Prediction of Currency Movements

by Michael F. Wilcox, CFA

The Challenge

"Almost all [international managers having EAFE mandates from US pension funds] have negative value added in the currency management function," according to a recent report by Intersec Research Corp.[1] Over the 8-year period ending December 31, 1995, this group of managers, on average, added more than 2% per annum to the returns of the EAFE index through active management *despite* having squandered more than 1% per year in currency losses. They could have outperformed EAFE by more than 3% per year, in other words, if they had focused exclusively on stock selection and market allocation.

The evidence thus suggests that most money managers (in the US, anyway) are either poor currency forecasters or have not been successful in integrating their forecasts into their portfolios. This chapter will address both these issues, focusing first on a practical, easy-to-implement method for forecasting currency exchange rates, and then will provide some suggestions for using these forecasts in conjunction with the management of portfolios containing foreign securities.

The Response

In the context of institutional money management, currency management most often involves making a decision whether or not to hedge exposure to foreign currencies. These exposures have been acquired in the normal course of business, by investing in foreign securities.

Most economists agree that the concept of purchasing power parity (PPP) describes an important component of how foreign exchange (FX) rates ought, in theory, to be determined. There is much less agreement as to how well the theory describes real-world conditions. We will describe a PPP-based methodology that has proven its efficacy over time.

The procedures we describe in this chapter, by selectively hedging portions of foreign currency exposure, added 3% per year to the returns of an EAFE-like portfolio in a simulation we performed covering the same time period described by the Intersec study just mentioned.

For the most part, this chapter addresses currency hedging as if that were the only application of interest to our readers. We realize that there will be those, particularly in the fixed-income area, who have the latitude to undertake more aggressive strategies and invest in currencies that appear undervalued in the hope of adding value through the appreciation of those currencies. For such readers, we offer assurances that the success of the system described here is symmetrical — that is to say, it is just as good at predicting currency appreciation as it is at warning of potential depreciation.

Results of our simulation

Chart 1 shows the pattern of growth experienced by two portfolios containing the same underlying equity investments.

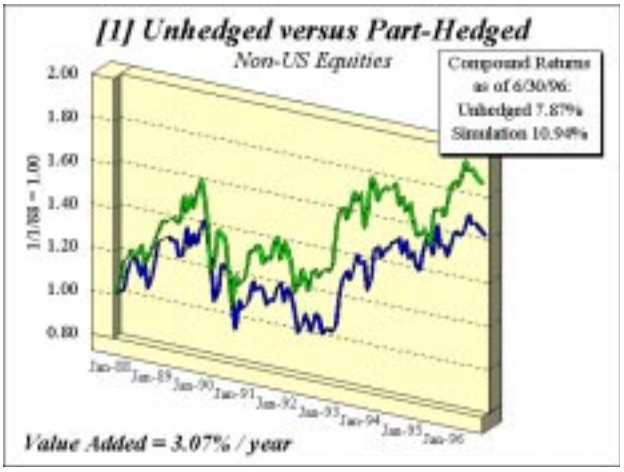

The dark line, which gained about 62% over the period from January 1, 1988 to June 30, 1996, represents a capitalization-weighted portfolio of the largest non-US equity markets: namely, Japan, the UK, Germany, France, Switzerland, Holland, and Italy. Because this group represents about 90% of the capitalization of EAFE, the performance of this portfolio will be, for our purpose of comparing with the Intersec results, very nearly the same as EAFE. We chose to use only markets where currency hedging is cost-effective for most money managers. Very large portfolios might benefit by hedging currencies in the smaller markets, but our goal was to demonstrate that the exercise will be beneficial even if done selectively.

The lighter line represents a portfolio that gained about 87% over this time period. The

superior growth of this "part-hedged" portfolio (which translates to a better than 3% per year advantage) is attributable to selective hedging of foreign (i.e. nondollar) currency exposures. When a currency appeared overvalued according to the algorithm described in this chapter, we hedged a portion of its exposure (again, in accordance with the methodology laid out below).

No transaction charges were assessed against the partly-hedged portfolio, but turnover is low under our guidelines, and spreads in the FX markets are tight, so hedging costs would, in reality, penalize performance by only a few basis points per year.

Our message here is that, in practice, a currency forecasting technique based on PPP works very well, though to be consistently successful, care must be taken to evaluate results only over long periods of time. By "long", we mean five years or more. The approach we describe is a valuation-based approach that is suitable only for those who have the temperament to make money, on average, only half the time. In the short run, there are many forces (such as international portfolio flows, central bank interventions, and governmental fiscal policy actions) that move the FX markets in what may seem, in retrospect, like perfectly understandable ways, but we have yet to discover a way to improve our PPP-based forecasts by taking into account these other economic forces.

Organization of the chapter

This chapter starts with a quick overview of relevant terminology, basic concepts, and some empirical observations, then reviews alternate theories for explaining currency movements, based on so-called structural models of exchange-rate determination. Most of the methods proposed by economists have proven to be impractical to implement in the context of money management, so the focus of the bulk of the chapter is on the frame of reference, already mentioned, that has proven its power in practice and in extensive backtesting — that of Purchasing Power Parity. The chapter concludes with some advice about applying forecasts to real-world portfolios to achieve the kind of results our simulation suggests are attainable. Practitioners with experience in the foreign exchange markets may well want to skip the review in the first part of the chapter and, after reading the next section relating to our terminology, turn to the section entitled "Implementing a PPP-based Forecasting System".

Forecasting Horizon and other Conventions used in this Chapter

For ease of exposition, all examples given in this chapter use a one-year forecasting horizon — this is a convenience that makes it easier to think about interest rates and inflation statistics, which are typically quoted as annual rates, but the reader should keep in mind that, in practice, the actual forecasting horizon of interest may be different (typically, shorter than a year). It is our expectation that forecasts made in the framework outlined here will come to pass in a matter of months rather than days, though there will of course be periods when the forecasts are directionally correct over time periods of only a few days or weeks. The principles enunciated in this chapter are best applied to time horizons of several months to several years, and it is our intention to provide a framework that is useful for money managers, whose investment horizon is fairly long, not necessarily for traders or speculators, who may have a more short-term orientation.

Also, in presenting examples of various investment strategies, we have in most cases rounded results off to the nearest percentage point or currency unit (cent or dollar, as appropriate, in illustrations involving the U.S.), unless more precision is needed for clarity.

We have used formulas and statistics sparingly, on the presumption that our readers are capable of grasping the essence of our methodology from the charts and descriptions. The process we describe is a highly quantitative one, but no one step is terribly complicated. In the few formulas we do provide (though not in all the illustrations and examples we provide), we have used the convention of quoting currency exchange rates in units of home currency per unit of foreign currency (usually called the "American convention").

$$FX = \frac{\text{Home Currency (\$)}}{\text{Foreign Currency}}$$

For dollar-based investors, this formula produces the dollar price of the foreign currency, but the reader should be cautious, because among the major currencies, only the pound sterling is typically quoted this way in the press, with most other currencies (except the Australian dollar) being quoted according to the "European convention" of units of currency per dollar. In the sometimes confusing terminology of the FX markets (akin to fixed-income investors who sometimes talk about

a bond going "up" by a certain amount of yield when they mean the price of the bond went up while the yield went down), a rise in the quoted rate of a currency usually means that the value of that currency (in dollar terms) went down.

The task at hand

Before we set out to forecast the future course of foreign exchange fluctuations, let us be clear about what it is that we are attempting to accomplish. If our job is to hedge foreign currency exposure, it is important to remember that predicting currency movements is of no value unless the actions we take in response to our forecasts result in returns that are superior to those attained by investors who undertake no hedging. Thus, our success should be measured in portfolio performance, not in statistics that tell us how closely our forecasts were correlated with actual rates at the close of our time horizon (though, of course, such measures can be useful in determining if we have produced usable information).

