A Prism into the PPP Puzzles:
The Micro-foundations of Big Mac Real Exchange Rates

December 2005

David C. Parsley
Owen Graduate School
Vanderbilt University
David.Parsley@Vanderbilt.edu

Shang-Jin Wei
International Monetary Fund
Brookings Institution and NBER
Swei@imf.org

Abstract

The real exchange rate has been called the single most important price in an economy, yet its behavior exhibits several puzzles. In this paper, we use Big Mac prices as a unique prism to study the movement of real exchange rates. Part of our innovation is to match these prices to the prices of individual ingredients. There are a number of advantages associated with the approach. First, unlike the CPI real exchange rate, we can measure the Big Mac real exchange rate in *levels* in an economically meaningful way. Second, unlike the CPI real exchange rate, for which the attribution to tradable and non-tradable components involves assumptions on the weights and the functional form, we know (almost) the *exact* composition of a Big Mac, and can estimate the tradable and non-tradable components relatively precisely. Third, we can study the dynamics of the real exchange rate in a setting free of several biases inherent in examinations of aggregate CPI based real exchange rates. These biases – the product-aggregation bias (Imbs, Mumtaz, Ravn, and Rey, 2005), the temporal aggregation bias (Taylor, 2001), and the bias generated by non-compatible consumption baskets across countries – are candidate explanations for the puzzlingly slow mean reversion alluded to by Rogoff (1996). Finally, we show that Engel's result that deviations from the law of one price are sole explanation for real exchange rate movements does not hold generally. We offer some evidence that departure from the Engel effect can be systematically linked to economic factors.

Key Words: Real exchange rates; TAR models; Real exchange rate decompositions; sigma-convergence; beta-convergence.

JEL Classification Codes: F31

This is a substantially revised version of NBER Working Paper #10074. We would like to thank Robert Cumby, Peter Pedroni, Bruce Reynolds, Kenneth Rogoff, Lucio Sarno, Moto Shintani, the Editor, two anonymous referees, and seminar participants at the ASSA 2005 meetings, the Bank of France, the European Central Bank, Louisiana State University, the University of Chicago, the University of Virginia, UC San Diego, Vanderbilt, and the IMF for useful comments. The views presented in this paper are those of the authors and do not necessarily reflect the views of any institution with which they are affiliated.

"In most economies, the exchange rate is the single most important relative price, one that potentially feeds back into a large range of transactions."

Obstfeld and Rogoff (2000).

1. General advertisement

The real exchange rate's central importance in an economy has long been recognized (see, for example, Milton Friedman, 1953; and more recently, Maurice Obstfeld and Kenneth Rogoff, 2000). According to a recent study by Lucio Sarno and Mark P. Taylor (2002), the concept of Purchasing power parity (PPP) has been propounded as a theory of real exchange rate behavior at least since the 1500's. Indeed, Jeffrey Frankel and Andrew Rose (1996), describe it as one of the most important theoretical concepts in international economics. Despite considerable academic attention, two key aspects of real exchange rate movements evade convincing explanation.

First, the estimated speed of mean reversion seems too slow (or, equivalently, the deviations from purchasing power parity seem too persistent). After surveying a long list of papers on the subject, Rogoff (1996) observed that a "remarkable consensus view" exists on the estimated half-life of deviations from PPP; which he concluded is on the order of three to five years. But this seems too long, based on economic theories with a plausible size of arbitrage costs (e.g., Chari, Kehoe, and McGratten, 2002).

Second, the role of the relative price of non-tradable goods across countries (e.g., through the Balassa-Samuelson effect) in accounting for medium- to long-run movements in real exchange rates seems too small.[1] A recent much-cited paper by Engel (1999) seriously undermines the conventional view; he finds that nearly 100% of real exchange rate variation is explained by deviations from the law of one price in tradables, and virtually none by differentials in the relative price of non-tradables across countries.

Four different explanations have been suggested for the persistence puzzle concerning CPI-based real exchange rates. First, the CPI baskets across different countries are not identical, and the constituents of each country's basket can change at a different pace over time[2]. Clearly, this diminishes the information content of observed differences in the prices of these consumption baskets. Hence it is not surprising that apparent price differences do not quickly disappear. Second, the common linear estimation of the persistence parameter may be

[1] Froot and Rogoff (1995) provide a comprehensive survey of studies investigating the long-run determinants of purchasing power parity.

[2] One component of the CPI basket is food. As an example, the French basket may have a relatively heavy weight on cheese, which the Chinese may not care much about; while the Chinese basket may contain lots of tofu, which may be a small portion of French consumption. It is not particularly meaningful to speak of arbitrage between cheese prices in France and tofu prices in China.

misspecified. Several authors (e.g., Obstfeld and Taylor, 1997; Taylor, 2001; Sarno and Taylor, 2002; O'Connell and Wei 2002) argue that arbitrage costs dictate a non-linear specification. Specifically, arbitrage costs lead to a band of no-arbitrage, within which the real exchange rate can behave as a random walk (i.e., the half-life can be infinite). Once outside the no-arbitrage zone, the force of arbitrage may drive the real exchange rate back at a relatively fast speed. In empirical work, once this non-linearity is taken into account, the real exchange rate is typically found to be much less persistent (the estimated half-life usually falls in a range of 1-2 years).

The third and fourth explanations of slow persistence have to do with two aggregation biases. Taylor (2001) shows that when price or nominal exchange rate data are averages of data collected at different points in time, the persistence of the real exchange rate may be over-estimated. This is called a time-aggregation bias.[3] Imbs et al (2005) argue that the estimated persistence of a composite (such as a CPI-based real exchange rate) is upwardly biased relative to the true average of the levels of persistence of the individual components of the composite. This is known as a product-aggregation bias. These four explanations are not mutually exclusive; collectively, they demonstrate the challenges that cloud interpretation in studies of CPI-based real exchange rates.

In this paper we adopt an alternative approach to study the movement of real exchange rates by using information in a panel of Big Macs. Like aggregate CPI-based real exchange rates, the Big Mac is a 'basket' comprised of both traded and non-traded components. Before extrapolating lessons from Big Macs, we need to demonstrate that the resulting time-series of 'Big Mac' and CPI-based real exchange rates are sufficiently similar. In particular, we show that statistically, Big Mac real exchange rates are highly correlated with CPI-based real exchange rates (both in levels and in first differences). This suggests that McDonalds' firm-specific pricing policies are not of first-order importance for our results. To strengthen the case that lessons from Big Macs have more general implications, we focus our analysis on that subset of bilateral real exchange rates most highly correlated with CPI based real exchange rates.

One of our key innovations is to match Big Mac prices to the prices of individual ingredients, e.g., ground beef, bread, lettuce, labor, etc., and to design the thought experiments in such a way as to mitigate the problems discussed above that have confounded much of the existing literature. More specifically, our approach offers five distinct advantages.

First, unlike the consumption baskets that go into the CPI calculation, which may not be comparable across countries, the Big Mac composite is nearly identical in all countries and

[3] Note that this is not the same as sampling at a frequency lower than the half-life. The 'problem' of sampling frequency may or may not produce a bias in mean reversion estimates (see Taylor 2001).

across time periods. In fact, due to McDonalds' global advertising strategy, millions of people world-wide can actually sing the exact combination of its ingredients.[4]

Second, unlike the CPI-based real exchange rate, we can measure the Big Mac real exchange rate in *levels* in an economically meaningful way. Thus, e.g., in studying real exchange rate persistence, we avoid the assumption that parity held in some base period.

