

Purchasing Power Parity: Even the Big Mac Can Predict FX Rates

By Dori Levanoni and Max Darnell

Many academics and practitioners have claimed that the principle models for measuring the fair value of a currency simply don't work. Purchasing Power Parity (PPP), in both its "strong form" and its "weak form" are said to fail, particularly over short-term horizons. Some published studies do support the notion that PPP works over very long-term horizons – as long as 5, 10 or even 50 year spans – but they do not tend, if ever, to support the notion that PPP is at work over horizons as short as a year or a month. Despite that, a measure as simple as the "Big Mac Index" published yearly by the Economist appears to work modestly well, and more sophisticated PPP constructions such as that from the World Bank and that belonging to First Quadrant appear to work quite well. We present the evidence and detail the most important issues.

The Big Mac Index

If you were to pay a dollar for a Big Mac in New York, and 100 yen for a Big Mac in Tokyo, at what rate would you expect to be able to exchange dollars for yen? If the answer was that 1 dollar buys 100 yen, then you could either buy a Big Mac in New York, or you could exchange your dollar for 100 yen and buy a Big Mac in Tokyo. The price of a Big Mac in New York would *effectively* be the same as the price in Tokyo. Formally, this is what is known as the *Law of One Price*.

Now Big Macs don't keep too well if transported across the sea, but other items such as bars of gold, cars, electronic equipment, rice and wheat do. If you found that the price of goods in Japan were lower than the price of goods in the US, then provided that either transportation costs were low enough, or the cost differential itself were large enough, then one could profit from buying goods in Japan and selling them in the US. This potential for arbitrage should ultimately insure that the prices of very similar goods don't get too far out of line across the ocean.

Arbitrageurs would not only drive the relative price of goods up in Japan as their demand for the good rises in Japan and falls on a relative basis in the US, but they will drive the price of yen up as well. In order to buy Japanese goods, they will have to sell dollars and buy yen. In other words, the demand for yen, just like the demand for Japanese goods will rise as well, driving the price of the yen up and the price of the dollar down. As this happens, the effective price advantage associated with the cheaper Japanese goods will narrow. Japanese goods will be rising in price, and for your dollar, you'll be receiving fewer and fewer yen with which to buy those increasingly expensive Japanese goods. In theory, this process should continue until the purchasing power of a dollar is equivalent to the purchasing power of the yen. In other words, simple supply and demand pressures will insure that Purchasing Power Parity holds approximately.

So what does the Big Mac Index do for us? Every year the Economist reports on the prices of Big Macs around the world. If, when translated through the exchange rate mechanism, the price of Big Macs appears to be lower in some countries than in others, then that is presumed to be an indication that the exchange rate is misaligned, i.e., that the currency is underpriced. The beauty of the index is that the Big Mac is very nearly the same good in every country in the world, so its price, according to the PPP theory, should be approximately the same everywhere.

Of course this is too simple, and with good humor behind its cause, the Economist knows that. If Big Macs are cheaper in Mexico City than they are in Los Angeles, Los Angelenos are not going to drive across the border to get lunch. But what if we were talking about automobiles in Tokyo compared with automobiles in Detroit? Then, yes, Americans would buy cars from the Japanese in part because of a favorable price advantage in doing so.

More importantly, while Big Macs (or automobiles) might appear to be underpriced in Tokyo, rice might appear to be overpriced there at the same time. What does this tell us about the exchange rate? Just looking at the relative pricing of one good might be quite misleading, and that's why serious efforts to evaluate exchange rates do so on the basis of some kind of "basket" of goods. While this helps by taking into account the "average" relative pricing of goods, it still suffers from some serious problems.

