

**THE REAL EXCHANGE RATE:
ISSUES OF CONCEPT AND MEASUREMENT**

Paper Prepared for a Conference In Honor Of Michael Mussa

Arnold C. Harberger

University of California, Los Angeles

June 2004

When the possibility was mentioned of my presenting a paper at this conclave, I jumped at the chance. First and foremost, I was motivated by the deep affection and respect that I have for Mike Mussa, whom I have known since his graduate student days, and whose career I have followed through every stage. We all know Mike as a leading figure in the firmament of international economics for the last 30 or so years. But I am not sure that everybody appreciates, at least sufficiently, what I have come to see as his most special, most defining trait and talent. To me he is among the least “derivative” economists of his generation. He belongs in no well-defined intellectual pigeonhole. He follows nobody else’s star. He is truly “his own man”, and as such he brings an element of originality and freshness to everything he puts his hand to. In an era in which too many seem guided or bound by models and methods, he has always put the problem and the evidence in center stage. For this, and for the diagnostic skills he has consistently displayed as he juxtaposed the problem and the evidence, we all are greatly in his debt.

While my purpose in being here is thus to honor Mike, my choice of topic was more guided by the attraction of the assembled company, which I pictured as consisting of much of the

rest of the firmament of contemporary international economics. Having all these stars in a single place at a single time constituted a temptation that I couldn't resist -- a temptation to roll out an agenda that I've had for years,¹ but that has left me feeling increasingly frustrated because the message wasn't anywhere near reaching its target. So here I go, making one more effort to win allies in the crusade.

The Concept

I take as my starting point the idea that the real exchange rate occupies a very important role -- it is the principal equilibrating variable of a country's international trade and payments. If the country imposes new import restrictions, the real demand for foreign currency will decline, and its real price should fall. If it imposes new export restrictions, the real supply of foreign currency will fall, and its real price should rise. Productivity advances in export goods should add to the real supply of foreign currency, those in import substitute goods should subtract from its demand. In both these cases the result should be a fall in the real price of foreign currency. But similar advances in nontradables productivity should lead to a rising real price of foreign currency. Rises in the real prices of export goods (e.g., oil for oil producing countries) lead to real appreciation of the currency (Dutch disease). Capital inflows have a similar effect so long as they are not spent exclusively on tradable goods.

I think that is a long enough list of how the equilibrium real exchange rate has to react to important kinds of real disturbance. The message here is that the RER has an extremely

¹Earlier efforts include: Arnold C. Harberger, "Economic Adjustment and the Real Exchange Rate," in Sebastian Edwards and Liaquat Ahamed, editors, Economic Adjustment and Exchange Rates in Developing Countries (Chicago: University of Chicago Press & NBER, 1986, pp. 371-423). Arnold C. Harberger, "Some Insights From Real Exchange-Rate Analysis," in Deepak Lal and Richard H. Snape, editors, Trade, Development and Political Economy: Essays in Honour of Anne O. Krueger (London: Palgrave, 2001, pp. 40-55).

important job to do in a market economy. And with so many sources of real disturbances, it's probably going to be busy most of the time, responding to one or another or yet another of the listed types of shock.

Some Measurement Issues

We are all familiar with the simplest of real trade models, in which we have just tradable and nontradable goods and where the real exchange rate is often represented as (p_t/p_n) or its reciprocal. This is an easy framework to work with when the disturbances are capital flows or changes in productivity in tradables or nontradables, but it runs into trouble with import and export restrictions, changes in world commodity prices, etc.

Maybe it is worthwhile to delve a bit into the cases of import tariffs and export taxes. We start with the fundamental proposition, first pointed out by Abba Lerner, that a uniform import tariff will have the same effects on imports and exports as a uniform export tax of equivalent strength. But they work in different ways upon the exchange rate. Consider a case where the free-trade exchange rate is 10 pesos to the dollar, and where the distorted situation has consumers of imports paying 12 pesos per dollar's worth while export producers get 8 pesos per dollar's worth. Under an import tariff of 50% this situation is generated by an exchange rate of 8 pesos received by exporters and a price of 12 pesos paid by importers consisting of an exchange rate of 8 pesos plus a tariff of 4. The same equilibrium in the export tax case has the importers paying a price of 12 pesos per dollar and the exporters receiving 12 pesos from the exchange rate but paying 4 in tax.