To be successful forecasters, we need not have perfect predictive ability -- we need only to be able to do a better job than a naive forecaster would do. At first blush, the benchmark for our task appears quite simple: like the weather, the best naive forecast would be for no change. This turns out not to be the case, however, because we are dealing with money, and in the world of investments, money is not left idle for a year, but is invested in instruments that will produce income.

Components of Expected (and Actual) Return for Currencies

There are two sources of return from owning a foreign currency. The first, and rather obvious, cause for a change in the value of one's holding is a change in the exchange rate over the holding period. The second involves the alternative investment results that obtain over the investment period in the countries in question.

It's easy to think about these two sources of return if we are talking about cash investments, because one can imagine exchanging dollars for another currency on the spot market (i.e. at today's exchange rate), investing that (foreign) cash in interest-bearing bank deposits or short-term government securities denominated in that currency, then exchanging the foreign cash for dollars when the holding period is over. In that scenario, the number of dollars obtained at the end of the holding period will be a function of both the then-prevailing exchange rate and the amount of (foreign) interest accrued during the investment period. When we translate our foreign holding back into dollars, however, we should not compare how much money we have with how much we started with, but rather with how much we would have had if we had kept our dollars invested in short-term dollar investments of similar risk. Thus, it is conceivable that we could make money, in an absolute sense, from both appreciation in the foreign currency and from the interest on our foreign investments, and still lose money relative to what we might have earned in dollar deposits if domestic interest rates were high enough. In that case, we would not claim to have added value through our cleverness in investing abroad, but instead would have to admit we had a negative return from our currency management.

Similarly, if we are dealing with other forms of investment, (such as bonds, stocks or even direct investment), it is also important to take into account the interest rate environment prevailing in the two countries over the relevant investment horizon. Any hedging to be considered (assuming it is opportunistic and therefore of a limited time horizon) will of necessity be done with short-term instruments (else it is not hedging, which requires isolation of the currency risk so that it can be neutralized).

In the case we address directly in this chapter, of hedging the currency exposure in a global equity portfolio, interest rate considerations are very much in evidence. Toward the end of the chapter, we will explain exactly how to take into account interest rate differentials when deciding how much to hedge, but for now it is important for the reader to understand that hedging is usually done in the forward market, and that forward contracts are explicitly priced to account for interest rate differences, as explained in the next section. For most institutional investors, the most cost-effective way to hedge foreign currency exposure is to sell that currency forward, on the presumption that it will decline in value (relative to the dollar), so that when the forward contract is settled, settlement can be made with units of the foreign currency that can be bought then at a lower price than is now available. Thus, if interest rates are properly taken into account, the investor will end up with more dollars than would have been the case if the hedging had not been done.

Covered Interest Arbitrage

The principle of covered interest arbitrage is

that currencies will be exchanged for each other in a forward contract at a ratio that reflects the difference in interest rates between the two countries.[2] Specifically, the premium (or discount) that must be paid to buy a foreign currency in a forward contract is

$$f = \frac{F}{S} - 1$$

where *f* designates the forward premium (or, if negative, discount), which is simply a function of the current spot rate for the foreign currency (S) and the forward rate for that currency (F). The forward rate itself is determined, under arbitrage arguments, by the difference in (default-free) interest rates between the foreign country (i_f) and the domestic market (i_d):

$$F = S \frac{(1 + i_d t)}{(1 + i_f t)}$$

where *t* is the length of time in the forward contract (as a fraction of a year, assuming interest rates are quoted as simple rates).

According to this description, it can be seen that the currency of the country with the higher interest rate can be purchased forward at a discount. The amount of the discount will exactly reflect the difference between interest rates in the two countries for the time period of the contract. If the discount is less than the interest rate differential, it would be possible to make riskless arbitrage profits by selling forward the higher rate currency and, with money borrowed at the lower interest rate, buying an appropriate amount (as determined by the discount) of that high-rate currency in the spot market, investing the proceeds in interest-bearing securities denominated in that currency. When the time comes to settle the forward contract, regardless of the then-prevailing exchange rate, there would have accumulated more than enough of the high-rate currency to settle the forward contract and (after exchanging back into the low-rate currency) repay the loan, and a riskless arbitrage profit would have been earned.

The arbitrage can be done in reverse if the discount is too great. In that case, the high-rate currency is bought forward, and interest-bearing instruments of the same maturity are sold short in that market (i.e. denominated in the high-rate currency) and the proceeds are invested in an interest-bearing note in the low-rate currency. When the position is unwound, there will be a residual profit that was earned at no risk to the arbitrageur.

In practice, the FX markets for the major currencies are so liquid that this principle of covered interest arbitrage does, in fact, describe the pricing of forward contracts very well. For developing or smaller markets, the expenses associated with even this simple arbitrage can be significant enough so that the forward premium may seem to violate the arbitrage argument until those costs are taken into account.

Doing better than a naive forecast

Let us reduce some of these abstractions to a numerical example. Our task is to predict the level of a given currency exchange rate at some point in the future (one year ahead, to use our standard time frame). As we know from the principle of covered interest arbitrage, a naive forecast would depend on the one-year interest rates prevailing in the two countries.

If, for example, a dollar can be exchanged for 110 yen at the present moment, how many yen will it take to buy back that dollar in a year's time? Suppose that at the beginning of the year, we had been able to purchase (default-free) one-year government securities that offered a yield of 1% in Japan and similar securities offered 5% in the U.S. Under those conditions, we would have to anticipate a 4% appreciation in the value of the yen (to 106 ¥/$) in order to remain indifferent between the two currencies. With a rise of 4% in the yen plus 1% interest on our yen holding, we would end up with $1.05 for each dollar invested in yen. If, on the other hand, we anticipated less of a rise in the yen, we could have simply bought US government securities and ended up with $1.05. Or, if we knew that the yen would rise more than 4%, we would buy yen with our dollar and earn the 1% available from Japanese government bonds plus that appreciation in the yen, to end up with more than $1.05.

The principle of covered interest arbitrage tells us that we should be able to enter into a one-year forward contract to buy yen for about 106 per dollar (reflecting the 4% interest-rate differential between the two countries). In other words, a naive forecast of the ¥/$ rate a year hence would be 106. Our task, as forecasters, is to do better than that, not better than predicting no change. Our goal, as money managers, is to use our currency forecast to make more than the 5% that is available in U.S. T-bills. If we correctly forecast a rise of more than 4% for the yen, we can add value by investing in yen- rather than dollar-denominated bonds. If we

own yen-denominated equities (or other securities) and correctly forecast that the yen will gain less than 4%, we can benefit by hedging some or all of our yen exposure back into dollars by selling yen forward. Note carefully that we might want to hedge a currency position even if we expect that currency to appreciate relative to the dollar, provided that interest rates in that country are sufficiently below those available in the US.

In practice, the opportunity set for capitalizing on prowess in the FX markets is usually limited to hedging currency exposure created by owning foreign securities. Most money managers are not permitted to establish short positions in currencies or to cross-hedge (use proxy currencies that provide more liquidity -- e.g. selling forward the deutsche mark to hedge a guilder position when no underlying mark-denominated assets are held). Thus, the only decision is to hedge or not to hedge (or to be a little more sophisticated about it: how much to hedge). Obviously, being limited by this one-sided approach lessens the value-added that could be achieved if one had the ability to fully capture good forecasts. Most of our comments focus on hedging, but keep in mind that there are many other ways to profit from making correct (on average) currency forecasts.

One natural question that arises is: how good is the naive forecasting guideline we have just mentioned? Is it truly a bogey that is hard to beat? As it turns out, forward rates have, at least in recent years, have not been good (i.e. unbiased) predictors of future spot rates.