Even more importantly, there is a lack of symmetry in the relevance of particular goods. If rice appears to be expensive in Japan relative to the price of rice in California, and wheat appears to be cheaper in Japan than in the US, it is not clear how price pressure should affect the exchange rate. Rice is a far more important item in the Japanese consumption basket than it is in the US, so the price of rice would seemingly put more pressure on the yen/dollar relationship from the Japanese perspective, but just the opposite would be true from the US perspective. Wheat matters more to Americans, and from the US perspective, the mispricing of wheat would seem to put more pressure on the yen/dollar exchange rate. This lack of symmetry makes comparing goods prices very, very complicated.¹

For all of its naivete, the Big Mac index works reasonably well. As the figure shows, a portfolio² that bought currencies that the Big Mac Index said were cheap, and sold short those that the Big Mac Index

said were expensive succeeds in generating a profit. No, the performance of this portfolio was not spectacular, but it has worked, and that's the point. Apparently, even such a simple, naive approach to PPP works despite the conventional claim to the contrary!

(1) Here's a semi-serious way of looking at the problem of symmetry. For the calendar year 1998, McDonald's got 50% of its sales from the US. The US was 51.86% of the MSCI world capitalization. McDonald's got 11.5% of its sales from Japan. Japan was 9.82% of the MSCI world capitalization.

<i>Ratio of McDonald's US sales to JPN sales:</i>	4.35
<i>Ratio of US MSCI weight to JPN MSCI weight:</i>	5.28

Looks like the Japanese are as likely as Americans are to choose between Big Macs and Stocks, though Americans seem to prefer Stocks to Big Macs slightly. However, PPP is not about Stock Price Levels, but overall Price Levels. So, let us look at GDP numbers instead.

<i>Ratio of US GDP to JPN GDP:</i>	161
---	------------

That number is a bit more troubling. It appears the Japanese are nearly three times more likely to buy items besides Big Mac's than their American counterparts. The similarity between U.S. and Japanese Big Mac's is weaker than one might hope for. And this is only one comparison of the many possible.

(2) The countries used in the simulation were Australia, Canada, Denmark, France, Germany, Italy, Japan, the Netherlands, Spain, Sweden, the United Kingdom, and the United States. We use the PPP measure to build paper portfolios to see how well they have predicted FX returns over time. The positions are based simply on the degree of mispricing that the respective PPP measure indicates a given currency is trading at (i.e. if the Franc is 10% underpriced relative to the simple average mispricing of all currencies, then we establish a long franc position of 10%. Because all positions are established on the basis of an average-relative mispricing, the portfolio is always market-neutral, i.e. the sum of all bets is made to be 0%. As the signals are annual, we simulate trading once a year, in April (the Economist makes the Big Mac data available at the end of March).

Figure 1
Big Mac PPP Simulations (1989-1997)



The Strong Form of Purchasing Power Parity

The *Law of One Price* as described above is the simplest form of PPP. It states that the prices of similar goods should carry the same effective price regardless of the currency in which that price is stated. This is often referred to as the *Strong Form of PPP*.

$$\frac{FX_A}{FX_B} = \frac{P_A}{P_B}$$

The central difficulty associated with implementing a PPP model is the problem of identifying which goods should be subject to the law of one price, how to weight the relevance of those goods *within* each market, and how to deal with the relative weights of those goods

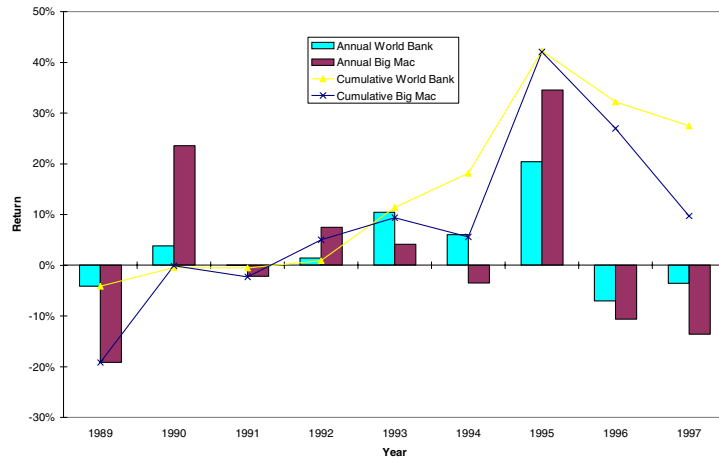
across markets. Complicating the issue further is that fact that the goods that should be included in the comparison will change over time. Goods that matter a great deal today, e.g., computers and various other electronic gadgets, would not have even existed when exchange rates were first allowed to float post-Bretton Woods.