This example is sufficient, I think, to reveal the sort of trouble one gets into trying to extend the (p_t/p_n) definition of the real exchange rate to even slightly more complicated cases. The average internal price of tradables (gross price paid of 12, net price received of 8) is equal to

10 in both cases, yet the exchange rate is 8 in the one case and 12 in the other. So the question boils down to whether we prefer a definition of the real exchange rate that leaves it the same as we go from free trade to a 50% import tariff to a 50% export tax, or whether we feel more comfortable with a definition of the real exchange rate that tracks its role as the fundamental equilibrating variable of a country's trade and payments. I think this choice is obvious here. We must recognize that the job of the real exchange rate is to close an excess-supply gap when we impose a general import tariff, and to close an excess-demand gap when we impose a general export tax.

It's pretty hard to concoct a (p_t/p_n) -type definition that will lead to the same answer. The closest one can come, for the case just treated, is to say that p_t refers to the price of an undistorted tradable good. Then it becomes the price of the export good when we impose an import tariff and it becomes the price of the import good when we impose an export tax. This is no road for us to follow.

Definition of the (p_t/p_n) -type also get us into trouble when the disturbances in question are changes in the international prices of particular tradable goods; of which one is a perfect real-world example. The economics profession has long recognized how a rise in world oil prices can cause "Dutch disease" in an economy with large oil exports. Ask for a definition of Dutch disease, and most people will tell you that it is a real appreciation of a country's currency, stemming from a great surge in foreign currency receipts, which in turn come from a rise in the world price of a major export commodity (in this case, oil). But if we use an index of tradables prices in a (p_t/p_n) definition of the real exchange rate, such an index can actually end up indicating a depreciation of the real exchange rate if in this case the oil price has a sufficient weight to make the tradables price index p_t go up. To me, the Dutch disease case is very clear.

An increase in the flow of export receipts leads to an excess supply of foreign currency. The equilibrating variable -- the real exchange rate -- has to appreciate as a result. We therefore should steer clear of definitions of the RER that can produce ambiguous results in such a case.²

The “Natural” Structure For An RER Measure ($E\bar{p}^*/\bar{p}_d$)

Rather than starting with (p_t/p_n) , we set out on a much safer path if we start from the nominal exchange rate, E . I apologize to those who like to look at the exchange rate from the outside, but I have lived too long in that half of the world that looks at it from the inside. Hence my nominal exchange rate, E , reflect the nominal price in pesos of, say, a dollar of foreign exchange. The litany is, then

- i) E = the nominal peso price of the nominal dollar;
- ii) $\frac{E}{\bar{p}_d}$ = the real peso price of the nominal dollar; and
- iii) $\frac{E\bar{p}^*}{\bar{p}_d}$ = the real peso price of the real dollar.

So, moving to step ii) we face the question of what to use for \bar{p}_d . I would argue, very strongly, that the only indexes worthy of serious consideration are indexes that we would use to express the relative price of just about anything -- i.e., to get the real price of hotel rooms, telephone calls, loaves of bread, strawberry jam, whatever. In short we are looking here for a “natural” numeraire for a national economy. The key criteria for a “natural” numeraire are: a)

²It is interesting, though of less practical importance, to pursue the case of a rise in the world price of an import good. Here the required equilibrating adjustment of the exchange rate can go either way, depending on whether we have an elastic or an inelastic import demand for the goods whose world price has risen. Yet while in reality the equilibrating move can be either up or down, the exogenous component of a p_t index would go up in all cases, again raising the possibility of our getting the “wrong” answer from an RER concept defined by (p_t/p_n) .

that be of well-defined and broad coverage, and b) that it be reasonably free of idiosyncratic movements and c) that it be easily available for the use in question. For example, it would be utterly absurd to measure everything (real wages, GDP, the nation's wealth, etc.) in terms of barrels of oil or kilos of sugar -- which would lead to the main movements of everything being dominated by those of the world price of oil or of sugar. Such candidates lose because their movements are too idiosyncratic. The only candidates that I feel meet all three criteria reasonably well are the GDP deflator and the consumer price index. The GDP deflator wins hands down on the criterion of generality, and the CPI wins on the criteria of being available easily, promptly, and usually on a monthly basis. Which one to use is thus likely to be a judgment call, with the winner varying as one shifts from problem to problem.