Forward Rate Bias

The principle of covered interest arbitrage (also called covered interest parity) ensures that forward currency contracts to exchange one currency for another are priced solely to reflect interest rate differentials between the two countries. While it may be possible for currency arbitrageurs to make profits when this relationship is momentarily violated, for all practical purposes it is impossible for money managers to achieve incremental returns by simply transferring funds from one country to another to take advantage of higher interest rates, provided that the currency risk is hedged away with a forward contract.

As outlined by example in the previous section, a naive forecast of FX movement is that currency exchange rates will change by an amount exactly enough so that the net return (i.e. including accrued interest) will be the same in all countries. If this condition were always to hold, then there would be no need to hedge currency risk, but history tells us this is a far cry from describing the real world. If, however, as many have argued, this proposition (called the theory of uncovered interest parity) holds, on average, over longer periods of time, this might persuade some investors that there is no need to forecast or hedge currencies in the context of managing an international portfolio of assets. In fact, the theory has fared very poorly in the real world over the past 20 years. Clarke and Kritzman[3] present empirical evidence that for seven major currencies (relative to the dollar) the forward rate failed, in a consistently biased manner, to predict the future spot rate. They call this the "forward rate bias", and observe that "when the forward rate was priced at a discount to the spot rate (foreign interest rates exceeded domestic interest rates), the spot rate did not depreciate to the level predicted by the forward rate... In contrast, during the months when the forward rate was priced at a premium to the spot rate, the spot rate failed to rise to the level predicted by the forward rate..." They conclude that "a hedger's best strategy is to hedge more of a low-interest-rate currency and less of a high-interest-rate currency exposure."

Why might this be so (that the forward rate underpredicts the appreciation that accrues to a high-interest-rate currency)? While we are not aware of any studies that have found a convincing explanation for this bias, it would be reasonable to assume that there is some underlying macroeconomic connection between high interest rates and the stimulation of forces that ultimately (even if indirectly) lead to a strengthening of a currency. Else how could this phenomenon have persisted for the past 20 years? Its remarkable consistency gives us some comfort that it will continue, perhaps not each and every year, but at least more often than not.

If this is true, adjustments such as Clarke and Kritzman suggest (by either lowering one's estimate of the return expected from a low interest rate currency, or simply increasing the amount of hedging done against such a currency) will improve portfolio results.

Nonrandom Behavior of the Foreign Exchange Markets

Several studies[4] have shown that FX markets violate the cherished assumption of many academics that prices change in random fashion. It is part of the folklore of the markets now that currency values change in discernible trends. If this is true, then profits can be made by implementing trading rules that require no forecasting ability.

Can these trends be used in forecasting?

Because of the regular appearance of trends in the FX markets, many traders have been successful at implementing trading rules that depend on the existence of these trends. We believe it is dangerous to use such systems in isolation, but when combined with the insights from a PPP-based valuation, the combination can be a powerful trading tool.

Is there an economic rationale for this seemingly "inefficient" behavior?

Many observers advance the explanation that there are players in the FX markets who are not profit-motivated, and so the markets do not "work" as traditional finance theory calls for them to. This, to many, is particularly surprising in light of the fact that the FX market is the largest market in the world when measured by the value of trades conducted during a given day, and spreads are razor-thin in comparison with other markets — symptoms of presumed efficiency. It is thought that there is enough opportunity for low-cost arbitrage that currencies should not stray too far from their "true" value, but this presumption is clearly violated in the real world.

In our view, the existence of central bank intervention is more of a psychological influence on the markets than an economic one; though, granted, its effect can be quite real in moving the markets. Many other participants in the FX market may be simply price-takers, such as corporations who hedge their contractual currency exposures or routinely translate their foreign sales into their home currency, and investors (in either physical or financial assets) who, without consideration of currency valuation per se, trade in one currency for another because they see better profit opportunities in the assets of another country. Nonetheless, these participants, over time, will slowly move the market price of a currency to a new equilibrium, as they become part of the structural forces responding to unequal purchasing power in different countries. The trouble, from the viewpoint of those who desire instantaneous adjustment, is that these very forces are slow-acting, because they require such things as building factories to take advantage of low wages, and other long-range activities. But, once set in motion, these same instruments of change will exert inexorable forces on the currency markets, as goods and services change hands in new patterns that respond to the slow, macroeconomic adjustments being made.

Thus we have, in our view, a case of turning around the proverbial ocean-liner; it takes time for the system to respond, but once it does, powers are set in motion that generate real trends in the way currencies are traded. So, a currency will move toward its fair value — but, lo and behold, these selfsame forces will cause the currency to move beyond its fair value, since they are not turned off any more easily than they were initiated in the first place. As a result, the currency will move far beyond its fair value before countervailing forces begin to reverse the process. All of this describes our view of the markets, and it is meant only as an anecdotal explanation for why currencies seem to trade in 2 to 4 year cycles. Whether this or any other explanation is wholly adequate may be irrelevant as long as the markets continue to behave in the fashion we have observed. Our system of forecasting does not depend so much on understanding what causes currencies to fluctuate around their fair value as it does on being able to measure that fair value and hence market deviations from that value.

Structural Models to explain foreign exchange movements

There are three[5] broad classes of economic theories that address exchange rate determination. While there is much overlap among the various theories, and it is often a matter of emphasis rather than an exclusive reliance on a single mechanism, it is still fair to say that theories can be divided into (1) those that view currency exchange rates as the relative price of money, which is known as the monetary approach, (2) those that view exchange rates as the relative price of goods, or the purchasing power parity approach, and (3) those that rely on a combination of trade, monetary, and financial flows to explain exchange rate behavior, which is termed the portfolio balance approach.

While in many ways, the latter (portfolio balance) approach might seem preferable, since it takes into account the broadest spectrum of economic activity, for our purposes it is not the easiest to use. Among other problems, this theory relies on risk preferences, which are highly conditional, potentially volatile, and subjective. It is equally hard to observe supply and demand schedules, but one must try to predict how exchange rates would respond to movements along those schedules, or to changes in these functions. While this may be all very well and good when the concern is understanding the way economies

operate and in devising government policies designed to optimize the economic well-being of a nation, these theories are, in our opinion, of limited practical value in forecasting exchange rates in the context of portfolio management. Much of the data needed to evaluate these theories is either not available (risk preferences and the like can only be inferred from behavior) or are collected with too much of a lag time to be of much value. This is not to dismiss these theories as being of no value to the international portfolio manager; quite the contrary, we suggest that being aware of their general outlines is quite beneficial to being able to evaluate whether the PPP-derived forecasts make sense, and what international events might trigger or encourage movement in the FX markets.

Similar comments can be made about monetary theories. The monetary approach attempts to explain currency exchange rates as a function of money stocks and flows. In a sense, PPP is not too much different, since presumably monetary policies have a strong bearing on inflation rates, but we believe it best to use something (like prices) that can be measured directly rather than something (like money supply) that more indirectly affects the FX markets. Also, it is not particularly comforting to observe the lack of universal agreement among economists as to what response FX rates should have to various combinations of monetary and fiscal policy.

Availability, Reliability, and Timeliness of Structural Data

Beyond the fact that there is no agreement in academia as to what impact fiscal and monetary policies are likely to have on FX rates, many academic empirical studies that address this issue use data that would not have been available in advance. The goal of these studies is generally quite different from that of the investment manager, in that the economists are often in search of policy prescriptions that could guide governmental actions. Unfortunately, though, this means that even if we became convinced of the validity of one of these models, we might not be able to use it in a forecasting role. Also, a common flaw of most academic studies is that they do not control for valuation. Instead, they generally (if only implicitly) assume that currencies are "efficiently" priced (i.e. fairly valued), whereas our experience warns us that results of many such studies would be critically dependent on the place in the valuation cycle that a currency is in when the study begins. A certain amount of exchange rate movement would have come about with no intervention or change in fiscal or monetary policy, so this should be factored out when looking at the result of such actions. No wonder nearly identical studies often come to opposite conclusions!

