

Board of Governors of the Federal Reserve System

International Finance Discussion Papers

Number 917

January 2008

Measuring U.S. International Relative Prices: A WARP View of the World

Charles P. Thomas, Jaime Marquez, and Sean Fahle

NOTE: International Finance Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. References to International Finance Discussion Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors. Recent IFDPs are available on the Web at www.federalreserve.gov/pubs/ifdp/. This paper can be downloaded without charge from Social Science Research Network electronic library at <http://www.ssrn.com>

Measuring U.S. International Relative Prices: A WARP View of the World

Charles P. Thomas, Jaime Marquez, and Sean Fahle*

January 2008

Abstract: In this paper we construct a new measure of U.S. prices relative to those of its trading partners and use it to reexamine the behavior of U.S. net exports. Our measure differs from existing measures of the dollar's real effective exchange rate (REER) in that it explicitly incorporates both the difference in price levels between the United States and developing economies and the growing importance of these developing economies in world trade. Unlike existing REERs, our measure shows that relative U.S. prices have increased significantly over the past 15 years. In terms of simple correlations, the relationship between our measure of relative prices and U.S. net exports is much more coherent than that between existing REERs and net exports. To explore this relationship further, we use our measure to construct an index of foreign prices relevant for U.S. export volumes and reexamine several export equations. We find that export equations with the new index dominate those with previous measures in terms of in-sample fit, out-of-sample fit, and parameter constancy. In addition, we find that with the new index of foreign prices the estimated elasticity of U.S. exports with respect to foreign income is a good bit higher than the unitary elasticity found in previous studies using other price measures. This has implications for U.S. current account adjustment.

Keywords: Automated Model Specification, China, Competitiveness, IMF, FRB, Geometric Aggregation, Penn World Tables, Real Effective Exchange Rates, Trade Elasticities

JEL Classification: C82, F41, C22

* We are grateful to George Alessandria, Tam Bayoumi, Menzie Chinn, Judith Dean, Neil Ericsson, Robert Feenstra, Michael Ferrantino, Joe Gagnon, Dale Henderson, Karen Johnson, Michael Leahy, Jaewoo Lee, Gian Maria Milesi-Ferreti, Alessandro Prati, Trevor Reeve, Luca Ricci, Nathan Sheets, Antonio Spilimbergo, and Eric van Wincoop for their detailed comments. A previous version of this paper was presented in the workshop series of the Federal Reserve, the International Monetary Fund, the U.S. International Trade Commission, and the Fall 2005 and 2007 meetings of the Midwest International Economics Group. The calculations use TROLL: Hollinger (2007); PcGive: Hendry and Doornik (1999); and PcGets: Hendry and Krolzig (2001). The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

1 Introduction

In this paper we assemble a new measure of international relative prices to gauge the average amount by which U.S. prices differ from foreign prices. Interest in developing such measures in international economics is not new.¹ What is new in this paper is the focus on the interactions between the dispersion of prices across countries and the increased trade with emerging economies. Recognition of these interactions yields a picture of U.S. international relative prices that is fundamentally different from the one given by existing measures of the real effective exchange rate. Indeed, unlike existing measures of relative prices, we find a significant increase in U.S. prices relative to its trading partners over the past 15 years. Further, most of this increase owes to greater trade with developing economies rather than increases in U.S. prices relative to individual countries.

Our measure differs from those currently available for two reasons. First, we measure bilateral relative price *levels*, as opposed to bilateral relative price *indexes*. Second, we use an aggregation method that retains the information embodied in those levels. In contrast, existing measures of relative prices are constructed by either chaining or averaging indexes – that is, they begin with price and exchange rate indexes constructed to have a value of 100 in a base year so that the value of the index in a given period indicates how much prices have changed since the base year. Thus, multi-country aggregates of these indexes measure the average change relative to the base year. Such methods are ideal if the purpose is to measure average changes in bilateral real exchange rates but not for measuring the level of U.S. prices relative to prices elsewhere.

Of course, others have recognized the importance of differentiating price indexes from price levels.² But the implications of combining that distinction with the increased role of developing countries in world trade has not received attention. In particular, the fact that prices in some developing economies are systematically below those in developed economies, combined with the fact that emerging economies' share of world trade has been increasing, has led to a decline in the average world price of traded goods even though prices in individual countries have not fallen. Aggregates based on price indexes cannot capture this interaction between price levels and trade shares. Our weighted average relative price (WARP) is designed specifically to capture this interaction and does so by using a geometric aggregate where the weights capture the change in the structure of U.S. external trade.

Section 2 reviews the evolution of several well known real effective exchange rate indexes. Al-

¹For a recent review, see Froot and Rogoff (1995). For early work on the importance of measuring relative prices, see Keynes (1925), Kravis and Gilbert (1954), and Kravis and Lipsey (1971). Other relevant papers include Lipsey, Molinari, and Kravis (1990), Hooper and Richardson (1991), and Turner and Van't dac (1993).

²For example, Turner and Van't dac (1993) examine this distinction using cross-sectional data.

though these indexes differ in source data and aggregation scheme, they generally paint a similar picture: U.S. prices relative to foreign prices have risen and fallen since 1975 but, on balance, they show no trend. Section 3 presents the WARP, discusses a few of its properties, and compares it to other measures. According to WARP, U.S. prices have risen significantly relative to its trading partners' prices since 1975 with most of the increase occurring since 1990.

This upward trend in U.S. international relative prices constitutes the main result of this paper. Section 4 examines several factors responsible for this upward trend: choice of price data, aggregation method, and currency basket. We find that the upward trend owes to the aggregation of relative price levels as such and to the shift in U.S. trade patterns away from the relatively high-price industrial countries toward the lower-price developing economies. Section 5 examines the sensitivity of this upward trend to both parametric structures and measurement errors; we find that the upward trend of U.S. international relative prices is robust.

Section 6 addresses whether WARP can be thought of as a measure of competitiveness. A point that comes clearly from the analysis is that any reasonable measure of competitiveness will necessarily incorporate the prices of non-traded goods and services as well as the prices of traded goods and services. Indeed, with analytical examples we show why the suitability of a measure of competitiveness to a particular application is largely an empirical question. With this in mind, Section 6 also looks at the relationship between WARP and the U.S. trade balance. We find that in terms of simple correlations, the relationship between relative prices and the U.S. trade balance (as a share of GDP) is much tighter when one uses WARP than when using conventional measures of real effective exchange rates. To explore why this might be the case, we examine several econometric specifications for the volume of U.S. exports. The focus is on assessing the implications for parameter estimates of using WARP-based and other measures of foreign prices to construct a relative price of exports. Our goal is not to offer detailed specifications for exports but, rather, to see if the WARP passes the "proof of concept" test. The evidence suggests that it does.