But before we go on, please take note that either of these two is a sensible deflator for the relative price variable that enters a demand function for imports or for tradables, or a supply function for exports or for tradables. Note also that each of these includes both tradable and nontradable components. Thus, in the demand for imports a rise in (p_m / \bar{p}_d) would reflect an incentive for demanders to substitute away from imports and toward nontradables and for the production of import substitutes to increase, and a rise in (p_x / \bar{p}_d) would reflect an incentive to draw productive factors away from nontradables and toward exports. In these aspects we are definitely retaining some of the flavor of the (p_t/p_n) framework, but hopefully without the weaknesses it reveals in extended applications.

Once we get through step ii) we have a measure E/\bar{p}_d which tells how many GDP baskets or how many CPI baskets it takes to buy a nominal dollar of foreign exchange -- certainly a simple, easily understandable, and very concrete and straightforward idea. It also makes thorough sense that this should be an equilibrating variable of a country's trade and

payments for any problem in which external prices are not changing too rapidly or too much. In this framework it fits perfectly with what I think we all learned about adjustment under fixed and flexible exchange rates. In principle, an economy can reach its equilibrium real exchange rate under either system, with the adjustment taking place in \bar{p}_d when E is fixed and in E (or in both E and \bar{p}_d) when it is flexible.

I think that for almost all short-run problems and even for some extending over as long as a decade, the definition E/\bar{p}_d is quite adequate for real exchange rate analysis. The need to incorporate a world price variable \bar{p}^* arises only when there is sufficient variation in the relevant world prices to warrant making an explicit correction for it. But, then again, why not make the correction in a standard, routine way, so nobody has to worry about whether or not it was really needed?

The Choice of the World Price Index \bar{p}^*

If only the answer were as simple as the above question makes it sound! The truth is that different real-world situations, different types of problems can lead one toward different definitions of \bar{p}^* .

Let us start with the situation of a particular country. What is more natural than to take the goods and services represented by that country's own trade, its own links with the world economy, and build an index out of those specific prices. This would be an index of the world prices of the country's export and import goods. But we already have seen the inconvenience that this would bring in cases of major movements in the price of the country's main export products.

This leaves us, at least in my own view, with an index of the world prices of the country's imports. This makes a great deal of sense, both analytically and intuitively. In this case

$E\bar{p}^*/\bar{p}_d$ answers the question, how many baskets of our imports can we buy with one of our production (GDP) baskets or with one of our consumption (CPI) baskets? This would be my nominee for a RER measure that was tailored to each country's trade pattern.

Why cannot this definition be made into the general canonical measure that everyone uses? The answer is that only a relatively few countries make serious efforts to calculate a reasonably comprehensive index of the world prices of their imports. Some have no import price index at all; others have highly unrepresentative indexes of a few, usually primary-commodity import prices. So, short of mounting a worldwide effort to do a good job for most if not all countries, I think we have to leave this genuinely country-specific measure of \bar{p}^* as a worthwhile luxury that only some can afford.

If we depart from the specific list of goods that each particular country imports, where do we then go? My candidate would be a world price index of tradable goods. The analogy is direct between such an index and the main national price indexes that we regularly collect and use. In general we do not have GDP deflators by state, by province, or by region. When we're studying a regional or local market and want to express its price in real terms, we typically deflate it by the national GDP deflator. It is the same with respect to the CPI, although a few countries have independent CPIs for different metropolitan areas or other subdivisions.

Thus, we are regularly making statements about how many national production baskets could be bought with the output of a given industry or region, or with the value added of a given industry or firm. Anyone who doubts this should consult the standard practice in ex post project evaluation. Here the profile over time of nominal "benefits minus costs" is virtually meaningless if there has been any substantial degree of inflation. At the same time it is utterly essential to deflate all nominal flows by the same deflator in order to be able to calculate a net present value

or an internal rate of return. And standard practice assigns this job to either of our two general indexes -- the GDP deflator or the CPI.

My question, then, is why do we not develop a world price index of tradable goods, for use as \bar{p}^* in our real exchange rate calculations, and for other useful purposes as well. Our real exchange rate question then becomes, how many local GDP (or CPI) baskets does it take to buy one standard basket of world tradable goods. A perfectly sensible question, and one that quite adequately reflects the role of the real exchange rate as the principal equilibrator of a country's trade and payments. When we analyze a rise in the price of oil for an oil exporter (or importer, for that matter) we express it in terms of \bar{p}^* ; i.e., $\Delta(p_O^* / \bar{p}^*)$. And when we trace the country's reaction to this change, we look for the process by which (E / \bar{p}_d) moves relative to \bar{p}^* to restore the country's payments equilibrium.