So, not only is it difficult to identify the relevant consumption basket by which to evaluate whether the law of one price holds or not, but identifying the right weights to apply to the goods within the consumption basket over time and across markets is spectacularly complicated and uncertain. The World Bank throws enormous resources at solving just this problem. Something on the order of 65 people in the United Nations' ICP group are employed in the effort of identifying what goods should be in the baskets, how those baskets change through time, and gathering and analyzing the data. Like the Big Mac Index, the World Bank's approach is a Strong Form approach, but rather than the basket containing one good (a Big Mac), their basket has over 200 constituents, and great care is taken to weight them appropriately for each economy.

As the table below shows, the World Bank's PPP measure has outperformed the Big Mac Index substantially over the period 1989-1997 (both measures have different data lengths, so for comparison purposes we have run the portfolios only over what are the common dates). The World Bank's PPP's information ratio is nearly three times better than that of the more simple Big Mac Index.

	Big Mac	World Bank
Mean	2.30%	3.03%
Standard Deviation	17.51%	8.47%
Sharpe Ratio	0.13	0.36

**Figure 2
 PPP Simulations (1989-1997)**



Though, both still have down years ('89, '96 and '97), they demonstrate that with naive construction you can use Strong Form PPP measures to predict subsequent currency returns, and if care is taken in the construction of the basket, you can do quite well!

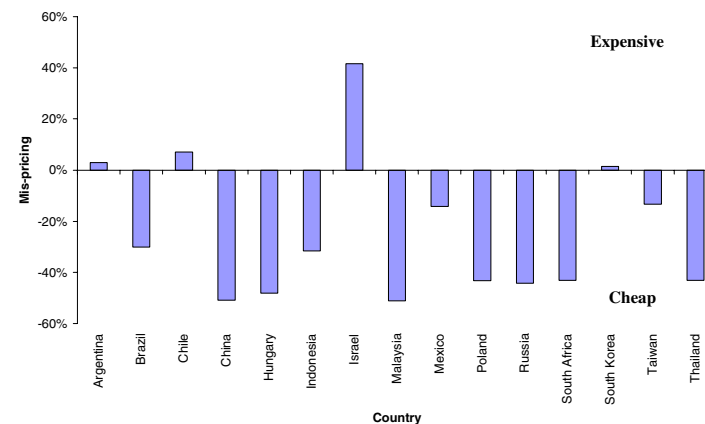
Is Fair Value Stable Over Time?

Given that few goods are strictly homogeneous, the “Parity” in Purchasing Power Parity may evolve over time. The Honda Civic of the late 1970’s more closely resembles Korea’s Kia in the 1990’s than it does the Honda Civic of the 1990’s. The quality associated with the Honda Civic has risen significantly, as has the quality of the vast majority of Japan’s exports. Changes in general prices don’t take this into account, so as the price of Japanese cars has risen through time

to reflect quality improvements we might be led to believe that the yen should be depreciating as a result of price inflation. Instead, the yen might not appreciate at all because, conceptually speaking, the quality adjusted price has not been inflated.