2 Existing Measures of Relative Prices

Existing measures of the dollar's real effective exchange rate (REER) are designed to reflect how much, on average, U.S. prices have *changed* relative to the prices of its trading partners.³ The top panel of figure 1 shows the measures constructed by the Federal Reserve, the OECD, and the IMF,

³The theoretical underpinnings of the REERs date back to work by Armington who, as McGuirk (1986, p. 3) points out, showed that an ideal weighting system is one in which an equiproportionate change in the product prices of all countries would leave the demand for any one country's product unchanged.

all of which are based on relative CPIs.⁴ Though they differ from one another in many important methodological respects, they all show two common features. First, over the past thirty years, U.S. relative prices have changed little on average, a property that is at odds with the growing U.S. current-account deficit. Second, over shorter periods, U.S. international relative prices deviate substantially from their long-term mean and indeed these prices reached a historical peak in 1985.⁵

These three measures are constructed by aggregating bilateral real exchange rate indexes. That is, they begin with bilateral nominal exchange rate indexes and adjust them by relative movements in U.S. and foreign consumer price indexes. These bilateral real exchange rates (indexed to 1973=100) are shown in the middle panels of figure 1. The left panel plots the indexes vis-a-vis selected industrial countries; the right panel plots the indexes vis-a-vis selected emerging economies.⁶ There is clearly a good deal of dispersion among these bilateral indexes, indicating that the CPI-adjusted value of the dollar has risen relative to some countries' currencies and fallen relative to others. On a bilateral basis, these real exchange rates can be interpreted as changes in relative prices.

Given the dispersion of bilateral real exchange rates across countries, it is hard to tell if there is a general pattern to the movements. This is the point of a REER: to distill these various movements into a single measure. To do so requires a weighting scheme. The aggregates shown in the top panels use weights based on trade shares. The weights used by the Federal Reserve Board in its Broad Real Index are representative; a selection of these is given in the bottom panels of figure 1.⁷ We note the increasing weight given to developing economies, especially China and Mexico, since 1990.

How has the increased weight of the developing economies affected the REERs? If one looks at both the increase (depreciation) in China's bilateral real exchange rate since 1973 (middle right panel) and the increase in China's weight in U.S. trade since 1990 (bottom left) one might conclude that China's real exchange rate has had a significant impact on the dollar's REER. However, in fact, China's real exchange rate has had a relatively small effect on the dollar's REER. The mechanics for this result vary with the particular REER used but, in general, the reason is that most of the

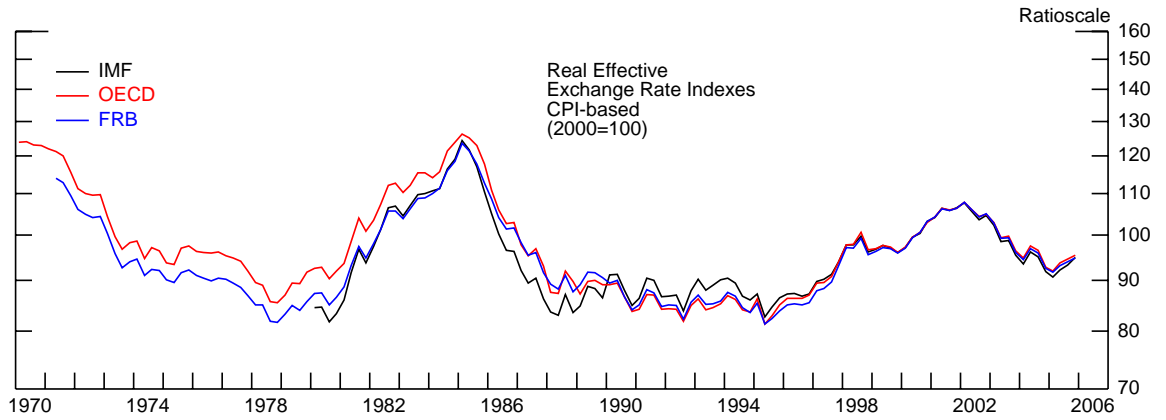
⁴Both the OECD and the IMF also report real effective exchange rates that are based on unit-labor costs; these measures show pronounced secular declines.

⁵Chinn (2005) reviews these measures. For the IMF, see Bayoumi et al. (2005), Zanello and Desruelle (1997), Turner and Golub (1997), Maciejewski (1983); for the Federal Reserve, see Hooper and Morton (1978), Pauls (1987), Leahy (1998), Loretan (2005); for the OECD, see Durand, Simon, and Webb (1992), Durand, Madaschi, and Terribile (1998). The BIS also constructs a real effective exchange rate comparable to that of the OECD and the Federal Reserve but the series starts in 1994 and so it is not suitable for our analysis; see Klau and Fung (2006).

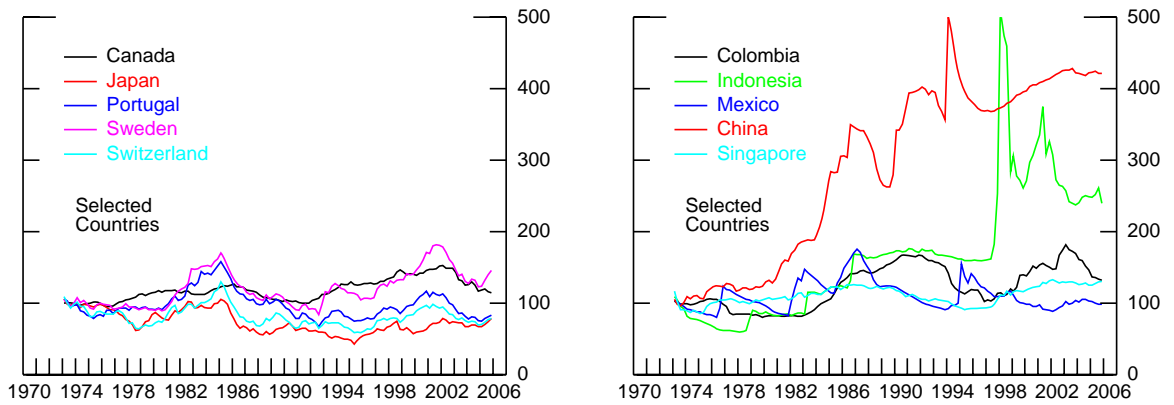
⁶The countries for each group were selected so as to encompass (form an envelope around) all the bilateral real exchange rates in our sample of 34 countries. The group of industrial countries corresponds to the Federal Reserve's classification of countries with major currencies; the group of emerging economies corresponds to the Federal Reserve's list of currencies of Other Important Trading Partners (OITP). These country groupings coincide, respectively, with the "High-price" and "Low-price" countries.

⁷The FRB's weight for a given country consists of trade shares for bilateral non-oil imports, bilateral exports, and a measure of the importance of the competition between that country and the United States in third-country markets. For further details on the construction of these weights, see Leahy (1998).