The IMF Is The Natural Entity To Develop A Standard Index Of World Prices of Tradables

We have seen how a world price index of tradables would be useful in calculating real exchange rates and in calibrating the movements of the relative prices of particular tradables. For these purposes alone, history and fate would tap the shoulder of the IMF. But the case can be made still stronger by considering some tasks that have long been normal functions for that organization. The Yearbooks of International Financial Statistics have from the outset given time series of the U.S. dollar values of exports and imports of member countries. These series cry to be converted into real terms. What meaning can be attached to the fact that Colombia's exports grew from U.S. \$538 million to U.S. \$6766 million between 1965 and 1990; when the general level of dollar prices more than tripled in this period? It would add significantly to the

meaning of these and other time series of trade values if they were expressed in units that were more relevantly comparable than simple nominal U.S. dollars or any other nominal unit.

A world price index of tradables would have another advantage in that it would naturally render “unit free” such items as the real value of Colombia’s exports. By this I only mean that it would not matter if we measured Colombia’s exports in dollars or pounds or Euros, for then, as we move to real values we would have to deflate by the world price index expressed in dollars, pounds or Euros. In moving from X_t / \bar{p}_t^* measured in dollars to $X_t' / \bar{p}_t^{X'}$ measured in pounds, we would simply multiply both numerator and denominator by the exchange rate giving the number of pounds per dollar. The real value of Colombia’s exports would not change at all.

Distinctly second best to a genuine world price index of tradables is what I call the SDR-WPI. This is an index that my students and I have been using for some 20 years as \bar{p}^* in the real exchange rate formula. $E\bar{p}^* / \bar{p}_d$. It is built up on the premise that on the whole, the price indexes that go under the label “wholesale price index” or “producer price index” are heavily weighted with tradable goods. This certainly was true in their origin and appears to be still true today. In any case, operating on this premise, we consider the wholesale or producer price indexes of the principal industrial countries to be independent indicators of the world price level of tradables, translated, of course, into true currency of the country in question. This leads to an index \bar{p}^* equal to

$$\bar{p}_{it}^* = \bar{p}_{i,t-1}^* \left[\sum_j f_{j,t-1}, \frac{E_{ij,t}}{E_{ij,t-1}}, \frac{P_{wjt}}{P_{wj,t-1}} \right].$$

Here $E_{ij;t}$ is the price of j 's currency in terms of i 's, at time t , and P_{wjt} is country j 's wholesale price index at time t . The f_{jt} are the weights that the IMF assigns to country j 's currency at time t in its calculation of the SDR. To deal with changes in SDR weights we have

found it convenient to allow the weights to change slowly over a 12 or 24 month period surrounding the actual date of the official change.

Note that when this definition of \bar{p}^* is used, we have a real exchange rate for country C that is totally independent of whether country i is the U.S., the U.K., or Japan. If we start with $E\bar{p}^*/\bar{p}_d$ with E being pesos per dollar and \bar{p}^* being an index of dollar prices, we can move to yen by taking $E'\bar{p}^*/\bar{p}_d$, where $E' = E \times \text{dollars per yen}$, and \bar{p}^* is an index of the yen prices of tradables = \bar{p}^* times yen per dollar. The real exchange rate is, as it should be, an attribute of the economy of country C, independent of which foreign currency is used to keep the accounts.

I do not really want to suggest that the IMF go down the road of “adopting the SDR-WPI, as its index of the world price level of tradables. Why? Because a tradables index directly calculated from the prices actually prevailing in world markets would be a great deal better. But failing this, I believe the SDR-WPI would provide a fully acceptable index for calculating real exchange rates (and for other applications where a world price level of tradables is needed).

The IMF's Current Measures of Real Effective Exchange Rates

This is the saddest section of my paper. The problem is this -- that in its excursions into real exchange rate measurement, the IMF has ended up with two main measures, neither of which appears to capture well the main function of the RER as the principal equilibrating variable of a country's trade and payments.