Third, apportioning the CPI real exchange rate into tradable and non-tradable components requires assumptions concerning weights and functional form that may vary by country. In contrast, we (almost) know the *exact* composition of a Big Mac and can estimate its tradable and non-tradable components relatively precisely.

Fourth, we can study the dynamics of the real exchange rate in a setting that is free of the product-aggregation bias or the temporal aggregation bias. To highlight the role of these and other biases that affect persistence estimation, we implement both linear, and non-linear, convergence specifications.

Finally, we use our estimated tradable/non-tradable shares and the large cross-section of country-pairs in our data set to re-examine Engel's (1999) proposition concerning the role of traded goods prices in real exchange rate movements. This data set provides us with an opportunity to explore whether departures from his result can be systematically explained as a function of country pair characteristics.

To summarize, this paper uses Big Mac real exchange rates to study these two well-known puzzles. This allows us to avoid several serious drawbacks shared by studies employing CPI-based real exchange rates. To be clear, we do not propose any new theory or new econometric method; we do however, reach several conclusions that would not be possible using CPI-based real exchange rates.

Aside from the literature on real exchange rates referenced above, there is a collection of recent papers that makes use of the Big Mac prices reported in the *Economist* magazine, including Click (1996), Cumby (1997), Lutz (2001), Ong (1997), and Pakko and Pollard (1996). They have typically showed that relative Big Mac prices between countries resemble CPI-based real exchange rates in many ways. However, none of these papers match Big Mac prices with the prices of its underlying ingredients. We use these matched data sets, which are part of our innovation, to decompose Big Mac real exchange rates into tradable and non-tradable components, and to address a range of questions beyond those typically studied using consumer price index based real exchange rates.

[4] We refer to the well known jingle "two all beef patties, special sauce, lettuce, cheese, pickles, onions, on a sesame seed bun". There are however, some small differences in Big Macs around the globe. For example, in India (not in our data set) no beef products are sold, and in Israel (in our data set) the beef is kosher.

The rest of the paper is organized as follows. In the next section, we provide a more detailed description of the data sets, including their sources and coverage across time, countries, and items. Section 3 contains the core of our statistical analysis, which is presented in four steps. First, we establish a connection between CPI-based and Big Mac-based real exchange rates. Second, we apportion the price of a big Mac into its constituent parts. Third, we study the dynamics of the Big Mac real exchange rates; in particular, by comparing its 'aggregate' convergence speed with those of its ingredients. Fourth, we examine the fraction of Big Mac price disparities attributable to deviations from the law of one price and we explore factors that may explain variation in this fraction across countries and over time. The final section offers some concluding remarks.

2. Data: sources and ingredients

Prices

The primary data we employ in this study are price observations for the Big Mac and for its various ingredients, in 34 countries over 13 years (1990-2002). The local currency data for Big Mac prices was obtained from various year-end editions of the *Economist* magazine. The countries reported in each edition of the "Big Mac Index" varies, hence we exclude countries with fewer than 5 years of data; the median number of years for the resulting sample is 11 years.

The second data set covers city specific local-currency prices of various ingredients of the Big Mac – ground beef, bread, labor cost, etc. – in the same set of countries and years, compiled by the *Economist Intelligence Unit (EIU)*. While both data sets refer to prices observed at a point in time (i.e., not temporal averages), the *EIU* data become available around October, while the *Economist* Big Mac data are available anytime between early December and early January. Consequently, we exclude periods of high inflation (described below) in order to minimize the impact of any timing discrepancies. Table 1 lists the countries reported in the *Economist*, and the corresponding cities reported in the *EIU* data set. A sense of the global distribution of countries is also highlighted in the table.

The *EIU* data comes from the *Worldwide Cost of Living Survey*, which is designed for use by human resource managers in implementing compensation policies. The *EIU* official description is at http://eiu.e-numerate.com/asp/wcol_HelpWhatIsWCOL.asp. Some of the goods in the *EIU* data set appear twice – differing by the type of establishment where the price was recorded. When there was a choice between two prices, we selected the lower price; these generally are supermarket prices. In particular, we selected local currency price data on the following five traded inputs: ground beef, cheese, lettuce, onions, and bread. We also include three nontraded inputs: hourly labor costs, rent (proxied by rent for a two-bedroom unfurnished moderate apartment), and electricity charges.

To ensure that our subsequent results are not driven by some peculiarities of the data sets, we undertake some basic "data cleaning." First, we exclude "high inflation episodes" from our analysis since the potential importance of timing mismatches is greater in these cases. The specific episodes we exclude are Argentina (1990-91), Brazil (1990-94), Mexico (1990-92), and Poland (1990-94). Second, we visually checked the data for possible coding errors via scatter plots. More concretely, we also looked for unreasonably large fluctuations in local currency prices, or price changes greater than 60%, which were subsequently reversed in the next period. We took the ten instances (lettuce (7), onions (2), and rent (1)) where this occurred in our data set to be coding mistakes and used the average (t-1, t+1) value instead. We have experimented with other cut-offs for coding errors, and found the results not very sensitive to the choice of the cut-off points.

Other variables

In addition to the price data, we use data on tariffs, sales and value added tax rates. The first source of tariffs is simple mean tariff rates, from Table 6.6 of the World Bank publication *World Development Indicators 2001*. For each country the tariff data are available for two years – once in the early 1990s and once for the late 1990s. We use the first reported value in our bilateral tariff rate calculations for the years 1990-95, and the most recent value for the years 1996-2002. Sales tax and VAT rates were collected from primary sources. For Europe, the European Commission publication: "VAT Rates Applied in the Member States of the European Community" (2002), was the main reference. For the remaining countries and cities, data were obtained from web searches, emails, or phone calls directly to national (and state) tax authorities.

3. Digesting the Big Mac

This section contains the core of our statistical analysis. We proceed in four steps. First, we check the connection between CPI-based and the Big Mac-based real exchange rates. We conclude that Big Mac and CPI-based real exchange rates share important time series and cross-sectional features. Second, we take advantage of the simplicity of the Big Mac structure to link its price to the costs of its underlying ingredients and to allocate traded and non-traded input shares of the Big Mac aggregate. Third, we examine the speed of convergence to the law of one price for the Big Mac real exchange rate and compare it with those of its ingredients. We examine convergence in the cross-sectional dispersion of prices as well as mean reversion in real exchange rates – employing both non-linear, and linear specifications. Fourth, we focus on the large cross-sectional dimension of bilateral real exchange rates. Specifically, we reexamine the Engel (1999) puzzle, with an emphasis on trying to identify factors that may systematically affect the importance of deviations from the law of one price in explaining observed real exchange rate movements.