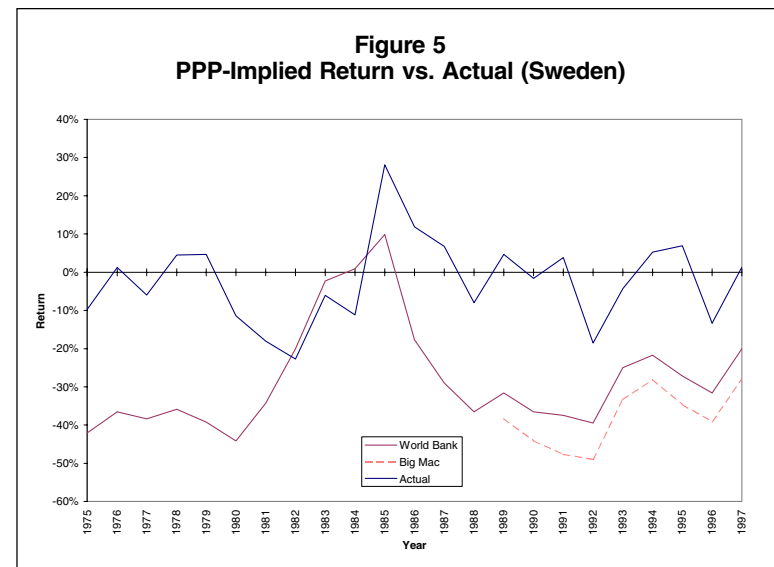
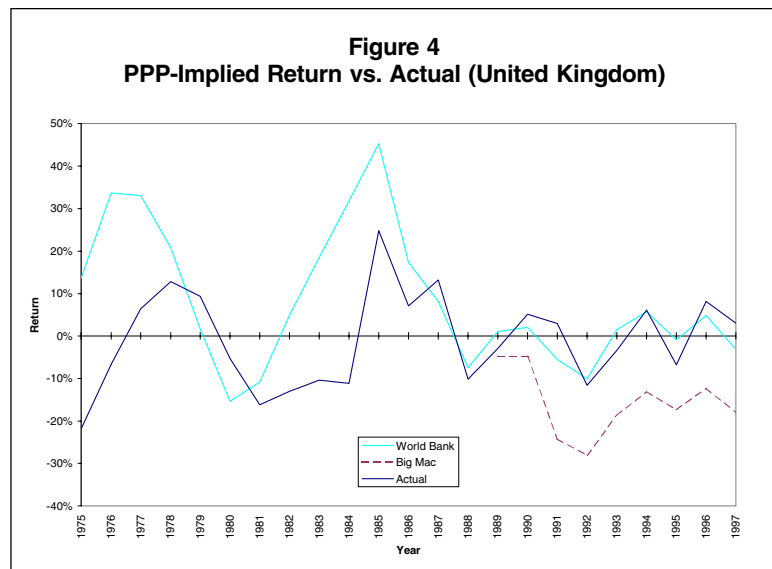
Markets where product quality is rising on a relative basis are therefore likely to be subject to the appearance of persistent mispricing in the exchange rate. There is also likely to be a relationship between productivity and biases in local prices. Where productivity and therefore wages and standards of living are low, some goods may benefit from cheap local labor prices, for example, and be priced lower than in markets where productivity is higher. It is interesting to note that in markets where quality and productivity are lowest, the Big Mac appears to be cheapest.

**Figure 3
 1999 Big Mac Index**



The implication is that changes in the prices of goods placed in the consumption basket may mislead us as to the value of the currency. The problem may reveal itself through a persistent “gap” between perceived fair value and the actual exchange rate. Persistent errors are, indeed, found in both the Big Mac and the World Bank measures of PPP-implied fair value.

The Big Mac Index shows an enormous bias against the Pound, which has been wrong in every year since the measure’s inception. The Big Mac is expensive in Britain as compared to the United States, and it has remained so for the past decade. In fact, the Big Mac shows this bias all throughout Europe, where Big Mac’s are consistently expensive (in 1999, the most expensive place to buy a Big Mac was Switzerland, followed by Denmark, Britain, Sweden, France and Germany). In this



case, the World Bank’s measure seems to avoid that problem, having agreed on average with the market exchange rate.