Real Effective Exchange Rates



Bilateral Real Exchange Rate Indexes - CPI-based (1973=100)



Aggregation Weights (Federal Reserve Board's Measure)

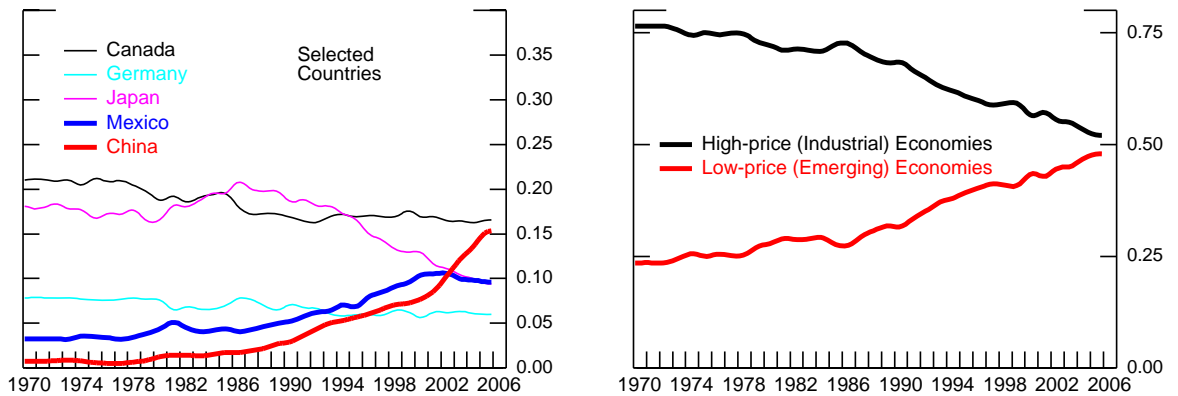


Figure 1: Real Exchange Rate Indexes: Effective and Bilateral – Selected Institutions

increase in the dollar real bilateral exchange rate vis-a-vis China occurred prior to 1990, a time when China's weight in U.S. trade was relatively small. REERs designed to show average changes do not get much of a boost from the Yuan's real depreciation prior to 1990 because China's weight in the index was small during that period. Conversely, despite the increase in the weight of China after 1990, there has not been much real depreciation of the Yuan during the period when the weight was large, so, again, the REERs do not get much of a boost. In general, what matters for existing measures of real effective exchange rates is whether the bilateral exchange rates are changing and, if they are not changing much, then increasing the weights on these countries does not cause the REER to change.

If the sole objective is to measure *changes* in the real effective exchange rates, then one can hardly improve upon existing measures. What we argue is that an exclusive focus on such changes carries a loss of information, that this loss is more than a theoretical possibility, and that the increased participation of low-cost producers in the world economy gives it economic significance.

3 The WARP

3.1 Intuition

The basic idea behind our aggregate is simple. Suppose, for expository ease only, that we have the foreign-currency price of a basket of goods in a foreign country i (call it P_i), and that we also have the dollar price of the same basket in the United States (call it P_{us}). As shown in equation (1), by multiplying the ratio of these prices by the market exchange rate we can define a bilateral relative price q_i as

$$q_i = \frac{P_{us}}{P_i} \cdot E_{i/\$}. \quad (1)$$

This relative price is unitless and easy to interpret: A value of 2 means that the basket is twice as expensive in the United States as it is in country i .

To combine these bilateral relative prices into an aggregate measure for the United States, we use a weighted geometric mean where the weights vary over time and reflect each country's importance in U.S. trade. Specifically,

$$Q_t^g = (q_{1t})^{w_{1t}} \cdot (q_{2t})^{w_{2t}} \cdot \dots \cdot (q_{nt})^{w_{nt}}, \quad (2)$$

where w_{it} is the time-varying weight associated with the i th country.⁸ Two features of Q^g are worth

⁸Because Q^g is a limiting case of the more general CES function, section 5.1 below examines the sensitivity of our results to alternative parameterization.

noting. First, the level of the aggregate has meaning: a value of 1.5 means that U.S. prices are on average fifty percent above foreign prices and this value is not arbitrarily determined by the choice of base year. Second, the aggregate can change even if all bilateral relative prices are fixed.

An obvious alternative to Q^g is the commonly used chained aggregate, which is a weighted average of the growth rates of bilateral relative prices:

$$\frac{Q_t^c}{Q_{t-1}^c} = \prod_{i=1}^N \left(\frac{q_{it}}{q_{i,t-1}} \right)^{w_{it}}. \quad (3)$$

By convention, $Q_{t=base}^c$ is set equal to 100 in a given base period and the level of the index for all other periods is defined recursively. Chained aggregation has two important features to recommend it. First, the index is independent of the levels of its constituent q 's, implying that we do not have to choose a meaningful base period for them. Second, changes in the aggregate index only reflect changes in the underlying relative prices. That is, if these rates do not change over a given period, then the aggregate index will not change, even if the weights do. Thus, Q_t^c may be ideal for measuring the average change in the dollar's bilateral relative prices.⁹

Given these aggregation formulas, how can, in the aggregate, U.S. prices rise relative to foreign prices? Holding all else equal, there are four channels:

1. The nominal dollar exchange rate can appreciate.
2. U.S. prices can rise.
3. Foreign prices can fall.
4. The weight of relatively low-price foreign economies can increase.

The first three channels operate through their impact on the bilateral relative prices – the q 's – and they are fully captured in both the geometric and chained aggregates. However, the chained index does not attempt to capture the fourth channel whereas the geometric aggregate does so explicitly. Specifically, logarithmic differentiation of equations (2) and (3) with respect to time yields

$$\begin{aligned} d \ln Q_t^g &= \sum_i w_{it} \cdot d \ln (q_{it}) + \sum_i dw_{it} \cdot \ln (q_{it}) \\ d \ln Q_t^c &= \sum_i w_{it} \cdot d \ln (q_{it}), \end{aligned}$$

which implies that

$$d \ln Q_t^g - d \ln Q_t^c = \sum_i dw_{it} \cdot \ln (q_{it}).$$

⁹Appendix A.2 documents the properties of Q^g and Q^c using numerical examples.

Thus the difference in growth rates between the geometric and the chained aggregate is $\sum_i d\omega_{it} \cdot \ln(q_{it})$. This term captures the interaction between each period’s distribution of the level of bilateral relative prices and the evolution of the weights; if the weights are constant, then the two growth rates are identical.