3.a. The Big Mac versus CPI-based real exchange rates

We begin by comparing the Big Mac, and the more standard, CPI-based real exchange rates. The idea is to see if Big Mac real exchange rates are informative about CPI-based real exchange rates, or alternatively, are too unique and narrow to be useful. Figure 1 shows that Big Mac real exchange rates are typically highly correlated with aggregate real exchange rates – both in levels, and in first differences. Note that if Big Mac and CPI real exchange rates both are integrated of order one (which is at the heart of much applied work on mean reversion), then correlation in levels could be spurious. For this reason we also examine the correlation after differencing. Thus, high correlation in first differences is a more stringent requirement.

The overall impression from the figure is that there are generally high positive correlations between aggregate and Big Mac real exchange rates in both levels and first differences. Nonetheless, for this study we make an effort to err on the conservative side and hence restrict our attention to only those bilateral cases where <u>both</u> correlation coefficients are greater than 0.65. In our sample, 61% (=343) of the 561 possible real exchange rates meet these two criteria simultaneously (the percentages for each of the criteria separately are: 74% in levels; and 80% in 1st differences). To convey an idea of what the restriction implies for the resulting sample, Table 2 presents the correlation coefficients for the twenty-one bilateral U.S. dollar real exchange rates that satisfy the restrictions. As is evident from the averages, the result is a sample of Big Mac real exchange rates that are very highly positively correlated with the traditional CPI based measures of the real exchange rate.

As an additional check, we test for cointegration between the underlying Big Mac price levels and the Consumer Price Indexes themselves.[5] We implemented the seven cointegration tests discussed by Pedroni (1999) on our panels of Big Mac and CPI data.[6] The tests (both parametric and non-parametric) are generalizations of conventional (i.e., non-panel) time series tests and have no cointegration as the null hypothesis, and thus have greater power than standard tests. The first four tests impose equal (but strictly less than one) cross-sectional autoregressive coefficients under the alternative hypothesis, while the remaining three tests allow country-specific heterogeneity in the long-run cointegrating relationship. All seven tests allow individual heterogeneity in short run dynamics. Six of the seven tests reject the null

[5] If Big Mac and CPI based real exchange rates are both stationary, they would be trivially cointegrated.

[6] Since each test has different relative strengths depending on the underlying data-generating process, we report them all. Due to missing values, the tests were performed on a reduced set (25) of the original 34 countries having a minimum of nine consecutive years of price data.

hypothesis of no cointegration at the 1-percent level or above, with the remaining test's significance level just shy of the 1-percent level.[7] We conclude that there is a stable relationship among the Big Mac and Consumer Price series that we examine.

Despite these findings, it is possible that the high real exchange rate correlations are driven entirely by high *nominal* exchange rate correlations. That is, high correlation may simply reflect the fact that under a floating exchange rate regime most real exchange rate variation is due to nominal exchange rate variation. This would make inferences from Big Mac to CPI-based real exchange rates more tenuous. We can investigate this possibility directly, by focusing on the correlation between Big Mac and CPI-based exchange rates for countries with low nominal exchange rate variability. If the correlation for this low volatility subset is substantially smaller than that for the rest of our sample, the inferences we could draw concerning aggregate real exchange rate behavior would be tentative.

To examine this we proceed in two steps. First, we computed the volatility of each bilateral nominal exchange rate as the standard deviation over 1990-2002 of first differences in the log of the series. Next, for the 'low volatility' group, i.e., those falling into the lowest quartile of the distribution, we compare the correlation of CPI-based real and Big Mac exchange rates with the rest. Since this correlation (0.693) is, in fact, greater than that for the remaining cases (0.626) no statistical test is necessary. Hence it is not the case that the correlations between Big Mac and CPI-based real exchange rates are driven solely by high nominal exchange rate volatility.

For completeness, we present the Big Mac – CPI real exchange rate correlations (as in Figure 1) for the low nominal exchange rate volatility sub-sample in Figure 2. These correlations are again higher, i.e., more than 80 percent of real exchange rates in the low nominal volatility sample have correlation coefficients (levels and first differences) above 0.8, compared with between 58 percent and 68 percent for the full sample. Since Big Mac and CPI real exchange rates are highly correlated in levels and in first differences, with or without high nominal exchange rate variability, we conclude that inferences from Big Mac real exchange rates should shed light on more traditional, highly aggregated CPI-based real exchange rates.

Finally, we cite supporting evidence in Cumby (1997), who demonstrates that deviations from relative Big Mac parity provide useful information for forecasting exchange rates. In particular, Cumby finds that, conditional on currency-specific constants, a 10 percent

[7] Using Pedroni's notation, the estimated test statistics for the log Big Mac price and Log CPI are (statistics for levels regressions are in parentheses): (1) panel v-statistic -5.40 (-4.76), (2) panel ρ -statistic = 3.29 (2.30), (3) panel Phillips-Perron statistic = 4.51 (5.97), (4) panel adf-statistic = 4.75 (7.79), (5) group ρ -statistic = 2.83 (-17.14), (6) group Phillips-Perron statistic = 52.7 (27.9), (7) group adf-statistic = 4.69 (4339). We thank Peter Pedroni for providing his RATS program to compute these test statistics.

undervaluation according the Big Mac real exchange rate in one year is associated with a 3.8 percent appreciation over the following year.

Taken together, these pieces of information suggest that the behavior of the Big Mac real exchange rate is very similar to that of the CPI based real exchange rate, and that it is not driven by the peculiarities of McDonalds' corporate pricing strategy.

3.b. Reverse-engineering the recipe

Our next task is to relate the price of a Big Mac to the cost of its ingredients. Suppose there are exactly n inputs; and the production function is Leontief:

$$1 \text{ Big Mac} = \min \{x_1, x_2, ..., x_n\} \tag{1}$$

Let $P_{k,t}^{\text{Big Mac}}$ be the price of a Big Mac in country k at time t, and $P_{k,j,t}$ be the price of input j in country k at time t. Then,

$$P_{k,t}^{\text{Big Mac}} = \sum_j P_{k,j,t}\, x_j \tag{2}$$

To be precise, here we use the term "input" broadly to also include an additive profit markup – which, without loss of generality, can be the last "input." That is, we could let $x_n=1$, and $P_{k,n,t}$ = the additive profit markup in country k at time t. Expressed in this way, equation (2) is an identity.

Suppose we observe $P_{k,t}^{\text{Big Mac}}$ and $\{P_{k,j,t}\}$ for a sufficient number of time periods and countries, (or, to be precise, when the number of locations times the number of time periods ≥ n), then it is a matter of simple algebra to solve for all x_i, i=1,2,..., n. In fact, under our assumptions, a convenient way to solve for $\{x_1, x_2, ..., x_n\}$ would simply be a linear regression of $P_{k,t}^{\text{Big Mac}}$ on $\{P_{k,j,t}\}$. The regression in this case is not a statistical tool, but an algebraic one. Since equation (2) is an identity, the R^2 =100%.

Of course, we do not literally have price information on every single ingredient of a Big Mac. For example, we do not have information on cooking oil, pickles, sesame seeds, or "special sauce" in the data set. However, we assume that, in terms of their shares in the total cost of a Big Mac, these missing items are relatively unimportant when compared with the items for which we do have information, such as labor, rent, bread, ground beef, lettuce, and three other inputs. This assumption will be verified later.