Our task is to identify the course of these valuation cycles, with the faith that structural changes and governmental policy actions will, in fact, not operate so quickly as to remove the profit opportunities that we identify. It is to that job that we now turn our attention.

Implementing a PPP-based Forecasting System

For purposes of the simulation and illustrations presented in this chapter, we have employed the simplest and most straightforward versions of the many possible data series that are available or can be constructed. We always use the CPI, for example, to calculate inflation, and we use a 60-month historical window (where appropriate). We chose to take this simplified approach to demonstrate that even with these parameters, our PPP-based system is robust enough to provide good results in forecasting currency movements.

Readers are invited to experiment with alternative measures (such as using export prices instead of consumer prices), or by varying the time window used as an historical reference base, but the methodology we describe will work very well, in our experience, in all markets, by employing only the most basic data. Naturally, improvements are possible by studying the unique characteristics of each country and market, but these incremental benefits are, in our perception, marginal in nature compared with the overall benefit of using the basic approach we advocate. Using export prices or fiddling with the time window will improve results in some countries. We have found, for example, that there is generally a correspondence between the length of the window and the volatility of the currency — more volatile markets need more observations to collect a representative statistical sample.

We do not mean to suggest that these refinements are a waste of time; quite the contrary, we are constantly searching for ways to improve our calculations. The incremental value added will come in basis points, however, not in percentage points. So these subtle modifications are only that — improvements at the margin; or, in other words, the kind of tinkering that quants love to do with their models, but that add just a few basis points per year to returns.

Our methodology is based on our belief that Purchasing Power Parity (PPP) theory describes

market behavior over long periods of time; that is to say: currency exchange rates will ultimately reflect relative changes in price levels between countries. Having said that, however, we admit to begging the questions: "how long?", and "which prices?"

The length of a currency cycle

As to the first of these issues: we are convinced, after several years of real-time experience with this methodology, and having done additional backtesting covering more than 15 years, that currencies do tend to gravitate toward their "true" PPP value (as we define it in this chapter). We must admit, however, that this observation is true only on average. Whether we look at monthly, quarterly, or annual periods, only about half the time does a currency move in the direction predicted by our PPP calculation. Fortunately, the moves toward fair value are, on average, larger than the moves away from it, so using the fair PPP value as a predictive measure is profitable from an investment standpoint. Be warned, however, that this is not an approach that will produce positive value added each and every quarter — we will describe, at the end of the chapter, some suggestions for reducing the inherent volatility of these results in a practical investment context, but the underlying process is a choppy one. The reasons for this seeming inefficiency in a market reputed to be the world's most liquid is a function of the mismatch of the speed of adjustment in the financial and physical markets. Real-world macroeconomic forces are slow and lumbering from the point of view of the fast-paced electronic investment world in which prices and information are broadcast about the globe instantaneously. These macroeconomic forces, nonetheless, create an effective negative feedback system that works to correct the excesses of both real and financial markets. The reader might ask, in response to this admission of erratic predictive ability, "Why not use a more stable approach, then?" We wish we knew a better alternative, but the PPP framework is the only system we have found that will, over time, consistently perform well in a real-time, investment-oriented environment.

Anecdotally, currency cycles (that is, the complete trip from fair value to maximum undervaluation, back to maximum overvaluation, returning to fair value) typically take anywhere from 2 to 4 years. Since about half of this time, a currency is moving in the "wrong" direction, one should not expect to consistently achieve profitable results over any period shorter than 4 years. Naturally, there will be many subperiods during that time when results will be favorable (but only slightly more than half of them, in our experience, as we mentioned above), and specific results will depend both on the length of a particular cycle, and at what point in the cycle an investment program begins. Given all of this, it should be clear that it would be dangerous to set a client's expectations too high. Nonetheless, for those willing to wait out the gyrations of the macroeconomic forces that work their magic (albeit slowly), the prospect of adding 3% per year to portfolio returns should provide ample incentive to stick with the process.

The best yardstick for inflation

On the second question ("which prices?"), our experience has been that any reasonable proxy of price levels within a country will serve well. In most cases, we use the consumer price index (CPI) [in the UK, consumer prices are measured by what is known as the Retail Price Index, or RPI], partly because it is, in most countries, the most readily available, the most carefully constructed and, usually, the most frequently updated. Then, too, the CPI is followed closely by many market participants, including central banks, and so may very well influence market prices directly, not just through its macroeconomic links. Also, the CPI closely tracks any other indicator of inflation over the time periods that are relevant to our analysis.

Presumably, though, currencies are exchanged for each other to buy things that can be physically transferred from country to country, such as automobiles, computers, and foodstuffs. Yet the CPI measures the price level of many things which cannot be transported, such as housing costs. Might it not be better to use a measure of the relative cost in each country of manufacturing export goods (e.g. automobiles), or perhaps use some other more specialized index of export prices? It is changes in the prices of *tradeable* goods that should influence changes in foreign exchange rates. After all, if a non-exportable commodity (a residence, for example) or service (such as a restaurant dinner) rises (or falls) in price, this would hardly induce someone in a foreign country to make a transaction in the FX market. On the other hand, if computer parts or agricultural goods or some other highly fungible product can be made more cheaply because of manufacturing efficiencies (e.g. keeping a tight cap on wages) or by technological innovation, then a country might expect to see its currency appreciate in response to demand for its products.

Ultimately, which measure of inflation is the

proper one to use is an empirical issue, but in most countries, we have yet to find an index that works better than the CPI. Also, keep in mind that it is relative prices between two countries that is of interest in valuing their currencies relative to each other. Although export prices may not always track closely with consumer prices when viewed from the context of one country, when looked at on an international scale, since export goods compete with other export goods (for budgets, if not directly) then it is unlikely that one country will be able to raise prices for its export goods faster than do other countries for extended periods of time. Then too, our experience is that currencies often get misvalued as much as 20% from their PPP value (as calculated by us), and it is unlikely that such large discrepancies would be significantly different even if we used an alternate measure of inflation.

In certain countries, it is true, using an index of export prices will improve results.[6] Japan, for example, has, unlike many other industrialized countries, over the past few decades, been a country dedicated to production rather than consumption. Much of the engineering and manufacturing talent of the nation has been devoted to improving sales abroad, and that emphasis has seen the price of exported goods decline in yen terms. But this is not the whole story. In almost every country, export prices have risen more slowly than consumer prices over the past few decades. The reason for this is intuitive: in a competitive world economy, only countries that succeed in keeping their production costs from rising will be able to find markets for their exports — technological innovation can create new markets or give a country a temporary edge, but in the long run, imitation and substitution will reward low-cost producers. Nonetheless, Japan has done a better job than most other industrialized countries in lowering the prices of its exports. Over the past 15 years, Japan's CPI has risen more slowly than in the US: the gap is more than 2% per year (2.1% in Japan versus 4.5% in the US). Over the same time period, Japan's export prices have actually fallen by more than 2% per year, while those of the US have risen by a like amount. This means that Japan has enjoyed nearly a 5% per year advantage over the US in terms of its export prices, which is remarkably close to the annualized appreciation of the yen relative to the dollar over that time span.