3.2 Implementation

The previous discussion assumed, for expository convenience, the availability of data for the price levels of the foreign and domestic baskets. Thus the first step in implementing our measure is to obtain the bilateral relative prices—the q ’s. Data for bilateral relative prices are particularly difficult to obtain because they require comparability of products across countries.¹⁰ To this end we use the Penn World Tables, which offer data on purchasing power parties.¹¹

Greatly simplified, Penn collects data on spending and prices for products that are comparable across countries to estimate bilateral purchasing power parities. To avoid the calculations being sensitive to the choice of base country, Penn introduces the concept of “international dollars.” This strategy generates a system of simultaneous equations—the Geary-Khamis system—in which the PPP estimates depend on the international dollar and vice versa.¹² Specifically, given the international-dollar price of the j th product, π_j , the purchasing power parity for the i th country is

$$PPP_{\frac{i}{s}} = \frac{\sum_{j=1}^m P_j^i \cdot Y_j^i}{\sum_{j=1}^m \pi_j \cdot Y_j^i}, \quad i = 1, \dots, n, \quad (4)$$

where P_j^i is the price of the j th product in the i th country, Y_j^i is the amount produced of the j th product in the i th country, and the j index runs over the list of goods and services included in GDP.¹³ The numerator equals the nominal GDP of the i th country expressed in local-currency terms whereas the denominator is the value of i th country’s GDP expressed in international dollars.

¹⁰See Vachrís and Thomas (1999) for the importance of comparability.

¹¹Penn’s purchasing power parties have been used extensively in empirical analyses for the last three decades. See Heston, Summers and Aten (2006). For an introduction to the Penn World Tables, see Summers and Heston (1991) and Gulde and Schulze-Ghattas (1993). For the associated details, see Kravis, Heston, and Summers (1978, 1982). Other institutions also report purchasing power parities, but we use those from Penn because they include a relatively long time series. For recent efforts to improve the Penn World Tables, see Feenstra et al. (2005).

¹²Our presentation follows closely that of Gulde and Schulze-Ghattas (1993); see Kravis, Heston, Summers (1978, 1982) for additional details.

¹³Note that the weights (the Y s) are the same for the international dollar and for the prices of the j th country. We want to emphasize that these Y s are not physical measures of output but notional quantities; see Kravis, Heston, and Summers (1978, 1982) for details of this concept.

Given $PPP_{\$/}$, the international dollar price for the j th product is computed as

$$\pi_j = \sum_{i=1}^n \frac{P_j^i}{PPP_{\$/}^i} \cdot \left(\frac{Y_j^i}{\sum_{i=1}^n Y_j^i} \right), \quad j = 1, \dots, m, \quad (5)$$

where the first term is the price of the j th product in the i th country expressed in international dollars and the second term is the i th's country share in world output of the j th product.

The system given by (4) and (5) consists of $m + n$ equations, of which only $m + n - 1$ are linearly independent. To address the over-determined character of the system, Penn uses the United States as the numeraire country meaning that the international dollar has the same purchasing power over total U.S. GDP as the U.S. dollar. Thus the average U.S. price relative to the average price of the i th country can be estimated as the market exchange rate divided by Penn's PPP:

$$q_i = \frac{E_{\$/}^i}{PPP_{\$/}^i}. \quad (6)$$

There are several drawbacks to the Penn data for studying the open-economy implications of movements of U.S. international relative prices. First, the data are released with long delays: the most recent release (release 6.2 in 2006) has data ending in 2004.¹⁴ Second, the data are annual. To address these two limitations, the paper develops a method to extend Penn's annual parities and to estimate the associated quarterly observations.¹⁵ Finally, the data are subject to errors and section 5 examines the implications of these errors for our measure of U.S. international relative prices.

With these considerations in mind, the top panels of figure 2 show the evolution of the levels of bilateral relative prices for selected countries.¹⁶ Among the industrial countries (left panel), U.S. prices are highest relative to Portugal and lowest relative to Switzerland with most measures near or a little below one. As shown to the right, among emerging economies, there is a good deal more dispersion with relative prices ranging between 1.5 and 6.

For aggregation we use the same trade weights as those in the Federal Reserve's Broad Real Dollar index (shown in lower left panel). Note that between 1980 and 1990 the total weight of emerging economies held steady near 25 percent, but since 1990 it has doubled to near 50 percent, reflecting

¹⁴An earlier version of this paper used data from the 6.1 release and the results are quite comparable to the ones reported here. Nevertheless, recent data from the World Bank's International Comparison Program (December 2007) point to data revisions for recent years. Section 5.2 below examines the implications of measurement errors. Further, in personal communication, Alan Heston informed us that there will be a new release of the PWT. Once the new data are available, we will update our calculations.

¹⁵Appendix A.1 documents these methods; for estimating quarterly parities we impose the constraint that the average of quarterly parities for a given year must be equal to the annual Penn parity for that year.

¹⁶The countries for each group were selected so as to encompass all the bilateral relative prices in our sample.