The most serious "missing input" is probably the profit markup, which might vary by country and year. This and other "missing inputs" would go into the residual of a regression. In subsequent analyses when the role of the "missing inputs" may matter, we experiment with various assumptions about them to ensure that our key results are robust. These robustness checks will be explained later when relevant.

With these points in mind, we regress the price of a Big Mac on the prices of the eight main inputs for which we do have information, and report the results in Table 3. We report only the coefficients from the random effects estimator since a Hausman test that the covariance between the independent variables and the error term is equal to zero is not rejected. Failure to reject this hypothesis indicates that the random effects estimator is the efficient estimator. As reported in the table, the computed value of the test statistic is $\chi^2(8) = 5.8$, with a significance level = 0.67.

All of the coefficients and the implied shares seem reasonable. What stands out in Table 3 is the importance of nontraded inputs – especially labor – for the price of Big Macs. According to the table, the total nontraded goods share is at least 55 percent, i.e., $\alpha = 0.456 + 0.046 + 0.051 \approx 0.55$. Alternatively, if we normalize the non-tradable share by the total amount explained by all observed inputs, then, non-tradables collectively explain 64 percent of the Big Mac price ($\alpha = 0.553/0.869 \approx 0.64$).

We will use these estimates when we explicitly allocate shares of real exchange rate movement to deviations from the law of one price for traded goods, and the relative price of non-traded goods. Before doing so, however, we estimate the persistence of 'aggregate' Big Mac real exchange rates and compare them with the persistence of the Big Mac ingredients.

3.c. Fast food: how fast is convergence?

In this subsection we address two fundamental questions regarding convergence. First, we ask how the cross-country dispersion of *prices* of the Big Mac, and of its ingredients has evolved over the sample period. Convergence in dispersion closely corresponds to the idea of '$\sigma-$convergence', as described by Barro and Sala-i-Martin (1995) and Sala-i-Martin (1996) in their studies of cross-sectional income dynamics. In our context, the fact that we observe price levels (as opposed to price indexes) makes an analysis of $\sigma-$convergence possible and informative. Second, we compare the persistence of deviations from the law of one price for the Big Mac and for its ingredients. This aspect of convergence is related to the concept of $\beta-$convergence in economic growth empirics. Unlike studies of $\beta-$convergence using CPI-based real exchange rates however, in this study it is not necessary to presume a base year when parity held.

In Figure 3a, we plot the cross-country dispersion (as measured by the coefficient of variation) of prices for the Big Mac and for the five traded inputs over time; Figure 3b presents the same information for its non-traded inputs. Figure 3c summarizes the information by averaging within each category (traded, non-traded) and for the Big Mac over time. The first thing to notice is that dispersion in Big Mac prices is lower than that of any of its ingredients. Second, nearly all traded inputs display a downward trend in price dispersion (the exception is onions). Third, price dispersion is often larger for non-traded inputs than for traded inputs and actually increased during the sample for rent and wages. Combined, this is consistent with an ongoing process of global integration in traded goods markets, and an absence of such a process among non-traded inputs. According to Figure 3c, average dispersion in prices of traded inputs declined by 8%, while dispersion in non-traded input prices rose by 10%. Big Mac price dispersion also rose by 6%, but remained well below that for all other inputs – especially non-traded inputs. We will return to these results as we interpret our tests of $\beta-$convergence below.

As existing studies of convergence focus on real exchange rates rather than prices, we now shift the analysis to bilateral price differences in U.S. dollars. Define the (log) real exchange rate at time t as: $q_t = s_t + p_t^* - p_t$, where s_t is the domestic currency price of foreign exchange, p_t^* is the foreign price of Big Macs, and p_t is the domestic price of Big Macs; all variables are expressed in natural logarithms. We begin, in Table 4 by providing OLS estimates of $\hat{\beta}$ from equation 3 for the Big Mac real exchange rate, and each of the eight input-based real exchange rates.

$$\Delta q_{i,t} = \beta q_{i,t-1} + \text{country \& time dummies} + \varepsilon_{i,t} \tag{3}$$

Country and time dummies are included to capture unobserved effects, and robust standard errors are reported in parenthesis beneath the estimates of $\hat{\beta}$ for each equation. We will subsequently report several alternative specifications. The final two columns of the table report F-tests (with p-values) whether (a) the country fixed effects, and (b) whether both country and time fixed effects are zero. These hypotheses are rejected in nearly every case – implying that these fixed effects are important. The only cases where we cannot reject the null are for the two inputs displaying the greatest trend in dispersion, i.e., Lettuce and Rent in Figures 3a-b. In both cases adding time dummies raises the significance level, however for Rent we still are unable to reject the null that both country and time fixed effects are zero. Since Figure 3b demonstrates that price dispersion actually increases during the sample for Rent, this failure to reject should not be interpreted as evidence of long run absolute price parity for rent.

Immediately apparent in the table is the fact that *Tradables*, as a group, have the least persistence and the shortest half lives.[8] Indeed, the median half life for *Non-tradables* (3.6 years) is more than three times that for *Tradables* (1.2 years) and the half life of Big Mac deviations (1.8 years) lie somewhere in between. Note that since country fixed effects are not zero, the long run mean of the cross-country price difference is not zero (prices are not equalized in the long run). While it is reassuring that estimated mean reversion for *Non-tradables* is slower than that for

[8] All half-life calculations make the simplifying assumption of a zero intercept.

Tradables, the magnitude of the half life is still somewhat surprising. We note that the analysis of $\sigma -$ convergence presented above indicates that price dispersion is actually growing for *Non-tradables* as a group. This is an area lacking comparable studies. However, in terms of relative rates of mean reversion, the general pattern of results presented in Table 4 holds in all of our subsequent regressions.

To gauge the sensitivity of the results to outliers and alternative specifications, the analysis was repeated in a number of different ways. To conserve space here, these robustness checks and extensions are available in a separate on-line appendix.[9] In Appendix Table 1, for example, we excluded observations associated with the largest 5 percent of the residuals from the corresponding regression in Table 4. To summarize, nearly all the half lives rise – an aspect we explore below. The general pattern across ingredients and the Big Mac composite remains; namely, the half life of Big Mac deviations is bounded by that of *Tradables* from below, and of *Non-tradables* from above.

We also report the results of a different estimation method (see Appendix Table 2), i.e., we use the random effects estimator. Though the Hausman test suggests the fixed effects estimator is efficient, (i.e., the null hypothesis is rejected at the 10% level in all cases) we report the random effects estimates for comparison. Again, the same general pattern remains. Specifically, the half life of Big Mac deviations is bounded by that of *Tradables* from below, and of *Non-tradables* from above.

We also estimated persistence using the GMM method proposed by Arellano and Bond (1991) that corrects for the bias inherent in estimation of equation 3 by OLS. We implement this estimation for the Big Mac and for each of the ingredients, (see Appendix Table 3). For each equation, two specifications are presented: (a) including time fixed effects (Arellano and Bond's preferred specification) and another, (b) excluding them. Our results do not appear too sensitive to these permutations. In the specification including time dummies, the median half life for *Tradables* is just below that for Big Macs, which is beneath that for *Non-tradables*.