Trade flows, portfolio flows, political, fiscal and monetary actions, and other considerations may move market values away from their PPP levels for periods of time (often for many months), and our challenge is to distinguish between temporary dislocations in the exchange rate (which present profit opportunities) and permanent changes brought about by structural changes in the global economy. A country running a trade surplus with another country, for example, may be perfectly content to invest its accumulated foreign currency in the capital markets of that other country, in which case the trade imbalance will have little or no effect on the currency exchange rate. Should, however, investors in the surplus country decide to repatriate their money or invest it in a different country, the FX transactions resulting from this shift in the demand schedule could have a large impact on the market value of the currencies involved.

Identifying and predicting the macroeconomic forces that determine the timing of exchange rate movements is difficult at best. Our technique of projecting PPP value and basing our forecasts on the idea that the market value will eventually be driven toward that level works well, and our research has failed to uncover any other method that is as successful at predicting returns in foreign exchange markets, or that can even be used to supplement this process.

Measuring Inflation

Chart 2 shows the pattern of consumer price inflation in France and the USA over a 15 ½ year period beginning in 1981. Notice that in the first half of this period, inflation in France was typically higher than in the US. After 1987, however, this pattern reversed itself and France experienced a rate of inflation that averaged about 1% less than that in the US.

[A technical note: consumer price inflation is normally reported by governments and therefore the press as the simple year-over-year percentage change in the consumer price index constructed by the government. We show here a slight

variation: we have computed the natural log of the ratio of the index to itself a year earlier. This computation produces a continuous rate of change, which in the lower end of the scale is nearly indistinguishable from the simple rate of change (e.g. a 3% simple rate is almost identical to its continuous version of 2.96%), but has the effect of compressing the scale in the early 1980s, (e.g. a 13% simple rate equals a 12.22% continuous rate). There are computational and statistical reasons, as well as visual ones, for preferring continuous rates, but either may be used provided one is careful about how they are used in calculations and in forecasting algorithms.]

Chart 3 shows the pattern of changes in export prices for the same period for the same two countries. Notice that, in both France and the US, the series is more volatile, and inflation was at times negative. The net compound result over this time period was that export prices in both countries rose less than consumer prices (about 70% in France and 45% in the US, compared with 120% and 100%, respectively). It might seem that there is a lot of information contained in these price series that is relevant to the franc/dollar exchange rate and that is not captured in the CPI.

To steal a look ahead, however, *Chart 4* depicts

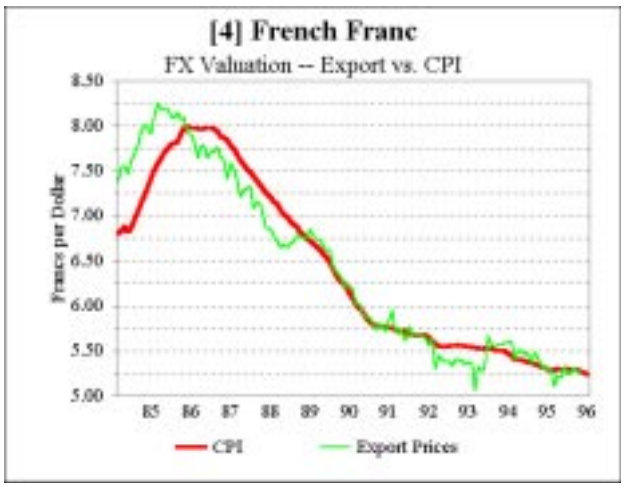

alternate valuations for the FX rate we would forecast with each of these price series. As might be expected, the export price series produces the more volatile forecast, but that does not make it better. In fact, our finding in France is that it provides a worse prediction of actual FX movement. Perhaps what is most notable is that the results are quite close to each other.

We will now look at how we derive these forecasts.

Purchasing Power Index

Our first step is to compute an index of relative inflation, based on the growth in the Consumer Price Index (CPI) in each of the two countries of interest. We can use any currency as the base (or numeraire) currency, but the examples given on these pages take the point of view of a dollar-based investor. For the sake of consistency, we will follow through and give all examples using the French franc, though the patterns seen here are very similar for most major currencies, especially the European ones.

Chart 5 relates the inflation experience in France to that of the US in what we have labeled a "Relative Purchasing Power Index" (RPPI).

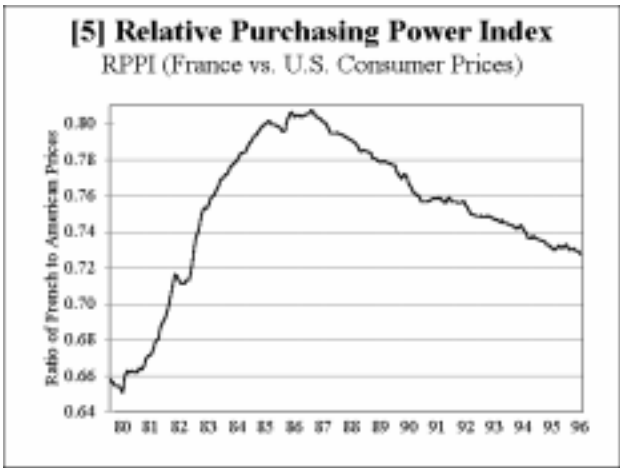

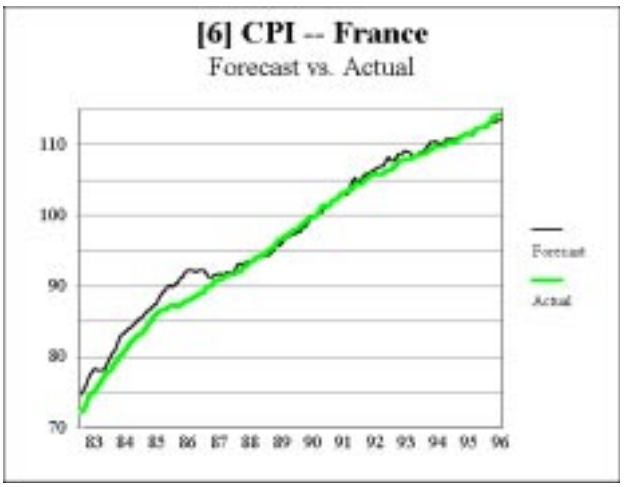

This index is simply the ratio of the CPIs in the two countries of interest. Note that, for France, this relative purchasing power index was rising in the early 1980s, reflecting the phenomenon shown in *Chart 2*: that prices in France were rising faster than in the US. Once France began to have lower inflation than the US, this RPPI began to fall. The scale on the left (Y) axis, by the way, has no particular meaning. As will become clear as we step through the analysis, it is only of interest whether this series is rising or falling, and by how much, in percentage terms. The level of each index as published by the governments of the two countries, and therefore the ratio of the two price indexes is completely arbitrary. Regardless of the scale that results from this calculation, the RPPI represents the pattern of relative price movements in the two countries, and that is the crucial piece of information we need.

Forecasting Inflation

Our second step is to forecast inflation in the U.S. and in the country of comparison. As a reminder: we have chosen, for illustration here, to forecast one year ahead. *Chart 6* shows the result of our methodology as applied to France:

In forecasting inflation, we are, in part, seeking a proxy for investor's inflation expectations. It is important to forecast what inflation will actually be, since that reality will ultimately influence the value of the currency exchange rate. It is also important, however, to understand what people's inflation expectations are, since these expectations will become impounded in exchange rates (along with a host of other relevant factors), thus helping to determine their near-term path.

For these reasons, we combine a mixture of an econometric forecast of inflation in each country with a measure of trailing inflation. We believe this provides the best blend of expectations -- anecdotally, people are heavily influenced in their inflation outlook by recent experience, but markets are also smart enough to take into account changes that will modify the future course of inflation. We use the average inflation experience of the past two years as a proxy for average inflation expectations.