In the case of Sweden, however, even the World Bank’s measure displays a persistent error problem, though no worse than the Big Mac’s. In neither of these cases can we identify the underlying cause of the persistent error. In fact, it is not inconceivable that these persistent gaps between the PPP-implied fair value and the market exchange rate are actually errors. What we can do is ask whether under the assumption that these cases do represent errors, can we improve on the predictive power of our models. We can seek to adjust for what may be described as “evolving biases in the data” by adjusting for persistent and drifting means in the mispricing of exchange rates. Such an adjustment is one of the distinguishing characteristics of First Quadrant’s PPP model, so at this point we should describe our approach to PPP.

First Quadrant's Approach to Purchasing Power Parity

To understand how First Quadrant has approached the PPP model, we must first step back and review how theorists responded to the principle difficulties of the Strong Form of PPP. Recognizing that the strong form of PPP was rife with difficulties that are confronted when applied to real world data, economists proposed a second form of PPP that allows one to drop the Law of One Price model from the effort. The *Weak Form of PPP* essentially assumes that exchange rates are already in equilibrium and evaluates future *changes* in exchange rates according to *changes* in their purchasing power. Here we ask how well changes in the exchange rate over time reflect relative changes in the prices of goods.

$$\frac{FX_A}{FX_B} = k \frac{P_A}{P_B}$$

The intention of a PPP measure is to determine the extent and the direction of the pressure on a currency to either appreciate or depreciate. However, there is no prediction of the time-scale in which the actual appreciation/depreciation will occur, and worse, this construction assumes that past exchange rates must have been approximately fairly priced. If you want to predict where a missile will land, you need to know the location of its launching pad. Assuming that a currency will depreciate simply because it has a higher rate of inflation ignores the possibility that it has "taken off" from a point of undervaluation in the first place. Such a currency might just as easily appreciate towards fair value rather than depreciate. Knowing the current relative inflation rate should not be sufficient, and yet this is a common approach to the problem.

Our rather unique approach to this problem was to try to put the concepts underlying the Strong and Weak Forms of PPP back together.

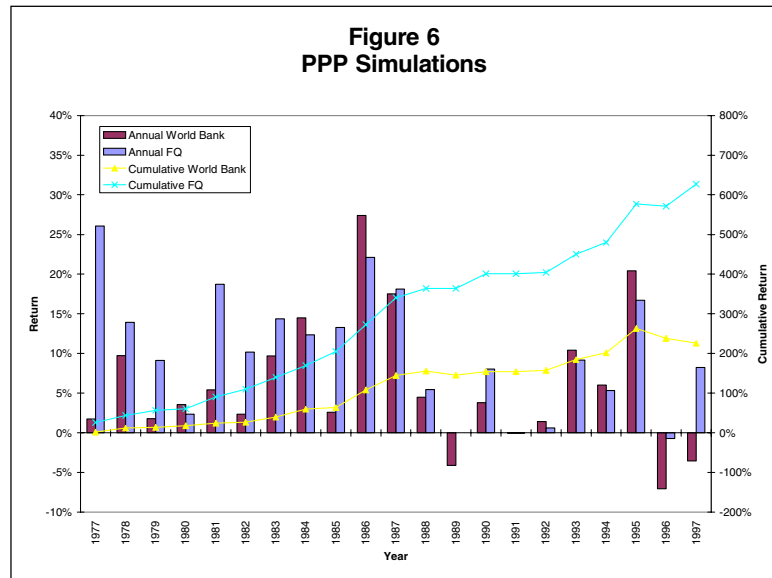
To do so, we took an approach towards identifying fair value that was quite different from the Law of One Price. Rather than identifying the weighted basket of goods that would allow us to compare prices through the exchange rate mechanism, and rather than trying to identify at which point in time the exchange rate was closest to its fair value, we took an entirely different route.