Earlier in this paper, I complained about trying to measure the real exchange rate by going directly to the prices of tradable and nontradable goods, but that does not deny in any way the general usefulness of what many of us have come to call the Australian approach to real exchange rate analysis. As I see it the essence of this approach is that the real exchange rate is

an attribute of each individual economy as it faces world markets. It is in no conceivable sense a bilateral measure between two countries. And it should in principle be unaffected by events in other parts of the world that have no direct impact on this country's trade and payments.

There is little meaning to be found from knowing how many Indian haircuts can be bought with the proceeds of one English haircut. Nor can much be made from knowing how many barrels of oil in the United States can be bought with one barrel of oil in the United Kingdom (both, presumably, net of taxes). In the first of these cases we are dealing with a nontradable good (haircuts) whose price should be set by the local supply and demand conditions of each country. In the second case we are dealing with a situation where the law of one price should work pretty well. Now, if an independent calculation of the "terms of trade" between haircuts in two countries would give information that is irrelevant for RER analysis, does it not also seem irrelevant to have haircut prices appear (along, of course, with other nontradables prices) in both the numerator and the denominator of a real exchange rate measure. Yet that is exactly what happens with an RER measure "based on relative consumer prices". At the opposite extreme, if it makes no sense to take the ratio of net-of-tax oil prices in two countries, does it not follow that an RER index based on relative wholesale prices would be seriously contaminated by many commodities which are subject to any significant degree to the law of one price? Yet International Financial Statistics presented real exchange rates based on relative wholesale price and relative export price indexes as late as 1997. I am happy to report, however, that it no longer does so.

Unfortunately, all the real exchange rate concepts that have appeared in IFS have always been symmetrical, i.e., using the same type of index to deflate both the numerator and the denominator of the nominal exchange rate. I claim, in contrast, that the fundamental

equilibrating role of the real exchange rate calls for, even cries for, quite distinct deflators to be used in numerator and denominator.

The second main real exchange rate measure currently reported in IFS is based on relative unit labor costs (or relative normalized unit labor costs). This measure strikes me as suspect, right from the outset. In the first place, it is thought of by its advocates as a measure of relative competitiveness, which I believe is quite different from a real exchange rate concept. For example, a productivity improvement in a country's export goods would improve the competitiveness of its exports in world markets, but would lead to an appreciation of the real exchange rate as I believe it should be defined, and certainly to an appreciation of the RER by any of the measures I am here suggesting (measuring \bar{p}^* by an index of the world prices of the country's own import basket, measuring it by a yet-to-be-created index of the world prices of tradables, or measuring it by what I have called the SDR-WPI). The IFS measure based on relative unit labor costs would be subject to two forces -- depreciation due to the productivity increase and an appreciation due to the exchange rate response. The normal result would be a net depreciation, using this measure, on the LeChatelier principle that the effect should not outweigh the cause. (The productivity increase shifts the supply of exports to the right, causing an excess supply of foreign exchange, which is alleviated by movements along both the export supply and the import demand curves). A net depreciation is the right answer if the question we're asking concerns export competitiveness, but it is the wrong answer if we are asking what happens to the real exchange rate.

It is a mystery to me how the literature on the real exchange rate ever got bogged down with the question of competitiveness. Quite obviously, a country's competitiveness will vary from product to product, and from market to market. Is it not more natural to measure

competitiveness directly in this sense, i.e., what is happening to country i 's share in country j 's imports of product k , or by direct comparison of prices of a given good from different sources?

The mystery compounds as one turns to the actual measure of unit labor cost, which is used to reflect a country's international competitiveness. The Fund's own Primer³ on the subject states:

Unit labor costs are calculated as the ratio of hourly compensation in manufacturing to measured labor productivity in that sector. The latter is obtained by dividing sector gross value added by the product of employment and average hours worked per person.
(p. 11)

Proceeding directly from this definition, I take it that hourly compensation is wages bill divided by total hours (Nh), and that productivity is gross value added divided by total hours (Nh). Hence unit labor costs are simply the share of wages in total value added!! I cannot imagine how this could be thought of as a measure of competitiveness!