To make an explicit comparison with Cumby (1997), we restricted the sample (see Appendix Table 4). In these regressions we examine only those countries in Cumby's sample, and we also restrict the time period to be closer to his by dropping the final three years from our sample period. Considering the reduced set of countries, the results are very similar to those for the full sample of countries – except that estimated convergence is generally slightly faster than for the full sample – a result similar to that found by Cumby. Overall however, the same general pattern emerges across the nine real exchange rates in the table.

---

[9] All appendix tables related to this paper are available in a single file at: http://www.owen.vanderbilt.edu/david.parsley/, or http://www.nber.org/~wei/

We also estimated a regression including one lag of the dependent variable (see Appendix Table 5). Given our relatively short time dimension, increasing the number of lags has a non-trivial impact on the sample size. The lagged dependent variables are generally insignificant, and our conclusions about relative convergence speeds are unaffected.

Finally, we consider the effects of taxes. It is well known that taxes and other transaction costs can create a wedge – within which, real exchange rates need not display convergence tendencies. Moreover, time variation in these transaction costs can, in effect, present a 'moving target' for mean reversion estimates. That is, in principle, arbitrage may be gauged on either a pre-tax, or tax-inclusive basis. In other words, the regressions presented in Table 4 (and Appendix Tables 2-5) may therefore embody considerable measurement error since they use prices inclusive of VAT and sales taxes. Hence, we repeated the analysis after subtracting VAT and sales taxes.[10] Our results (in Appendix Table 6) are similar; the adjustment for VAT and sales taxes seem to matter little for estimated convergence rates. The most notable aspect of the regressions is that the estimated standard errors always rise, and the adjusted R-squared nearly always declines.

Several authors, e.g., O'Connell (1998), Obstfeld and Taylor (1997), Taylor (2001), Sarno and Taylor (2001), and O'Connell and Wei (2002) argue that estimates of real exchange rate persistence obtained from a linear regression are biased upward, since such estimates are essentially averages of two regimes: very low speed of convergence for deviations smaller than transaction costs, and possibly much faster convergence for larger deviations. These authors have addressed the problem of lumping data from two regimes by estimating a threshold autoregression (TAR) model. As O'Connell and Wei (2002) note, if transaction costs create a band of no-arbitrage, TAR models provide a more powerful way to detect global stationarity – even if the true price behavior does not conform to the TAR specification. We consider two such models of non-linear price adjustment – an Eq-TAR (for "equilibrium threshold autoregressive model"), and a Band-TAR – both of which can be represented by restrictions on equation 4.

[10] It should be noted that this adjustment may introduce error into the estimation since the sales tax data has been taken from a number of sources – many of which present the information in 'simplified' form only. For example, some countries tax 'agricultural products' while others tax them at a reduced rate, while others do not. Moreover, 'agricultural products' may include beef for some countries, while in other countries 'agricultural' may be taken to be 'vegetable'. While we have made considerable effort in compiling accurate data, we recognize the potential for error such ambiguities introduce. Parsley and Wei (1996) is the only study we know of that considers the effects of taxes on convergence rates. In their study of intra-national (U.S.) real exchange rates, they find that taxes have virtually no effect on their persistence estimates since there is simply not much variation over time in sales tax rates within the United States.

$$\Delta q_t^* = \begin{cases} \rho\left(q_{t-1}^* - b\right) + \varepsilon_t, & \text{if } q_{t-1}^* > c \\ \varepsilon_t, & \text{if } -c \le q_{t-1}^* \le c \\ \rho\left(q_{t-1}^* + b\right) + \varepsilon_t, & \text{if } q_{t-1}^* < -c \end{cases} \tag{4}$$

We often reject the hypothesis that country fixed effects are zero, hence we remove the long run means from q prior to estimation, and designate the de-meaned variable as q^*. According to the Eq-TAR model, convergence occurs toward the center of the band, hence the implied restriction is b=0. On the other hand, mean reversion in the Band-TAR model is assumed to be sufficient to push the price differences only toward the outer edge of the bands, hence this model imposes b=c. These models allow the real exchange rate to have a unit-root inside the transaction cost band. Once the real exchange rate exceeds the transaction cost parameter (c), the real exchange rate reverts at rate, $1-\rho$. In the Eq-TAR model, reversion is toward the center of the transaction cost band [-c, c], while in the Band-TAR model reversion is toward the edge of the threshold. The Eq-TAR model would characterize behavior if fixed costs are an important part of impediments to arbitrage. Similarly, if the impediments to arbitrage take the form of variable costs only, then the Band-TAR model would be appropriate. Currently, there is no consensus as to which model is uniformly 'best', and there does not exist a good way to estimate a general model that would nest both as special cases. It turns out that our conclusions are similar for either model.

Estimation of these models can be done via maximum likelihood or sequential conditional least squares. Franses and van Dijk (2000) demonstrate the equivalence of the two methods. Procedurally, we estimate the pooled model using the fixed effects panel estimator by performing a grid search over possible values of c. Starting with an initial value of c at 0.003, the search adds 0.003 in each successive round until c reaches the 75th fractile of the distribution of $|q^*|$. This results in roughly 100 estimations per good. The model with the minimum residual sum of squares is reported in Table 5.

For comparison, we present the Eq-Tar and Band-Tar results in the two sets of columns. Overall, the estimates of convergence are faster in these non-linear specifications, as one would expect. However, in both specifications the same pattern prevails as before. Namely, the median tradable good converges fastest, while non-tradables have the greatest persistence, with the Big Mac 'sandwiched' in between. Note that while the thresholds for *Non-Tradables* often appear smaller than those for *Tradables*, the results in section 3a suggest *Non-tradables* prices are actually becoming more disperse. Obstfeld and Taylor (1997) report thresholds of between 8 and 10 percent – while those in the table (for de-meaned q) are generally closer to those reported in Sarno and Taylor (2004), who examine more disaggregated data.

In summary, the speed of convergence for the Big Mac aggregate is bounded by the (faster) convergence rates of traded inputs and (slower) convergence rates of non-traded inputs, regardless of the estimation method chosen. We emphasize that such an inference is possible due to availability of both input and Big Mac prices that do not suffer from the problems of time aggregation or unknown index composition. We now turn to a formal decomposition of movements in Big Mac real exchange rates into parts attributable to movements in tradables and non-tradables separately.

3.d. Two for the price of one: new accounting versus old theory

Up to this point, the investigation has relied mostly on time series aspects of the data; we now turn to a question that crucially depends on the extensive cross-section dimension of the data. In most models of the real exchange rate, the relative price of non-tradable goods in terms of tradables plays a key role. For example, according to the well-known Harrod-Balassa-Samuelson effect, currencies from countries experiencing relatively faster tradable goods productivity growth will tend to appreciate.[11] Productivity growth however, is not the only source of movements in the relative price of non-tradables across countries. For example, Dornbusch (1989) and Froot and Rogoff (1991) argued that different government macroeconomic policies can also be important in explaining real exchange rate movements via a traded/non-traded goods price channel.

This view of the role of the relative price of non-tradables in real exchange rate determination has recently come under assault. In an influential and much-cited paper, Engel (1999) concludes that movements in relative prices of nontraded goods appear to account for essentially *none* of the movements in aggregate U.S. based CPI real exchange rates. Instead movements in real exchange rates are almost completely due to deviations from the law of one price for tradable goods. In subsequent discussion, we refer to this stark result as the Engel effect. The nature of the challenge is clear; namely, under this view, neither the Harrod-Balassa-Samuelson effect, nor the Dornbusch-Froot-Rogoff effect, help to explain movements in real exchange rates.