We are aware that other techniques could be employed here that might provide a better inflation forecast. We believe, nonetheless, that errors introduced by any reasonable approach will be minor compared with deviations in market exchange rates from true economic value. In other words, even having a perfect inflation forecast would not go very far in explaining why exchange rates vary so much from their PPP values, as we calculate them. As mentioned earlier, our experience is that currency exchange rates often get as much as 20% or more away from the value our model ascribes to them. Errors in estimating inflation expectations *differentials* are a small fraction of that. *Chart 7* expresses our forecast for the CPI as an inflation rate. Except for the mid-80s, when the inflation was rapidly deteriorating, the error in our forecast has been about ½% per

year in France.

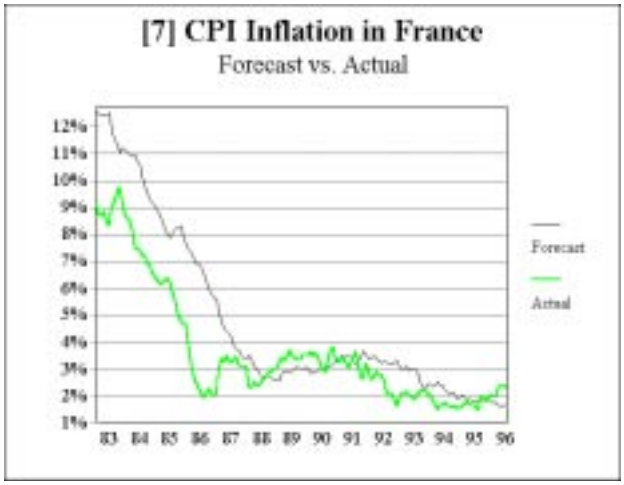

Completing the circuit, we substitute our inflation forecast for actual experience in computing a forecast for relative inflation (the RPPI). *Chart 8* gives the picture for France.

Note that the errors in the forecast are a combination of errors in forecasting inflation in the

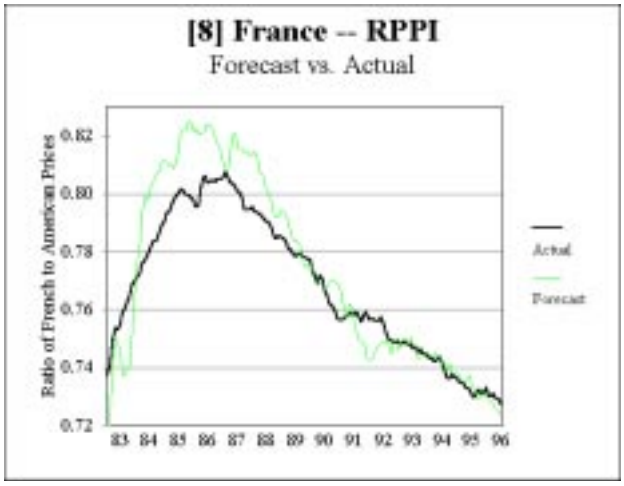

US as well as in France, since this chart is a comparison of the two. Of course, in some periods, these errors will offset each other, but at other times the error will be compounded. On balance, the forecast is not too much different from the actual outcome, at least in the context of actual exchange rate fluctuations, as mentioned in the previous paragraph.

Given all our caveats about how difficult it is to fathom inflation expectations and to forecast inflation accurately, combined with our disclaimer that the errors are relatively insignificant, the reader may very well wonder why we bother at all to go through the exercise of forecasting inflation. Keep in mind that in projecting exchange rates in the context of the PPP paradigm (the goal of this chapter), one must (if only implicitly) forecast inflation. If no conscious projection is made, the unstated assumption is that inflation differentials between the two countries involved will remain unchanged. In our experience, making a stab at an inflation forecast improves results, if only at the margin, so we do it; but we would not find serious fault with an approach that skipped this step.

Pricing the Index: Purchasing Power Price

The third step in our process is to compute and then forecast the price of the Relative Purchasing Power Index in units of the foreign currency.

This number represents the contemporaneous market price of the relative price index (RPPI) computed in the first step. If this price were always a “fair” economic value, we would expect it to remain constant. If that held true, changes in the PP Index (that is, relative inflation) would flow through directly to the FX rate. In reality, however, this price ratio will fluctuate, and we can never be sure of when it is in “equilibrium”. In truth, we don’t attach much significance to the concept of equilibrium as defined in its classical economic sense (determined by the interaction of supply and demand schedules). Although it may be true that, at any point in time, an FX rate is in fact the balancing of supply with demand, that observation only begs the important question of what causes these schedules to shift. Clearly, preferences change over time in response to developments in the macroeconomic environment. It is a hopelessly complex task, for our purposes here, to try to determine what influences come to bear on these supply and demand schedules, and to try to predict their precise shape and drift. We will be content to observe that, whatever their cause, changes in supply and demand seem to have the effect of a negative feedback system. When the price of a currency gets too high, forces come into play that drive that price back down. These agents of change take time to work, and once they are at work, are not easily turned off, so the price tends to drift beyond what would be its long-run equilibrium level. As the price becomes too low, countervailing forces come into play and reverse the decline, sending the currency on its final leg of a cycle. All of this may sound like magic, and in a sense it is wondrous how the world does work, but we are talking about the forces that are commonly recognized to come to bear on FX rates: portfolio adjustments (both in the financial markets and in direct investments), governmental monetary and

fiscal policy, central bank intervention, speculation and arbitrage in the currency markets, and hedging activities by owners of real and financial assets and contracts. Because of the complexity of this system, the FX market is an imperfect negative-feedback mechanism, and therefore tends to overshoot the mark. Because the adjustments required are slow-acting and interrelated, the exact path and timing of the adjustment cannot be predicted, but the general pattern can by used to advantage if one thinks of the price of a currency as a mean-reverting mechanism.

See *Chart 9* for an illustration of this relationship in France.

Note that in the early 1980s, the franc clearly suffered a precipitous fall relative to the dollar in terms of its purchasing power. In fact, this was not unique to the franc; most other currencies showed a similar pattern compared with the dollar. The dollar enjoyed a very long run, pushed up (it is generally agreed) by a combination of a restrictive monetary and expansive fiscal policy in the US. High real interest rates and an expansive domestic economy created an environment in which PPP did not hold very well for a few years. For the past ten years, however, the franc has traded in a fairly narrow channel — for the last 120 months shown on this chart, the mean of .135 was only slightly higher than for the entire time period shown (.128). In that most recent 10-year period, more than 53% of all observations fall within plus or minus 10% of the mean (if one assumes that the observations shown are representative of an underlying process that produces normally distributed values). For the entire period shown on the chart, fewer than 25% of observations fall within that same band, and for the first 10-year period, only 22% of observations are that close to the mean. Clearly the market has become more orderly in recent years, but this only highlights the need to be cautious about using too short a period to define "equilibrium": if only the mid-80s had been used, the mean value used as a standard would be about 10% below what seems in retrospect to have been a fair price.

Forecasting the price of the RPPI

Chart 10 shows two different attempts to forecast the RPPI price, representing two philosophically different views of the FX market. One forecast (the light line labelled "Regression") is based on a typical time-series econometric approach, which implicitly assumes that past patterns will repeat themselves and that trends, once in place, will persist. The other approach (the darker line identified as "Reversion") is based on a mean-reversion assumption, and will work best if the series does have a tendency to fluctuate around a constant level. As a reminder, we illustrate here the use of a moving 60-month window. In other words, the forecast is based on the premise that the actual price will revert to its mean of the last 60 monthly observations. A shorter window would produce a more volatile forecast series, but the pattern wouldn't be too much different. Most currencies complete a full valuation cycle in less time than that, and it is important to include at least one such full cycle. Over time, however, structural changes in the world's economies and markets make older information less and less relevant to today's conditions, so including data that is too old can impair the predictive ability of the analysis. In the end, it is a judgement call as to how long a window of history is relevant.