Bilateral Relative Price Levels

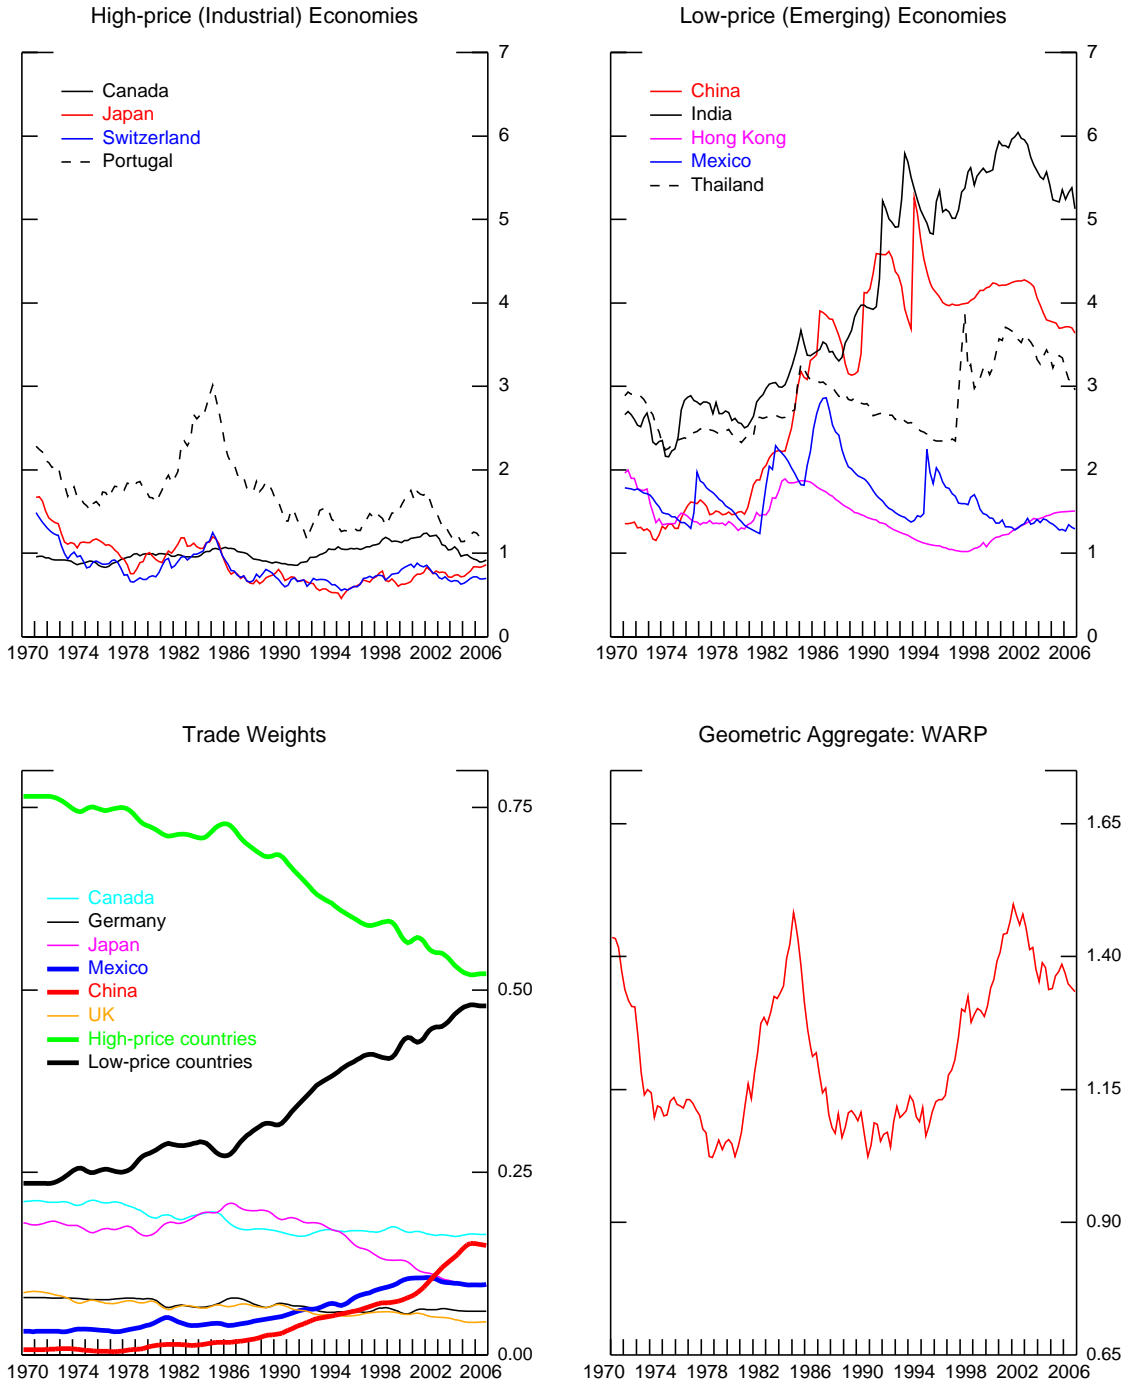


Figure 2: U.S. International Relative Prices: Bilateral and Aggregate

rising weights for China and Mexico. The weight for industrial countries has declined, with Japan's weight declining the most. The weighted average of the 34 bilateral relative prices is constructed using equation (2) and shown in the lower right panel. The aggregate of U.S. international relative prices shows an upward trend since the end of the Bretton-Woods period. Indeed, by this measure, U.S. prices are roughly 40 percent above those of its trading partners.

Figure 3 compares the evolution of this measure to the real effective exchange rates from the Federal Reserve and the IMF, rescaled by their own 1971-1991 sample means. The three measures move in near lockstep between 1971 and 1986. As such, neither the choice of aggregation method nor the measure of bilateral price has a noticeable effect on the aggregate measure of U.S. international relative prices through 1986. Since then, however, the aggregates tend to diverge. Specifically, the WARP shows a sustained increase and by 2002 it reaches the same value it had in 1985. In contrast, the other measures remain well below their 1985 peaks. This more recent divergence of U.S. international relative prices might be of interest in assessing the likelihood of a dollar depreciation large enough to address the U.S. external imbalance. Specifically, if one were to apply the 1985-1987 dollar depreciation to the 2006 values of the aggregates based on bilateral price indexes, then these aggregates would fall to levels not recorded in history. In contrast, applying the same depreciation to WARP would bring it to its 1986 value and, by this historical standard, such a depreciation would be consistent with previous experience.

4 Explaining the Rise in U.S. International Relative Prices

We now look at why our WARP has risen much more than the other measures since the late 1980s. To ease the exposition, we abstract from differences involving country coverage and weighting scheme to focus on the measurement of bilateral prices and aggregation methods.

4.1 Aggregation Methods and Price Measures

WARP differs from existing measures in both the choice of aggregation formula and the measure of bilateral prices, raising the question of which of these two factors explains the different trends in the aggregates. To address this question, we construct similar aggregates to those reported by other institutions where their bilateral relative prices are replaced with ours. This strategy ensures that any difference can be interpreted as due to the choice of aggregation method.

