambda\Sigma^{-1}\mu$, with the remainder of the portfolio invested in the riskless asset. Further insights can be obtained by decomposing w^* into a core portfolio and currency portfolio. Partition μ and Σ into components that correspond to underlying assets, represented by x, and currency forwards, represented by f:

$$\mu = \begin{pmatrix} \mu_x \\ \mu_f \end{pmatrix}, \ \Sigma = \begin{pmatrix} \Sigma_{xx} & \Sigma_{xf} \\ \Sigma_{fx} & \Sigma_{ff} \end{pmatrix}. \tag{A2}$$

Then define

$$\beta = \Sigma_{ff}^{-1}\Sigma_{fx}$$

as the regression coefficients of the assets on the hedges and

$$\Sigma_{x \cdot f} = \Sigma_{xx} - \beta'\Sigma_{ff}\beta$$

as the covariance matrix of underlying asset returns conditional on the hedges. The optimal portfolio is then derived from the partitioned inverse of the matrix Σ, which can be written as:

$$\Sigma^{-1} = \begin{pmatrix} \Sigma_{x \cdot f}^{-1} & -\Sigma_{x \cdot f}^{-1}\beta' \\ -\beta\Sigma_{x \cdot f}^{-1} & \Sigma_{ff}^{-1} + \beta\Sigma_{x \cdot f}^{-1}\beta'. \end{pmatrix} \tag{A3}$$

We can now detail the optimal positions in currencies and underlying assets for three separate implementations of currency management. In *a joint optimization*, where w_x and w_f are determined simultaneously:

$$\begin{pmatrix} w_x^*(x,f) = \lambda(\Sigma_{x\cdot f}^{-1}\mu_x - \Sigma_{x\cdot f}^{-1}\beta'\mu_f) \\ w_f^*(x,f) = \lambda(\Sigma_{ff}^{-1}\mu_f) - \beta w_x^* \end{pmatrix}. \quad \text{(A4)}$$

Note that the optimal hedge positions w_f^* depend on the optimal asset positions w_x^*, which themselves are affected by the presence of currencies in the portfolio. Also, the hedge positions have a speculative component, driven by non-zero expected returns in currencies, as well as a variance-reduction component related to β.

A partial optimization is conditioned on predetermined underlying asset positions. When these are determined optimally without regard for the hedges, the weights are:

$$\begin{pmatrix} w_x^*(x) = \lambda(\Sigma_{xx}^{-1}\mu_x) \\ w_f^*(f|x) = \lambda(\Sigma_{ff}^{-1}\mu_f) - \beta w_x \end{pmatrix}. \quad \text{(A5)}$$

Note that, as before, the optimal hedge positions w_f^* depend on the positions in the core portfolio. Because the core positions do not account for hedges, neither core nor currency positions are globally optimal.

Partial optimization is efficient under the following conditions. Positions A4 and A5 will be identical if returns on underlying assets, measured in dollars, are uncorrelated with exchange rates ($\Sigma_{fx} = 0$, $\beta = 0$, $\Sigma_{x\cdot f} = \Sigma_{xx}$). In this situation, the globally optimal weights simplify to

$$w_x^*(x,f) = \lambda(\Sigma_{xx}^{-1}\mu_x)$$

and

$$w_f^*(x,f) = \lambda(\Sigma_{ff}^{-1}\mu_f).$$

In *a separate optimization*, asset and hedge positions are determined independently. The weights are:

$$\begin{pmatrix} w_x^*(x) = \lambda(\Sigma_{xx}^{-1}\mu_x) \\ w_f^*(f) = \lambda(\Sigma_{ff}^{-1}\mu_f) \end{pmatrix}. \quad \text{(A6)}$$

Here, neither the core nor the currency position depends on the other asset class.

FOOTNOTES

1. As reported by *Pensions & Investment Age*, May 17, 1993.
2. See, for instance, A. F. Lee, "International Asset and Currency Allocation," *Journal of Portfolio Management*, Fall 1987.
3. A. Pérold and E. Schulman, "The Free Lunch in Currency Hedging: Implications for Investment Policy and Performance Standards," *Financial Analysts Journal*, May/June 1988.
4. Indeed currency-hedged benchmarks have recently appeared in response to market demands. For instance, Morgan Stanley Capital International started to report currency-hedged international stock indexes in June 1989; Salomon Brothers has constructed currency-hedged bond indexes since March 1988.
5. See F. Black, "Equilibrium Exchange Rate Hedging," *Journal of Finance*, 1990, and also Adler and Prasad, "On universal currency hedges," *Journal of Financial and Quantitative Analysis*, 1992.
6. R. Roll, "A Mean-Variance Analysis of Tracking Error," *Journal of Portfolio Management*, 1992.
7. Pérold and Schulman, "The Free Lunch," *op. cit.*
8. See B. Solnik, "An Equilibrium Model of the International Capital Market," *Journal of Economic Theory*, 1974.
9. Black, "Equilibrium Hedging," *op. cit.*
10. Changing the observation interval in order to increase the number of observations would have no effect on these tests. The numbers reported in the table are annualized from monthly data, with returns multiplied by 12 and standard deviations multiplied by the square root of 12, assuming independence over time. Going back to monthly data, the t-test expressed in monthly terms is $(3.01/12)/((12/\sqrt{12})/\sqrt{14\times12})$, which is the same as before.
11. The bias resulting from the optimization has been noted by R. Michaud, "The Markowitz Optimization Enigma: Is Optimized Optimal?" *Financial Analysts Journal*, January/February 1989.
12. The statistic that represents the increase in performance is

$$F = (T - N_2)/(N_2 - N_1)(\theta_2^2 - \theta_1^2)/(1 + \theta_1^2),$$

where T is the number of observations, N_1 the restricted number of assets, N_2 the enlarged number of assets, and θ_1 and θ_2 are the maximum Sharpe ratios with N_1 and N_2 assets, respectively. The methodology employed here is explained in P. Jorion, "Portfolio Optimization in Practice," *Financial Analysts Journal*, January/February 1992. Formal statistical tests are presented in J. Glen and P. Jorion, "Currency Hedging for International Portfolios," *Journal of Finance* 48 (1993), 1865–86.
13. The latter number is obtained from the algebraic sum of positions in four foreign stock markets plus four currencies.
14. This was first shown by J. F. O. Bilson, "The 'Speculative Efficiency' Hypothesis," *Journal of Business* 54 (1981), 435–52. Recent attempts to explain these results in terms of a rational risk premium are reviewed in Hodrick, *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets* (Harwood Academic Publishing, 1987).
15. For instance, a typical two-way spread on the pound is 5 points, or about 0.03% in relative terms. So, rolling over 12 monthly swaps involves an annual transaction cost of 12(0.03%/2) = 0.18%.
16. See, for instance, R. J. Sweeney, "Beating the Foreign Exchange Market," *Journal of Finance* 41 (1986), 163–82.
17. I thank Michael Adler for his useful comments.