Just to illustrate the weakness of the unit labor cost measure consider that the big differences in real product per worker between developed and developing countries stem even more from differences in the production function than from differences in factor proportions. Now suppose two Cobb-Douglas functions, $A_p L_p^\alpha K_p^{(1-\alpha)}$ and $A_r L_r^\alpha K_r^{(1-\alpha)}$. These refer to a particular product. The labor share α is the same, but the constant terms differ between a poor and a rich country. Yet as I read the Primer's definition, the unit labor cost variable would be the same in both countries, being precisely equal to α in both cases. True productivity, of course, would differ in the proportion A_r/A_p .

³IMF Policy Development and Review Department, "A Primer on the IMF's Information Notice System," prepared by Alessandro Zanello and Dominique Desruelle (IMF Working Paper, May 1997), see esp. pp. 6, 11.

The “normalized unit labor cost” concept does not get around this problem, as it simply is what results after passing the non-normalized series through a Hodrick-Prescott filter.

I have little more to say on this subject. My first plea is to get rid of the confusion caused by thinking that the real exchange rate concept is somehow supposed to reflect a country’s competitiveness in world markets. What can be said is that when for some other reasons (e.g., capital inflows), real appreciation occurs. This reduces the country’s competitiveness in world markets. But when the disturbance comes from real cost reductions, we get increased competitiveness together with RER appreciation. And please do not forget that real cost reductions are the principal wellspring of most economic growth.

My second plea is, if the IMF is going to publish an index of competitiveness, that it should surely label it as such, and at the same time try to find a better index than one that identifies competitiveness with the ratio of labor costs to total value added.

The Need For At Least One Asymmetric Index of the Real Exchange Rate

I find it very hard to understand why so much of the literature continues to focus on symmetric measures of the real exchange rate when the theory of the subject gives such a central role of the distinction between tradable and nontradable goods.⁴

Edwards makes extensive empirical applications of asymmetric real exchange rate measures. He states:

In the theoretical chapters of this book the real exchange rate was defined as the relative price of tradables to nontradables: $RER = \hat{E}p_T^* / p_N$In the empirical chapters ... p_T^* has been proxied by the wholesale price indexes (WPIs) of the country’s trade partners.

⁴See, for example, Rudiger Dornbusch, Exchange Rates and Inflation, (Cambridge: MIT Press, 1988).

...With respect to the domestic price of nontradables we have proxied it by the country's consumer price index (CPI).⁵

I agree fully with Edwards's choice of a measure for the RER. I would only modify his explanation slightly. Rather than consider the consumer price index to be a proxy for the domestic price of nontradables, I would simply treat it as a general index of the form $\bar{p}_D = \lambda \bar{p}_T + (1 - \lambda) \bar{p}_N$. Under this approach, the real exchange rate measure $E\bar{p}^* / \bar{p}_D$ becomes

$$\text{RER} = E\bar{p}^* / [\lambda \bar{p}_T + (1 - \lambda) \bar{p}_N].$$

If one is ready to accept the approximate equality of $E\bar{p}^*$ with \bar{p}_T , the internal price level of tradables, then the RER becomes

$$\text{RER} = 1 / [\lambda + (1 - \lambda)(\bar{p}_N / \bar{p}_T)],$$

a measure whose movements vary directly with (p_T/p_N) .

It is worth noting here that it is virtually impossible to actually create a proper index of the price level of nontradables that really fits the theory. The reason is that, under the small-open-economy hypothesis, the prices of tradables are determined in world markets. It does not matter whether domestically produced tradable goods use or do not use nontradable inputs, nor does it matter what proportions these occupy in the total cost of a tradable good. The world price of that final good, p_j^* , is all that matters.

It is not so, however, in the case of nontradable goods. The prices of these goods are determined by the forces of supply and demand in the domestic marketplace. Yet those domestic prices p_i are typically mixtures of the prices of tradable inputs as well as those of nontradable

⁵Sebastian Edwards, Real Exchange Rates, Devaluation, and Adjustment, (Cambridge: MIT Press, 1989), p. 88.

inputs. A true index of tradables prices for the country as a whole would incorporate not only the prices of tradables that were final goods, but also the tradable-inputs component of the prices of “nontradables”. The true nontradables prices index would incorporate only the nontradables-inputs component of the nontradables final products of the economy. If this sounds complicated, simply work through the mental experiment of what happens when a country’s currency is devalued. Obviously all final-product tradables go up in price to reflect the devaluation. But so also do many prices that are set by supply and demand in the local marketplace. These rise simply to reflect the increase in the local-currency prices of their tradable inputs.