We implicitly rather than explicitly identified fair value. We began with the simplifying assumption that currencies probably vary randomly around their fair value. We wanted to see whether or not the data supported the view that the "average" exchange rate reflected fair value. To test this notion, we looked to see if the average exchange rate served as a useful guide to forecasting future exchange rate movements. If it did, then we assumed that the data supported this view.

Now the average nominal exchange rate cannot be useful to us because if a yen/dollar exchange rate was a fair value in 1980 at 200/1, then given the relative rates of inflation in Japan and the US since then (Japanese inflation has been lower), then the fair value should be somewhere less than 200/1, let's say 120/1, today. This is where the *Weak Form of PPP* is therefore integrated into our construction in that we adjust all previous nominal exchange rates by the changes in relative prices, effectively stating each exchange rate in today's terms. This leaves us moderately exposed to the problems associated with comparing baskets of goods, but far less exposed to those problems than the Law of One Price construction.

While First Quadrant's own PPP measure is a hybrid of the Strong and Weak PPP measures, it also seeks to take into account the possible presence of persistent errors. A simple way to describe what is done is that we subtract the mean historical error from the measure of mispricing. In other words, if the model has a persistently held view that a currency has been overpriced by 10%, our model will subtract that average "error" from the current measure of mispricing.

The relative performance of our approach as compared to the World Bank shows the advantage of this method.



It is interesting to note how applying our own methods to the World Bank's measure results in an impressive improvement in their own model. By applying an ex ante adjustment to the mean error of their PPP, the World Bank's information ratio rises nearly as high as the information ratio on First Quadrant's measure. The modification does not make the measures the same, of course, as the underlying construction is quite different, but the performance of their measure improves considerably as a result.

Another statistical issue that we face regards the relative degrees of confidence we have about measuring fair value accurately. Exchange

rates that swing widely from cheap to expensive may make it harder for us to gauge the fair value accurately (i.e., more volatility in the exchange rate will tend to produce a higher standard error around the mean – fair value – which we have measured). Where a measure is less reliable, we might want to assume less certainty about our measures and take proportionally less risk. A simple approach to this is to “standardize” the variance for each currency so that we effectively shrink the measures of mispricing where mispricing appears to swing widely relative to currencies where the mispricings vary much less. In an optimization context, this will lead to more similar levels of risk-taking in those currencies that appear to reach extreme mispricing as in those currencies that achieve only more modest mispricing.

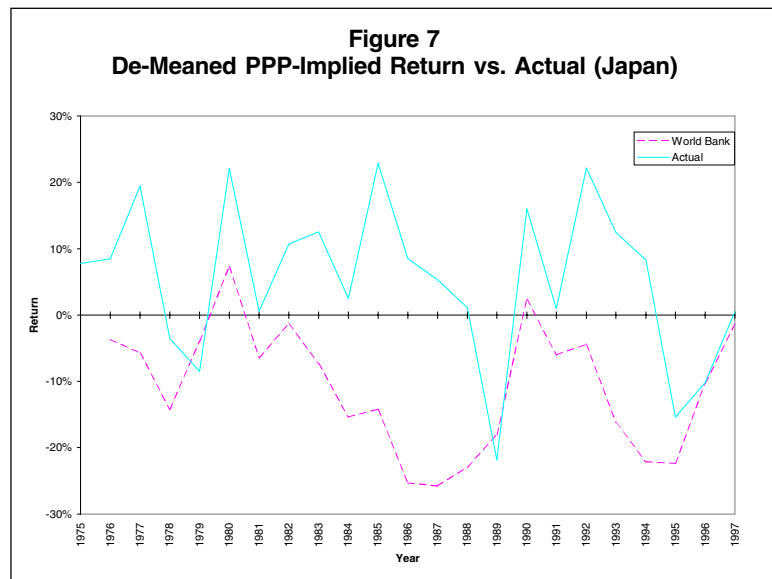
	<u>Raw</u>	<u>De-mean</u>	<u>Variance -adjust</u>	<u>Standardized</u>	<u>FQ</u>
Mean	6.62%	5.58%	86.83%	84.50%	89.30%
Std. Dev.	8.20%	4.52%	157.93%	63.82%	66.87%
Sharpe Ratio	0.81	1.23	0.55	1.32	1.34

By de-meaning the World Bank measure, the quality of the signal improved by 56%! Getting the means adjustment produces an enormous gain. With just variance adjustment, the World Bank measure *loses* power (by 22%). The change in the mean that occurs by Variance-scaling the signals (as the mean is not zero, scaling the series has the effect of scaling the mean as well) is forcing the measure *further* away from the correct value, costing performance.