In this subsection, we examine whether it is possible that the Engel effect is important under some conditions but less so under others. Mendoza (2000) provides one hint that the Engel effect may sometimes be less than 100 percent. In a study of the Mexican peso/US dollar real exchange rate, Mendoza found the Engel effect is present when the country's nominal exchange rate was on a floating regime; but the effect declines to between 30% and 50% when the nominal exchange rate was tightly managed. A reasonable conjecture from Mendoza's case

[11] For textbook treatments, see, e.g., Caves, Jones, and Frankel (2002, p. 372-3), or Obstfeld and Rogoff (1996, p. 210-214).

study is that exchange rate volatility and/or the nominal exchange rate regime may play a role in determining the relative importance of international deviations in traded goods prices in explaining real exchange rate movements.

An important empirical drawback to the Engel (1999), Mendoza (2000), or Parsley (2001) accounting exercises is that they require a decomposition of the aggregate price index into traded and nontraded components. Engel acknowledges that the resulting decompositions are imprecise. One indication of the scope for such measurement error is provided in the resulting estimates of the non-traded goods weight in the six industrialized countries' CPI indexes he studies; that share varies from 24 percent in Italy to 46 percent in the U.S., a range that implies substantial variation in CPI baskets across countries. Of course, in order to obtain these estimates, one must make an assumption on the specific functional form that combines tradable and non-tradable prices into the aggregate price index. The conventional practice makes the simplifying assumption that traded and non-traded components are combined in a Cobb-Douglas fashion. In contrast, in this study we examine the robustness of the Engel effect for the case of Big Mac (aggregate) real exchange rates, where there is very little room for substitution across inputs either within or across countries. Hence, the decomposition is arguably more straightforward.

Another drawback to using aggregate CPI data is the well-known price index problem, i.e., price indexes tell us something about the change in prices from the base period. Assuming PPP held in the base period, observing changes in price indexes would convey the same information as examining price levels in each period. However, if PPP did not hold in the base period this link is severed, and movements in price indexes may not convey useful information about the level of the real exchange rate. In contrast, we present decompositions of movements of Big Mac real exchange rate levels, into shares attributable to traded- and non-traded inputs directly. An important goal of this analysis is to examine the robustness of these earlier studies in the context of a single "aggregate" good, where we know its composition reasonably precisely. Our methodological approach differs from previous studies in that we explore a much greater cross-section dimensionality (though shorter time series with lower frequency).

We begin by describing the decomposition of real exchange rates into traded and nontraded components. Express the Big Mac real exchange rate (Q^{BM}) as:

$$Q^{BM} = \frac{SP^{BM^*}}{P^{BM}}, \tag{5}$$

where, P^{BM^*} is the foreign currency price of a Big Mac abroad, and P^{BM} is the U.S. dollar price of a Big Mac in the United States. The nominal exchange rate (U.S. dollars per foreign currency) is designated by S, and we have suppressed time subscripts. Since $P^{BM^*} = P^{T^*} + P^{N^*}$ and $P^{BM} = P^T + P^N$, we can write the log real exchange rate as:

$$q^{BM} = \left[\log(S) + \log(P^{T^*}) - \log(P^T)\right] + \left[\log\left(1 + \frac{P^{N*}}{P^{T*}}\right) - \log\left(1 + \frac{P^N}{P^T}\right)\right]$$

The first part of this expression (in the first square bracket) is simply the deviation from the law of one price for traded inputs (*x*), and the second part is the relative-relative price of non-traded goods, i.e.,

$$\begin{aligned} &q^{BM} = x + y, \text{ where} \\ &x = \log(S) + \log(P^{T^*}) - \log(P^T) \text{ , and} \\ &y = \log\left(1 + \frac{P^{*N}}{P^{*T}}\right) - \log\left(1 + \frac{P^N}{P^T}\right) \end{aligned} \tag{6}$$

Unlike previous studies, a distinctive feature of this study is that traded goods prices can be computed directly as $P^T = \sum \hat{\beta}_i^T P_i^T$, where the summation is over the *i* traded inputs (*beef*, *cheese*, *lettuce*, *onions*, *and bread*) and the $\hat{\beta}$ estimates are computed previously in Table 3. A similar computation can be made for P^{*T}, P^N, and P^{*N}. Here, as in Engel (1999), the log Big Mac real exchange rate is the sum of deviations from the law of one price among traded ingredients, and the relative-relative price of nontraded inputs abroad and at home.

Armed with empirical counterparts to *x* and *y*, Engel's (1999) approach is to decompose movements in aggregate real exchange rates to shares attributable to movements in each. Using more than thirty years of monthly data he focuses on (among other measures) the mean squared error of changes in the real exchange rate at all horizons, e.g., 1-month, 2-months, up to the highest n-month difference the data would allow. In our case we have annual observations for thirteen years. The annual frequency and relatively short time span necessitate a different approach – one that nonetheless uses the total variation in the data set. Hence, we propose an alternative, time-dependent, way to construct the shares attributable to *x* and *y*. Since we observe prices (and not price indexes) we construct absolute (i.e., levels) measures of *x* and *y*, as well as for aggregate Big Mac deviations, at each point in time. We have a potential cross-section of 561 real exchange rates with 13 time series observations each (without missing values).[12] Our approach has the advantage that we can systematically relate these shares to

[12] Since we study 34 countries, we have 561 (=34x33/2) real exchange rates. However, we continue to focus on only those 343 good level real exchange rates highly correlated with aggregate real exchange rates. As described above, the restriction requires the correlation coefficients of both the level and the first-difference of the good-level real exchange rates with their corresponding aggregate real exchange rates must exceed 0.65. For robustness however, we also report results using all 561 real exchange rates.

observable country-pair and time-specific factors. For comparison with previous studies, we also present results using annual changes.

Generically, we construct the time-dependent measure of the share of Big Mac real exchange rates at time t attributable to *x* as the ratio of the squared deviation of *x* from its country-pair specific mean, to the sum of that for *x* and *y* together, i.e.,

$$x\text{–}share_t = \frac{\left(x_t - \bar{x}\right)^2}{\left(x_t - \bar{x}\right)^2 + \left(y_t - \bar{y}\right)^2}\,, \tag{7a}$$

We label this as 'share in variance' since it most closely approximates Engel's variance decomposition, though equation 7a preserves the time-series dimension; below we also consider an approximate mean-squared error version. As in Engel's formulation, the denominator of equation 7a does not equal the total variation of the Big Mac real exchange rate. This is because our cost share regressions did not allocate 100% of the variation of Big Mac prices to the ingredients we included. Hence we must also account for this unexplained portion for completeness. We adopt an agnostic view and report three separate approaches, namely, (a) ignoring the unexplained portion, (b) attributing the entire unexplained portion to x, and (c) attributing the entire unexplained portion to y. As it turns out, the three approaches yield qualitatively similar results with regard to our key results. We therefore conclude that how we attribute the unexplained portion is not crucial to our discussion below.