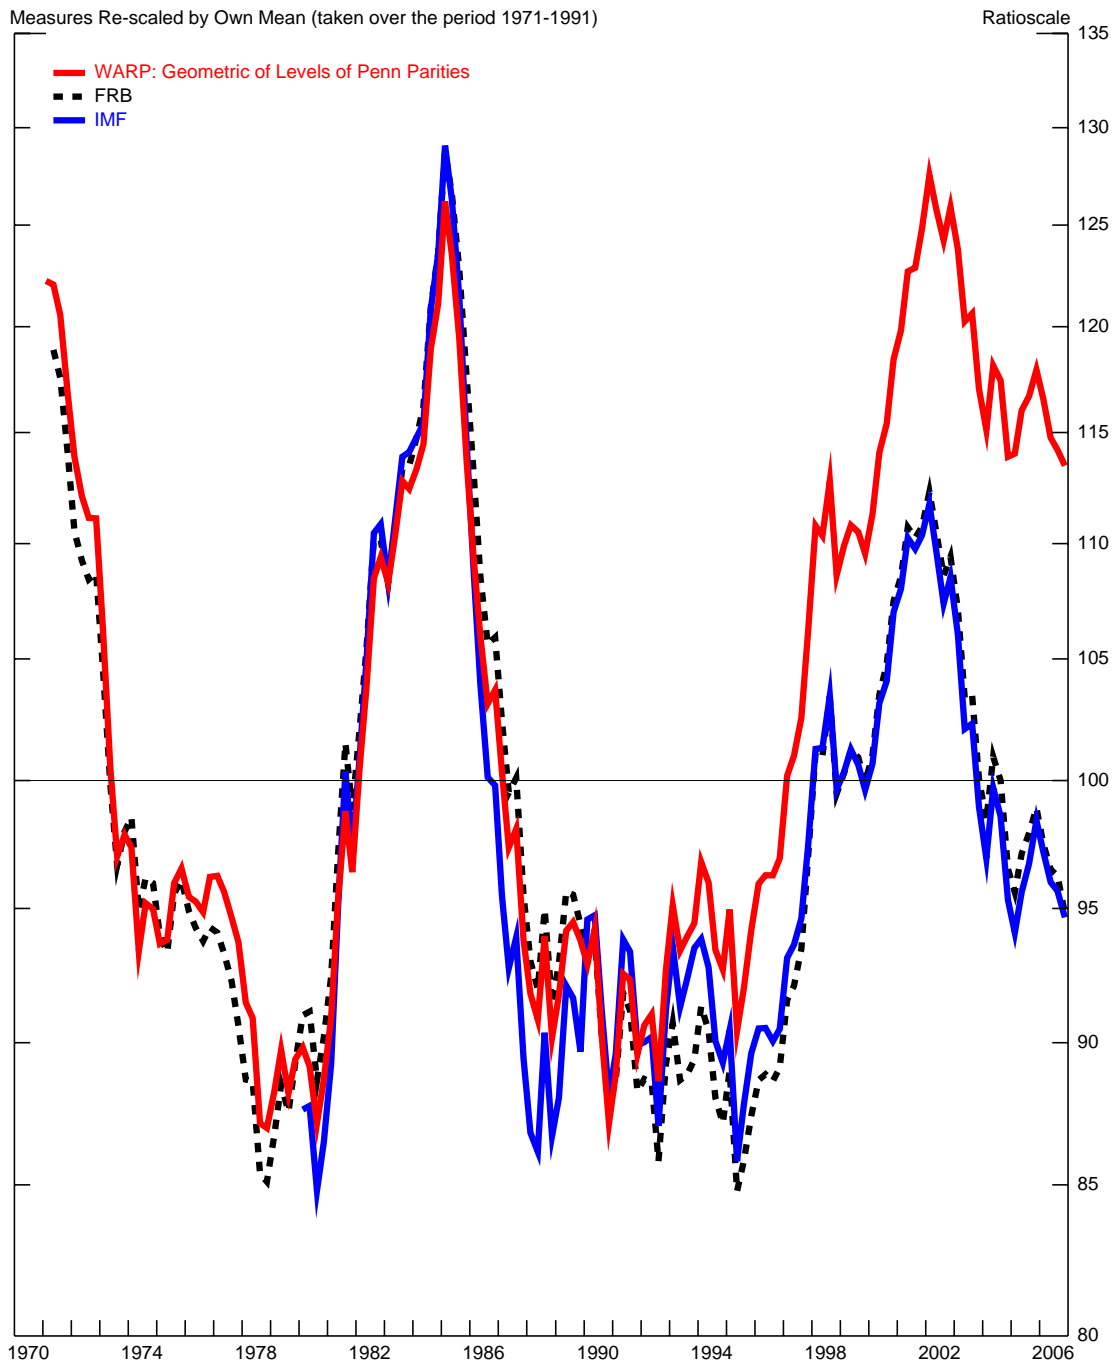


Figure 3: Alternative Measures of International Relative Prices

The Federal Reserve reports chained aggregates of bilateral CPI-adjusted exchange-rate indexes:

$$\frac{Q_t^c}{Q_{t-1}^c} = \prod_{i=1}^N \left(\frac{r_{\frac{i}{s},t}}{r_{\frac{i}{s},t-1}} \right)^{w_{it}}, \quad (7)$$

where

$$r_{it} = \left(\frac{CPI_{us,t}}{CPI_{i,t}} \right) \cdot \left(\frac{E_{\frac{i}{s},t}}{E_{\frac{i}{s},t_0}} \right),$$

t_0 represents the base period, CPI_{us} is the U.S. consumer price index, and CPI_i is the consumer price index for the i th country. The IMF reports a fixed-weight geometric aggregate of bilateral CPI-adjusted exchange rates:

$$Q_t^{gI} = \prod_{i=1}^N \left(r_{\frac{i}{s},t} \right)^{w_i}. \quad (8)$$

There are several differences between q_{it} and r_{it} that are potentially relevant for explaining differences between WARP and existing measures of real effective exchange rates. First, q_{it} measures the level of bilateral relative prices whereas r_{it} measures the percent change in bilateral relative prices. Second, the basket used for q_{it} refers to GDP items and thus includes consumption, investment, government purchases, and exports. The basket used for r_{it} is limited to consumption items both from domestic and foreign sources. Finally, the baskets embodied in q_{it} are the same for U.S. and foreign prices whereas the baskets embodied in r_{it} are not the same for U.S. and foreign prices.

To examine whether the upward trend in our WARP is due to differences between q_{it} and r_{it} , we construct the same aggregates reported by other institutions while using q_{it} instead of r_{it} . Following the FRB's methodology, we construct a chained aggregate of bilateral relative prices, Q_t^c , substituting the Penn parities (the $q_{i's}$) into equation (7). We also report the geometric average of *indexes* of bilateral relative prices, similar to the IMF's methodology:

$$Q_t^{gI} = \prod_{i=1}^N \left(\frac{q_{it}}{q_{it_0}} \right)^{\omega_i},$$

where, following the IMF, we set q_{it_0} as the mean of the values of q_{it} in 2000 and ω_i as the sample mean from 1989 to 1991.

Figure 4 shows the evolution of the real effective exchange rates from the Federal Reserve and the IMF along with Q_t^g , Q_t^c , and Q_t^{gI} ; for comparison purposes, we rescale these measures by their own 1971-1991 sample means. The results indicate that aggregates based on Penn's bilateral relative price *indexes*, Q_t^c and Q_t^{gI} , show a downward trend meaning that the upward trend in WARP is not due to differences between q_{it} and r_{it} but, rather, to the choice of aggregation method.