The full standardization (i.e. de-meaning then variance adjusting) produces the best results, as we de-mean first, which allows the variance adjustment to avoid the pitfalls of moving means and capture the gains from helping the standard error problem.

Based on this examination, we find that, for Strong Form measures like the World Bank PPP Index, adjustments for the means problem are far more valuable than the adjustments for the standard error problem. Applying both methods, however, produces the best results.

Unfortunately for the World Bank measure, the Law of One Price doesn't say, "The prices of similar goods, *ex-ante standardized for the history of the country*, should carry the same effective price regardless of the currency in which that price is stated." In fact, if you review some of the individual countries' data, even with the "correction" of *ex-ante* standardization, the World Bank still appears to have problems with the mean.



In Japan, the World Bank measure still demonstrates a persistent bearish bias. Unfortunately for the World Bank, the actual returns do not demonstrate this bias. As these measures were de-meaned for these charts, this does clearly demonstrate the means problem that pure *Strong Form* PPP measures can have. Even de-meaning the measure can't remove all of the problem, which implies a significant structure problem inherent in the measure.

There are other, more sophisticated ways to "de-mean" PPP measures that try to take account explicitly of the underlying causes of drift in the real exchange rate. For example, the work of Balassa and Samuelson suggests that productivity shifts might lead to evolutions in fair value, while Stockman points to the changing composition of domestic and foreign goods in the demand curve as leading to a similar evolution. These are areas of current research at First Quadrant.

Conclusion

We should point out that our currency models do not rely solely on PPP to drive the signals. Though it does appear to have even some medium to short-term predictive power, other factors with a shorter-term horizon (such as relative interest rates, changes in yield curves, etc.) provide even stronger short-term predictive power. The combination of the two can be quite powerful as both the simulated results below indicate as well as our live performance since April of 1992.

January 1982 - December 1998

	Purchasing Power Parity (Long-term Value) Only	PPP and Short-term Factors*
Value Added	1.65	3.87
Tracking Error	4.56	3.08
Information Ratio	0.36	1.25

These information ratios were calculated in historical simulations run over the period 1982-1998 (our full model signals begin in 1982), using signals generated from the use of a PPP model only, and from our complete model that includes all of the factors we look at. The signals from April of 1992 represent First Quadrant's live signals. As you can see, the PPP model accounts for 29% of the information ratio over this specific period, leaving shorter-term factors to contribute 71% of the predictive information.

We have seen how the different philosophies of PPP construction change the measures in ways that are very difficult to correct. Given the data at hand, constructing an idea of the "fair-value" of a currency, with the intent of using it in a practical manner, turns out to be next to impossible, at least, if you try to hold to the dogma of either the *Strong* or *Weak* version. First Quadrant's own methodology relies philosophically on both, and, much like a hybrid in genetics, benefits from the strengths of each, and seems to avoid the pitfalls of either. It is also of great significance that First Quadrant's model is available real time, and not with a lag measured in months or years. We find today that the World Bank's "predictions" are not yet available for either 1998 or 1999. In fairness, they are not in the prediction business. We are.

Dori Levanoni
Associate Director

Max Darnell
Partner

Our research on Purchasing Power Parity is on going, and this paper represents a snapshot of our current views on PPP. Dori Levanoni is responsible for research on our currency models, and can be contacted at (626) 683-4136, or dori@firstquadrant.com.