Figure 4 plots the histograms of these three measures of *x-share*, using all available cross-section and time series data points. That is, without missing values there will be 13 observations for each of the 343 'highly-correlated' real Big Mac exchange rates that we have been focusing on previously, i.e., those with correlation coefficients > 0.65 between CPI and Big Mac real exchange rates in both levels and in first differences (i.e., nearly 4500 observations). The x-axis records the share of traded-goods deviations in the aggregate Big Mac real exchange rate. The x-axis labels indicate the lower bound of each bin, e.g., 80% stands for the percent above 80%. The height of the bars measures the percentage of real exchange rates meeting that criterion. The figure indicates that there is considerable heterogeneity across the 343 real exchange rates. In particular, in less than 40% of the cases do we get the result that *x* accounts for more than 80% of real exchange rates. This is true whether we attribute the unexplained portion to *x*, or to *y* or whether we ignore it and focus on the 'pure' version of equation 7a. Moreover, it is apparent that *x* accounts for a relatively small portion of real exchange rates for a non-trivial proportion of the real exchange rates in our sample. Thus using direct measures of the size of traded goods deviations relative to overall real exchange rate deviations (as opposed to changes in deviations used in previous studies) we see that the Engel effect is substantially smaller than has been documented in studies using price index data.

We now turn to a more systematic panel-regression analysis using both the cross-sectional and time series information in our data. Inspired by Mendoza's (2000) findings for Mexico, we explicitly consider the effect of the exchange rate regime.[13] We begin by incorporating a dummy variable for the U.S. dollar pegs of Argentina, and Hong Kong. This dummy (*$peg*) takes the value one corresponding to these four country-pairs for all time periods except 2002, for pairs involving Argentina. We also include a dummy variable (*Euro*) for the Euro countries during the 1999-2002 time periods. However, a more general (i.e., continuous) way to capture exchange rate effects is to incorporate exchange rate variability directly into the specification.

Thus, the basic specification we report in Table 6 includes the three variables (*$peg*, *xrvol*, and *Euro*). In the second column of the table we add time and city dummies. In the final specification we add controls for membership in a trade bloc, sharing a common language, the level of tariffs between the country-pair (= the sum of tariffs in countries *i* and *j*), and the (log) distance between their capital cities. The most general specification is shown as equation 8 below.

$$\begin{aligned} x\text{–}share_t = {} & \beta_1 xrvol_{ij,t} + \beta_2 \$peg + \beta_3 Euro + \beta_4 \ln(dist_{ij}) + \beta_5 Tariff_{ij} \\ & + \beta_6 \text{Common Language} + \beta_7 \text{Bloc}_{ij} + city\ and\ time\ dummies + \varepsilon_{ij,t} \end{aligned} \tag{8}$$

Distance is calculated using the great circle formula using each country's capital city's latitude and longitude. Exchange rate variability is defined as the standard deviation of changes in the monthly bilateral exchange rate (between the country-pairs involved) during each year. Tariff is defined as the sum of the two average tariff rates in countries i and j, unless the two countries are both in the same free trade area or customs union (such as within the United States, or within the European Union). In these cases the value for tariff is set equal to zero. The first two columns of Table 6 (labeled *pure*) report the results where the variation in the unexplained portion of Big Mac prices is ignored. In the second group of columns (labeled *over-attribution to x*) the variation in the unexplained portion of Big Mac prices has been attributed to *x*, and in the third group of columns, this variation has been attributed to *y*.

The results in the table are quite stable across all specifications. First, higher exchange rate volatility is associated with a larger *x-share*, i.e., higher exchange rate volatility exaggerates the importance of traded goods deviation. Second, having a peg to the U.S. dollar lowers the contribution of deviations from the law of one price in traded goods to movements in

---

[13] Exchange rate regimes have also been found to be important for real exchange rate behavior by, e.g., Frankel and Rose (1996), or Parsley and Wei (2003).

'aggregate' real exchange rates, as hypothesized by Mendoza (2000). Results for the Euro, however, are generally weaker – though also in the same direction. Tariffs are negative and statistically significant, which suggests that tariffs diminish the impact traded goods have on real exchange rates, since higher tariffs reduce the scope for arbitrage. Distance is strongly statistically significant across all specifications, suggesting that arbitrage is less important for more distant locations. Having a common language does not seem important. The trade blocs we include have some mixed results. The European Union dummy is negative (but insignificant) when the *y-share* is over-attributed (i.e., it includes the entire unexplained portion), but positive and insignificant when *x-share* is over-attributed. Surprisingly, Mercosur, APEC, and ASEAN all seem to be positively associated with *x-share*. This may reflect the overall size of traded goods price disparities among these countries.

One may wonder if our results are specific to the subset of real exchange rates we are studying. Hence, we repeated the estimation including all Big Mac real exchange rates – i.e., even those with correlations with CPI real exchange rates below 0.65. These results (reported in the on-line appendix as Appendix Table 7) hardly change, suggesting they are not limited to our specific subsample.

We also examine a more comprehensive measure of variation in the real exchange rate. Equation 7b is approximately the share of the mean squared error (MSE) of the real exchange rate attributable to x.[14] In our case however, the MSE of each term (x and y) is computed as the sum of the time t squared deviation plus the time t deviation from the mean squared. As before, we present three different measures of 7b depending on how we treat potential covariation between x and y.

$$x\text{–}share_t = \frac{\left(x_t - \bar{x}\right)^2 + x_t^2}{\left(x_t - \bar{x}\right)^2 + \left(y_t - \bar{y}\right)^2 + x_t^2 + y_t^2}, \qquad (7b)$$

Results using this as the dependent variable (see on-line appendix Table 8) are largely similar to those using the share in variance as the dependent variable (i.e., Table 6). Namely, (1) higher exchange rate volatility is associated with a larger *x-share*; (2) having a peg to the U.S. dollar lowers the contribution of deviations from the law of one price in traded goods to movements in 'aggregate' real exchange rates; (3) x generally accounts for a higher proportion for countries that are farther apart; and (4) tariffs are negative and statistically significant. Finally, these results suggest that (when significant) x accounts for a higher proportion for countries that have adopted the euro, even after controlling for the European union, as well as countries in Mercosur and Asean.

[14] Our equation 7b corresponds to Engel's (1999) equation B1.

One potential statistical problem is that the dependent variable, a share, is constrained to lie between zero and one. Strictly speaking, the normality assumption of the error term in the OLS specification is incompatible with this. We address this issue by taking a logistic transformation of *x-share*, which allows the dependent variable to take any positive or negative value (see Greene 1997, p.228). For the definition of *x-share* given in equation 7a, the new dependent variable becomes:

$$x\text{–}share_t = \ln\left(\frac{\left(x_t - \bar{x}\right)^2}{\left(x_t - \bar{x}\right)^2 + \left(y_t - \bar{y}\right)^2}\right) - \ln\left(1 - \frac{\left(x_t - \bar{x}\right)^2}{\left(x_t - \bar{x}\right)^2 + \left(y_t - \bar{y}\right)^2}\right), \qquad (9)$$

Results using this specification (corresponding to equations 7a and 7b, respectively) are reported in (on-line) appendix Tables 9 and 10. Statistical significance generally rises using this specification, but other conclusions remain qualitatively the same. The only notable changes are that distance is often not statistically significant, the dummy for Common language is negative, though it is not generally statistically significant, and the trade bloc dummies (APEC and ASEAN) become statistically insignificant. All other conclusions hold under this transformation.