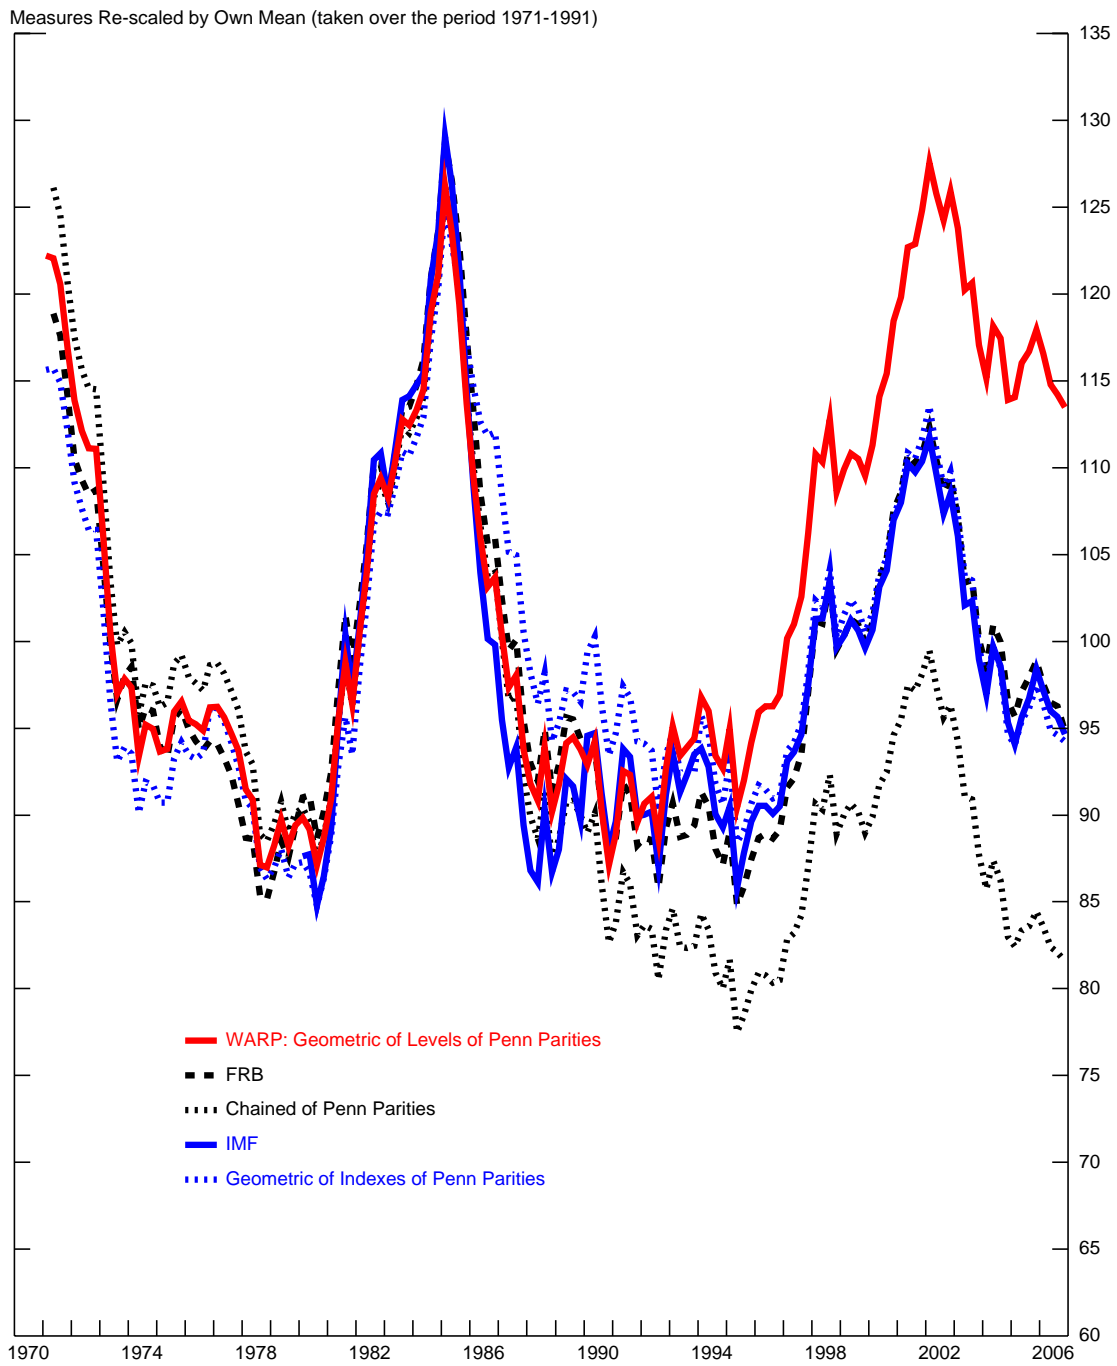


Figure 4: U.S. International Relative Price – Sensitivity to Price Data and Aggregation Method

4.2 Currency Baskets and Weighting Schemes

We now ask what factors in our aggregation method are responsible for the upward trend in Q^g . Is this trend due to the composition of the currency basket or to our weighting scheme? To address these questions, figure 5 reports separate geometric sub-aggregates for industrial countries and for emerging economies; the w'_{it} s for each group are renormalized to add up to one. The thick blue line is the aggregate of U.S. prices relative to other industrial countries alone. It has been trending down slowly, and it indicates that in 2006 U.S. prices were on average 10 percent below those in other industrial countries. The thick black line plots U.S. prices relative to the prices of emerging economies alone. It has been trending up sharply. These calculations suggest that a key factor accounting for the upward trend in our WARP, the thick red line, is the shift in U.S. trade patterns. Specifically, within the overall aggregate, trade has shifted away from the relatively high-price industrial countries toward the lower-price emerging economies, which tends to raise the overall measure of U.S. prices relative to our trading partners. Within the emerging economies sub-aggregate (the black line), trade has shifted toward the lowest-price economies, such as China; this shift tends to raise U.S. prices relative to the group.

To illustrate the importance of the increased weight of the low-price economies, we construct a counterfactual where we ask what would have happened if the weights after 1991 were fixed at their 1991 values. As shown by the dashed red line, in this world, our measure would have U.S. prices only about 10 percent above those of our trading partners—roughly unchanged since 1975 and near the 30-year average. Further, this fixed-weight aggregate has a downward trend with a historical peak in 1985, quite similar to the pattern of the standard aggregates shown earlier. The key question, however, is whose weights exert the strongest influence. To address that question, the figure reports the fixed-weight aggregate for the industrial countries. Appearing as the dashed blue line, this aggregate exhibits a downward trend with a historical peak in 1985, similar to the case of variable weights and to the associated aggregates reported by other institutions. This result implies that fixing industrial-country weights does not change the evolution of that sub-aggregate. In contrast, fixing the weights of each emerging economy induces a *downward* trend in the associated sub-aggregate (dashed black line), which leaves the historical peak back in 1985, unlike its variable-weight version which peaks in 2003. Therefore, the upward trend in our measure of U.S. international relative prices is due to the increased weight of the low-price economies in the U.S. basket.




Figure 5: U.S. International Relative Price – Sensitivity to Currency Basket and Weighting Scheme

5 Sensitivity Analysis

5.1 Sensitivity to Parameterization

A well known result is that our Q^g is a particular case of the CES function

$$Q_t^{ces} = \left[\sum_{i=1}^n w_{it} \cdot (q_{it})^{\sigma-1} \right]^{\frac{1}{\sigma-1}},$$

where σ is the elasticity of substitution among purchases of foreign products and $Q^g = \lim_{\sigma \rightarrow 1} Q_t^{ces}$. In the absence of econometric evidence supporting $\sigma = 1$, a relevant question to pose is how sensitive is the upward trend in U.S. international relative prices to alternative values of σ .¹⁷

Figure 6 shows the sensitivity of Q_t^{ces} to values of σ ranging from high substitutability ($\sigma = 2.5$) to near complementarity ($\sigma = 0.05$). The calculations reveal three findings. First, there is a direct association between the value of σ and the slope of the trend of U.S. relative prices. Second, the 2006 level of WARP is sensitive to extreme values of σ (2.5 and 0.05); using less extreme values of σ (1.1 and 0.9) yields values of Q_t^{ces} quite close to the values taken by Q^g . Finally, if one interprets the large swings in U.S. bilateral trade shares as suggesting high substitutability among foreign products ($\sigma > 2$), then the upward trend in Q^g understates the extent to which U.S. international relative prices have been increasing.