As it turns out, of the two lines, the reversion method does a better job of forecasting, in the sense defined earlier. Namely, a portfolio strategy based on that assumption has better returns than one using the econometric approach, although in many countries, the regression line has a better "fit" with actual outcomes.

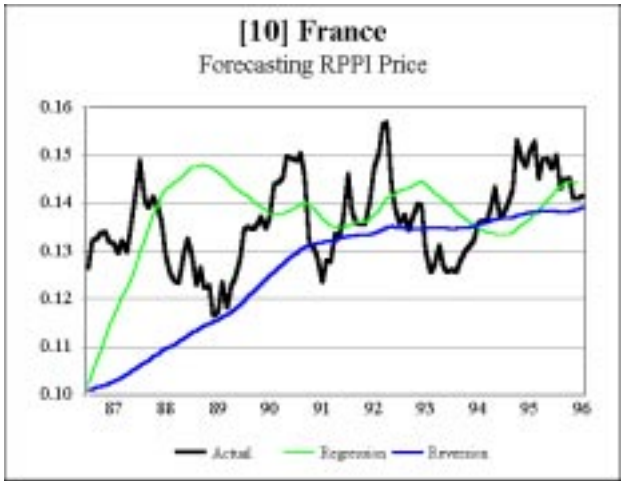

Projecting Exchange Rates

The fourth step in our process is to project a future (in this case, one year out) FX rate by assuming that the PP Price will return to some average level as observed in or projected from the time series computed in our second step.

Our forecast of inflation in the two countries over the desired horizon is used to forecast a future PP Index. A future PP Price can then be estimated as described in the previous section. As mentioned, we have often obtained better statistical results with projections based on econometric regressions, but almost always better financial results by using a reversion assumption. *Chart 11* shows what happens when an exchange rate forecast is made with the alternative forecasting methodologies described in the previous section.

As can be seen, the forecasting methods don't

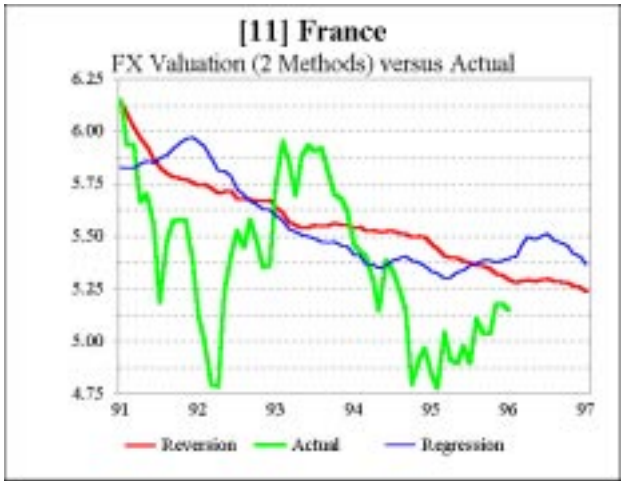

produce dramatically different results. Both forecast time series exhibit more stability than the actual FX rate, and both forecasting methods visually seem to be about equally right (or wrong). The discerning eyeball can perhaps detect that the reversion method is a bit better over this five-year period (ending June 30, 1996), and in fact it does have the better R^2 — .2615 versus .1037 for the regression approach. Beyond its better statistical fit, we find the reversion line more appealing because it changes more slowly, which is accordance with our view of how the world works, and because it more closely tracks the pattern of the RPPI (*cf. Chart 5*). In other words, the reversion forecast conforms more closely to what we would expect to see if we believe in the workings of the PPP paradigm.

Chart 12 focuses exclusively on the use of the reversion methodology, and here we show a longer (10 ½ year) period. To a statistician, this chart may not look like a very good forecasting model (and, in fact the R^2 of the entire period shown is only .3153 — actually a little better than over the past 5 years alone of .2615), but remember that our goal is not to forecast FX rates per se. Our job is to make a decision as to whether to hedge a currency or not, and if so how much, and our prize is whether those decisions prove to be profitable.

In this case, the financial success of this model is actually better over the most recent 5-year period than it was in the early years of the chart shown, so looking exclusively at R^2s or other measures of modeling veracity may give us misleading information. We can, of course draw some comfort from knowing that we have discovered useful information, but it is only if we can translate that information into concrete results (i.e. superior performance) that we should be congratulated. Let us complete, therefore, the description of our framework for using these forecasts in the context

of hedging a global equity portfolio.

Taking a Valuation Perspective

In making portfolio decisions, we find it easier to convert everything to the present tense. Somehow, we'd rather deal with spot rates and current interest rate differentials than to keep track of forward rates, although because of the principle of covered interest parity, we realize these are equivalent. We also find it helpful to think of our forecast value as a valuation measure rather than a prediction of a future market level. Again, the two concepts are equivalent, but by translating everything to the present, we find it simpler to think about what action we should take now. By simply shifting the time axis back one year for the forecast series from the previous chart, we can compare our valuation (where we think the FX "ought" to be) directly with the spot rate (where it is). *Chart 13* shows this transformation.

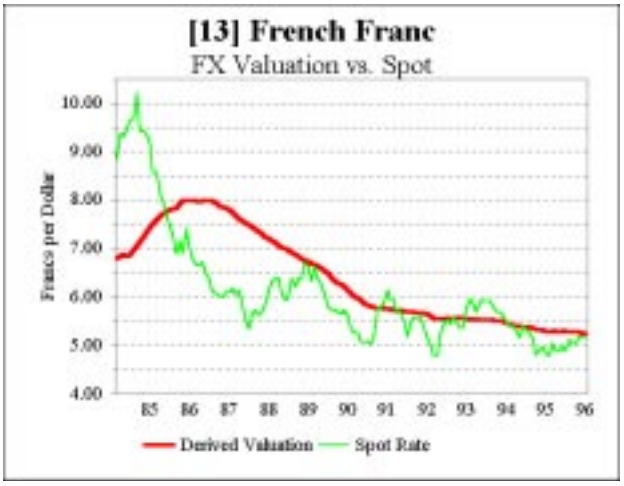

Computing an Expected Return

The final element in our forecasting process is to take into account interest rate differentials as well as the valuation just described. It could be that even if a currency is overvalued in accordance with our method of calculating its PPP level, it may not be a candidate for hedging. Remember that the cost of putting on a hedge is reflected in the forward rate that we will obtain when we hedge by selling some of our currency exposure off in the forward market. The price of the forward contract will, under the principle of covered interest parity, be determined by the difference in interest rates between the US and the country whose currency we want to hedge. It may be that if interest rates are low in the US (at least by comparison with the other country), we may have to sell that currency forward at a discount that is more than we anticipate it will depreciate over our investment hedging horizon. In that case, we should forego hedging. We therefore have to keep track of both these elements of expected return.

Chart 14 shows the components of expected return. Clearly, misvaluation (the difference in our valuation and the spot, expressed in percentage terms) represents the bulk of the potential for FX gains or losses. Nonetheless, there have been times, such as the period around the middle of 1992 when the French central bank was forced to aggressively raise interest rates to defend the franc, when interest rate differentials have been wide enough to change the sign of the expected return for the franc.

In calculating our expected returns, we combine the percentage misvaluation identified in the previous steps with the interest rate differential over the relevant time horizon. *Chart 15* shows the net result of combining these two sources of expected return.

Mapping Expected Returns into Hedging Ratios

The payoff from all of our previous work comes from interpreting our expected returns in a way that can aid portfolio construction. Managers who are using and are comfortable with an optimization framework can take the expected returns derived so far and integrate them directly into a country allocation analysis that has a full covariance matrix of both local asset returns and currency returns.