Thus, the nontradable price index \bar{p}_n is by its nature not directly observable. We could extract it indirectly from $\bar{p}_d = [\lambda\bar{p}_t + (1-\lambda)\bar{p}_n]$, by making assumptions or guesses as to the size of λ . My response to all of this is, why try? We have a perfectly good index in the form of $E\bar{p}^*/\bar{p}_d$, and \bar{p}_d is easily measured and readily available. It is easy to simply accept and live with the fact that $E\bar{p}^*/\bar{p}_d$ does not measure (\bar{p}_t/\bar{p}_n) , but does quite clearly move up and down in sympathy with it.

Why has not the profession and most especially the IMF itself, not taken the “natural” road as exemplified by Dornbusch and Edwards? Why has the empirical measurement of real exchange rates been so dominated by symmetric approaches rather than the asymmetric ones dictated by the theory of the subject?

In trying to answer these questions I can only speculate as to the reasons. Two elements come to mind here. One is the long tradition in the international trade literature of framing examples in terms of two countries (A and B) facing each other in a sort of symmetric way. This is how my generation was taught, with reciprocal demand curves (offer curves) playing a big role in the exposition. This way of thinking naturally leads to a symmetrical treatment, because it

renders the rest of the world (country B) as being “just like” or “quite akin to” country A. I think that the development of small-open-economy macroeconomics represents a big advance over the traditional treatment. The distinction between tradable goods and nontradables is an integral part of this approach, which also generalizes quite easily to a multi-good framework with many importables, many exportables, and many nontradables. Once we are here, the real exchange rate becomes a key variable not only in the process of adjustment but also in determining which goods will end up as exports, which as imports, and which as nontradables. That is, the real exchange rate becomes the arbiter determining how a country’s comparative advantage changes in response to different disturbances.

My second answer to the above questions revolves around purchasing power parity (PPP). PPP is not intrinsically a bilateral concept but I believe that is the way most of the literature has treated it. When people worked with pesos per dollar they deflated by a local peso index and a U.S. price index. With pesos per pound, the U.S. index was replaced by a U.K. index, etc. On the basis of this tradition, people thought of real exchange rate movements as being deviations from PPP, often thinking of them as reflecting out-of-equilibrium “misalignment” rather than equilibrium responses to real shocks. I truly believe that most observed movements of the real exchange rate are best interpreted either as equilibrium movements (with the economy moving from one equilibrium to another) or, more realistically, as observations of the dynamic process by which the economy is always trying to go to a new equilibrium, determined in turn by the successive shocks to which that economy has been subject.

I think that just about everybody would agree with Peter Montiel when he says:

Conceptually, monetary neutrality implies that relative PPP should hold in the face of nominal shocks, but theory does not

require that PPP hold in response to real shocks. Empirically, the evidence suggests that PPP may or may not hold ... over spans of time [of more than 50 years]. [But] over shorter periods of time -- including the complete span of data ... available for developing countries -- ... PPP provides a poor approximation of the equilibrium exchange rate.⁶

In a subsequent chapter, Montiel supplies examples of how the equilibrium real exchange rate responds to foreign transfers, terms of trade changes, and differential productivity shocks. His results here are quite in line with the simple examples given earlier in this paper. And in this chapter he defines the real exchange rate in the traditional theoretical manner, i.e., the price of tradables relative to that of nontradables.⁷ Yet in the earlier chapter, speaking of tests of purchasing power parity he reverts to the symmetrical definition “where p and p^* are ... the domestic and foreign price levels, respectively (including both traded and nontraded goods)” (p. 234). Perhaps this will help to confirm my hunch that the tendency to use symmetric definitions of the real exchange rate has its roots in the concept of purchasing power parity. I cannot think of an argument which would favor a symmetrical definition when the purpose is to reflect the fundamental role of the RER as the principal equilibrating variable of a country’s trade and payments.

⁶Peter J. Montiel, “The Long-Run Equilibrium Real Exchange Rate: Conceptual Issues and Empirical Research,” in Lawrence E. Hinkle and Peter J. Montiel, editors, Exchange Rate Misalignment (Oxford: Oxford University Press and the World Bank, 1999, p. 262).

⁷Peter J. Montiel, “Determinants of the Long-Run Equilibrium Real Exchange Rate: An Analytical Model,” in Hinkle and Montiel, op. cit., pp. 264-90.