Overall, we interpret these results as suggesting that the upward trend in U.S. international relative prices associated with WARP is not due to our reliance on a unitary elasticity of substitution among foreign products. This conclusion, however, abstracts from the quality of the data of relative prices, an issue that we examine next.

5.2 Sensitivity to Measurement Errors

We now ask how sensitive is the upward trend of WARP to measurement errors. Other things equal, a lower estimate of $PPP_{i/s}$ raises q_i , our measure of U.S. prices relative to country i . Thus a relevant question is whether we are over-estimating q_i , and hence WARP, because Penn is underestimating $PPP_{i/s}$ for emerging economies. We consider three approaches to address this question: examining alternative measures of purchasing power parity for China; adjusting our measures of relative prices by imputing correction factors larger than those of Kravis and Lipsey (1990); and comparing the WARP to calculations based on the price data for the Big Mac.

¹⁷ Appendix A.3 shows the derivation of Q^{ces} .

Sensitivity to Degree of Differentiation




Figure 6: U.S. International Relative Prices – CES Aggregates with Alternative Values of σ

5.2.1 Alternative Estimates of PPP for China

A focus on China's purchasing power parity can be motivated in two ways. First, the weight for China has experienced the largest increase and it now has the second largest value in our weighting scheme. Second, the price data for China are of questionable reliability. Figure 7 compares the estimates for China's $PPP_{i/s}$ from Penn (solid dark-blue line) to the IMF's estimates from seven recent vintages.¹⁸ Prior to 1994, Penn's estimate is never more than eight percent below the IMF estimates. For all the post-1994 period, Penn's estimate is at least as large as any of the estimates from the IMF. Thus there does not seem to be a systematic undervaluation of Penn's parities relative to those of the IMF. The one estimate we could find that is above that of Penn is that of the OECD.¹⁹ For 2004, the OECD estimate for China is 2.3, compared to Penn's estimate of 2.1. Thus the OECD estimate is roughly 10 percent above the Penn estimate.

We do not interpret this scant evidence as suggesting that Penn's estimates are relatively error free but that, perhaps, comparing inaccurate measures across institutions is not informative. Thus we examine below the implications of imputing large measurement errors to the relative prices of emerging economies. Indeed, we find that even if the relative prices for all emerging economies were 20 percent below what Penn estimates, this is still not enough to overturn the basic upward trend of WARP.

5.2.2 Imputation of Correction Factors

We now impute measurement errors to the relative prices of emerging economies to examine the sensitivity of WARP to such mismeasurement. Specifically, we denote \tilde{q}_{it} as the error-free but unobserved bilateral relative price and postulate that $\tilde{q}_{it} = (q_{it})^\theta$ where θ is the imputed correction factor. We could impute the value of θ using either the 13 percent estimated by Kravis and Lipsey (1990) or the 10 percent wedge implied by the OECD estimate for China. To encompass these sources and to allow for even larger errors, we apply $\theta = 3/4$ to the relative prices of China, of Latin American countries, of all emerging economies excluding both China and Latin America countries, and of all emerging economies. Note that the magnitude of the error is directly related to the value of q . If $q = 2$ and $\theta = 3/4$, then $\tilde{q} = 1.68$ implying an error of 19 percent.

The top panel of figure 8 shows how U.S. prices relative to those of emerging economies respond to the imputed correction factor. We find that if $\theta = 3/4$ is applied to all developing countries, then

¹⁸The data come from the *World Economic Outlook*. The IMF's calculation starts with the PPP exchange rate in year 2000 from the World Bank; this value is then extended forwards and backwards by the growth in relative GDP deflators. See <http://www.imf.org/external/pubs/ft/weo/faq.htm#q21>

¹⁹See *OECD Economic Surveys China* Volume 2005/13 September 2005, page 9.




Figure 7: IMF PPP \$ Rates for China – Sensitivity to Data Vintage

reliance on Penn data overstates U.S. relative prices vis-a-vis this countries by about 20 percent. The bottom panel of figure 8 shows how these measurement errors affect WARP.²⁰ When the correction is applied only to China, the upward trend in U.S. international relative prices remains in place. Applying the correction factor to either Latin America or other emerging economies (except China and Latin American countries), leaves the trend rate of WARP largely unchanged. Finally, applying the correction factor to *all* of the emerging economies dampens the upward trend of WARP but by no means eliminates it.

5.2.3 Consistency with Big Mac Prices

We now evaluate whether the results from using Penn’s parities are unique by comparing them to the prices of McDonald’s Big Mac reported by *The Economist*. This alternative is of interest because *The Economist* reports the absolute dollar-price levels for the Big Mac. Figure 9 shows the cross-country dispersion of dollar prices for the Big Mac from 1986 to 2007 for 31 countries.²¹ The data reveal that the number of countries with prices below the U.S. price has increased markedly over the years. As for the range of prices, Switzerland tends to have the highest price whereas China generally has the lowest price.

Given these prices, we construct the U.S. bilateral relative price of a Big Mac, $q_{B,i}$, as

$$q_{B,i} = \frac{P_{B,us}}{P_{B,i}},$$

where $P_{B,us}$ is the dollar price of a Big Mac in the United States and $P_{B,i}$ is dollar price of a Big Mac in the i th country. Given $q_{B,i}$, the associated geometric aggregate is

$$Q_{Bt}^g = \prod_{i=1}^{N_B} (q_{B,i})^{\omega_{B,i}}, \quad (9)$$

where N_B is the number of countries included in the aggregate and $\omega_{B,i}$ is the trade weight for the i th country; we construct this aggregate for countries that are included in the Federal Reserve’s Broad Real Dollar index.

Because the list of countries reported by the *Economist* varies across time, we construct Q_{Bt}^g

²⁰ Assuming that the measurement error is concentrated in countries $i = n_0 + 1 \dots n$, the WARP is constructed as

$$\tilde{Q}_t^g = \left(\prod_{i=1}^{n_0} (q_{it})^{w_{it}} \right) \cdot \left(\prod_{i=n_0+1}^n (q_{it})^{\theta \cdot w_{it}} \right) = Q_t^g \cdot \left(\prod_{i=n_0+1}^n (q_{it})^{w_{it}} \right)^{\theta-1}.$$

²¹ The data for this section was collected by Jeffrey Traczynski. Note that for countries that adopted the euro, *The Economist* reports prices beginning in 1999.