An alternative that, for us, produces superior results is presented here. Our experience is that this approach, which we term adaptive affine

mapping, creates lower turnover and less volatility than does a traditional optimization methodology while at the same time preserving (and in some cases even enhancing) the incremental returns. This is the process we used in simulating the portfolio presented at the beginning of the chapter.

Note that, for hedging purposes, we are only interested in the negative half of the distribution. When the expected return for a currency is below its historical mean (and here again, we used a 60-month moving window to define "history"), we want to hedge part of our exposure. The question is, "How much?" and the answer is determined by looking at a statistical interpretation of how significant is the current deviation from the mean.

The essence of the approach is a simple one — we transform the expected return directly into a hedge ratio. For each currency, we translate the current expected return into a number between -1 and +1, based on the statistical history of the expected return time series in that country. We recommend using a sigmoidal function (a process that will be familiar to users of neural networks), although a linear transformation could be used to good advantage as well. In our simulation, we used a normal probability distribution to assess the likelihood of any given expected return being significant.

Chart 16 shows the results of the transformation we made in running the simulation depicted in *Chart 1.* Again, France is used as an illustration, but we did the same thing in each of the other countries as well. When the "desired exposure" line was negative in any country, we hedged the portion of that currency exposure indicated by the level of the transformed series.

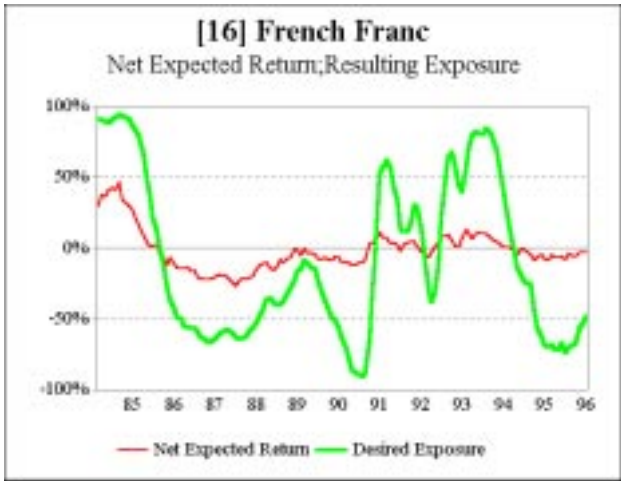

Additional methods to reduce volatility or further enhance returns

Two areas that have already been touched on are the forward rate bias and the nonrandom behavior of FX markets. While there are conflicting views as to whether these phenomena are transient or permanent features of the FX landscape, our own reasoning and the weight of historical evidence convinces us that these tendencies should be taken seriously — to the point that we think it makes sense to modify expected returns to reflect their existence.

Note that the simulated portfolio results reported in this chapter do not take advantage of these two observations, nor of our view that currency valuations tend to move in cycles. Instead of beginning to hedge the franc in early 1995, for example, when the expected return started to become negative, we could have waited, knowing that the normal tendency of the currency would be to become much more overvalued before turning around. By implementing decision rules such as these, the results of our simulation can be made less volatile and even more productive.

Another feature of real-life investing is that many portfolios will not be able to do as much hedging as would be required to replicate our results. Managers (or clients) may be uncomfortable hedging the majority of a currency exposure. In the case where there are limits, our strategy can simply be stopped when the limit is reached, or the numbers can be rescaled to fit within the constraints. If, for example, there was a restriction against hedging more than 50% of exposure to any one currency, the scale of the "desired exposure" series in *Chart 16* could simply

be changed to be half of what we show at every point on the curve.

Since there are an infinite number of such overlays of decision rules and constraints that could be used, we decided to present only the plain vanilla version. Volatility could be reduced with constraints, but at the expense of giving up a portion of the excess returns. Decision rules, such as delaying implementation, can improve returns while reducing volatility, but as such rules become less mechanistic and more judgmental, they are hard to simulate.

There are many other potential applications besides the one we illustrated here. If one can do cross-hedging, for example, or take long positions in currencies, the framework presented here could be used to compare the potential returns of two nondollar currencies, or to establish a portfolio of bonds or short-term instruments that has a higher expected return than a domestic-only portfolio of similar duration. For now, we will be content to have conveyed the message that it is possible to add value through the use of disciplined currency management in a global portfolio.

Suggested Areas for Future Research

Our experience in using the expected returns derived in our methodology suggests that they have a shelf life of several months. This is the result of already observed tendencies of the FX markets to exhibit nonrandom behavior; to wit, that monthly changes in FX rates are positively serially correlated, and that currency valuations move in cycles. There are many alternatives besides simple averaging to computing a reversion target that take advantage of this fact. For example, an exponential weighting scheme could be used to give more importance to recent experience.

When these phenomena are analyzed with the many other parameters that have been mentioned in this chapter (such as the length of the historical window to be used) it can readily be seen that the search space of all possible prediction methods becomes quite vast. It is not feasible to conduct research by trial and error, nor is it possible to find the best combination of parameters by optimizing one and then testing another. Because of the complex interaction of the assumptions, there is a high degree of danger of locating only a local maximum in the search space. For this reason, we are experimenting with two other approaches that have shown some promise in similar applications: genetic algorithms and simulated annealing. While we can't claim to have made any profound breakthroughs by using these techniques, we do believe they show promise, and they have helped sharpen our thinking on a number of issues.

In addition, because of the complex dynamics of the FX market, some research tools inspired by the study of chaos and nonlinear dynamics may be useful to our efforts. If, as we believe, the FX market is a nonlinear dynamical system that evolves through feedback mechanisms, the path of the FX rate itself or, more likely, of one or more of the intermediate forecasting calculations we make, may profitably be dealt with as a response to a strange attractor. We are currently investigating some recently developed time series forecasting techniques such as time-delay embedding and are pursuing the use of neural networks.

To date, our research suggests that all of these techniques are overkill. One problem may be that many of the underlying mathematical models were developed to understand how natural nonlinear systems operate. Unfortunately, in the world of human finance, the rules are always changing, and the past may not be prologue. The good news is that the very basic approach we have outlined in this chapter is hard to improve upon. Nonetheless, we are hopeful of gaining insights through alternate techniques such as those just mentioned that will prove useful in future versions of our analysis.

Michael F. Wilcox, CFA

President, Alford Associates, Inc.

Consulting and Advisory Services to the Professional Investment Community

———————————————————————

217 East Road • Alford, Massachusetts 01266 USA
Tel: 1-413-528-8098 • Fax: 1-413-528-8099
Email: Alford Inc@aol.com

Michael F. Wilcox, CFA, founder and CEO of Alford Associates, has more than 20 years of investment management experience. During that time, he has developed investment research and management solutions for major institutions in Hartford and New York, and most recently was an international equity portfolio manager and the head of Global Asset Allocation at a leading Boston investment house. He spent 10 years in the Wall Street environment, the last 5 as head of Quantitative Research in the Equity Research Department at Morgan Stanley. Wilcox holds a Master's degree in Economics from Trinity College in Hartford. He received the CFA designation in 1981, and became a registered securities agent in 1982.

Mr. Wilcox is an active member of several professional groups, and has served as an adviser to the Board of Directors of INQUIRE Europe and as a member of its program committee. He is an internationally recognized authority on quantitative investment methods and speaks frequently at international investment conferences. He is a specialist in the areas of asset allocation and currency management, and has broad experience in the valuation and trading of all types of securities and derivatives.

Mr. Wilcox is also the editor of the *Thrift Savings Investment Quarterly*, a publication of The Center for Retirement Research. This newsletter is oriented toward federal employees who are members of the federal government's Thrift Savings Plan, and offers advice on financial planning for retirement and on choosing a proper fund allocation.

It was intended that this page be left blank. Instead, this disclaimer appeared.