So far, we have studied the share of deviations from purchasing power parity attributable to deviations from the law of one price in traded goods. In contrast, previous studies have focused on share of *changes* in real exchange rates attributable to *changes* in deviations from the law of one price in traded goods. In previous studies, this emphasis was necessary since the level of the real exchange rate using index data (e.g., CPI) is arbitrary. Thus the measure we study here more directly accounts for movements in the level of real exchange rates. We have shown that deviations from the law of one price in traded goods generally account for a much smaller portion of real exchange rate movement than previous studies would have led us to expect. We conclude that differences between our results and those in previous studies are partly attributable to our much larger cross-section, our ability to focus on real exchange rate levels, and to our ability to decompose the real exchange rate relatively precisely.

We have also shown that exchange rate variability is strongly positively related, and exchange rate pegs (especially the US dollar pegs in this sample) are strongly negatively related, to the fraction of absolute PPP deviations one can attribute to traded goods price disparities. Finally, we have found that the importance of law of one price deviations is often higher for countries participating in regional trading blocs.

To check whether our findings are driven by our focus on the level of the real exchange rate, we also estimated first-differenced versions of 7a and 7b. These results are reported in (the on-line) appendix tables 11 and 12. We also perform a logistic transformation of the dependent variable so that the transformed dependent variable is not bounded by zero and one. The relevant regressions results are reported in Appendix tables 13 and 14 (also on-line). The message that exchange rate variability raises the importance of deviations from the law of one price in real exchange rate movements continues to hold. However, other conclusions are less apparent in this weaker version of the decomposition. That is, the formerly robust conclusions concerning the dollar peg, distance, the European Union, and tariffs, are no longer apparent. Since the level of real exchange rate can be meaningfully measured in our thought experiment, we regard the analyses on (7a) and (7b) as more informative.

4. Thoughts at the checkout counter

This paper has studied one particular 'aggregate' real exchange rate – i.e., the Big Mac real exchange rate – where we know a great deal about how that aggregate is constructed. We have shown that Big Mac real exchange rates are generally highly correlated with CPI-based real exchange rates both in first differences as well as in levels. Thus, the lessons learnt from the Big Mac real exchange rates are relevant for the CPI-based real exchange rates. Our main innovation is to match these prices to the prices of individual ingredients (ground beef, bread, lettuce, labor cost, rent, etc.) in 34 countries during 1990-2002, which allows us to conduct a number of useful thought experiments.

As a result of our focus on *prices* and real exchange rate *levels* we uncover a number of interesting findings. First, we find that the non-traded component of Big Mac prices is substantial, i.e., between 55% and 64%. Second, the non-traded component displays greater price dispersion than the traded component of Big Mac prices; the Big Mac itself has lower price dispersion than any of its ingredients. Moreover, cross-country price dispersion ($\sigma-$convergence) actually increased for non-traded inputs, while falling for traded inputs over the time-frame of this study. We also examined the persistence of the real exchange rate ($\beta-$convergence). Our setting is arguably free of a number of possible biases induced by non-comparability of consumption baskets across countries, product aggregation bias (Imbs, Mumtaz, Ravn, and Rey, 2005), and time aggregation bias (Taylor, 2001). We find that the speed of convergence for tradable inputs is sufficiently fast to be compatible with economic theories (Chari, Kehoe, and McGratten, 2002), and that for the Big Mac real exchange rates is slower than the speed for its tradable inputs, but faster than its non-tradable inputs. Finally, we show that Engel's result that all movements in real exchange rates are attributable to deviations from the law of one price in traded goods is inversely related to exchange rate volatility. That is, reduced exchange rate volatility, exchange rate pegs, and tariff barriers generally weaken the Engel effect.

Table 1: Geographic Coverage

Countries and Regions

Europe	***Western Hemisphere***	***Asia, Pacific, and Africa***
Austria (Vienna)	Argentina (Buenos Aires)	Australia (Sydney)
Belgium (Brussels)	Brazil (Sao Paulo)	China (Beijing)
Czech Republic (Prague)	Canada (Toronto)	Hong Kong, SAR
Denmark (Copenhagen)	Chile (Santiago)	Indonesia (Jakarta)
England (London)	Mexico (Mexico City)	Israel (Tel Aviv)
France (Paris)	United States (Chicago)[1]	Japan (Tokyo)
Germany (Berlin)		Malaysia (Kuala Lumpur)
Hungary (Budapest)		New Zealand (Auckland)
Ireland (Dublin)		Singapore
Italy (Rome)		South Africa (Johannesburg)
Netherlands (Amsterdam)		South Korea (Seoul)
Poland (Warsaw)		Taiwan (Taipei)
Spain (Madrid)		Thailand (Bangkok)
Sweden (Stockholm)		
Switzerland (Zurich)		

[1] To correspond with the Economist's Big Mac Index, data for the U.S. is an average of Atlanta, Chicago, San Francisco, and Washington, D.C.

Table 2: Correlation of Big Mac and CPI based Real Exchange Rates
(In-Sample U.S. Dollar Bilateral Real Exchange Rates, 1990-2002)

Country	Correlation in levels	Correlation in changes
Australia	0.938	0.893
Austria	0.992	0.986
Belgium	0.657	0.886
Brazil	0.895	0.967
Denmark	0.966	0.866
France	0.941	0.704
Germany	0.956	0.878
Indonesia	0.727	0.967
Japan	0.886	0.927
Malaysia	0.912	0.846
Mexico	0.827	0.860
Netherlands	0.759	0.851
New Zealand	0.947	0.895
Singapore	0.783	0.732
South Africa	0.925	0.882
South Korea	0.932	0.909
Spain	0.954	0.778
Sweden	0.993	0.942
Switzerland	0.971	0.987
Taiwan	0.841	0.917
Thailand	0.906	0.670
Medians:		
U.S. bilateral rates	0.891	0.873
All bilateral	0.889	0.915

Table 3: Cost Function Estimation for Big Mac Production
(1990 – 2002)

Ingredient	*Regression in Levels* Coefficient Estimates[1]	Implied Cost Share (%)[2]	*Change Regression* Coefficient Estimates[3]
***Traded*:**			
Beef	3.010	9.0	2.257
	(0.645)		(0.669)
Cheese	2.530	9.4	1.995
	(0.592)		(0.625)
Lettuce	1.546	0.7	6.017
	(3.645)		(3.476)
Onions	1.156	0.5	4.411
	(3.610)		(3.239)
Bread	13.428	12.1	11.256
	(3.053)		(3.200)
Nontraded:			
Labor	9.245	45.6	11.823
	(0.832)		(1.069)
Rent	0.008	4.6	0.010
	(0.003)		(0.004)
Electricity	0.085	5.1	0.078
	(0.027)		(0.039)
		Total = 86.9%	
# of observations	318		284
Adjusted R-squared	.95		.66

[1, 3] Coefficient estimates and standard errors are multiplied by 100. Estimation method is random effects. Hausman test statistic for levels regression is $\chi^2(8) = 5.8$ (significance level =0.67), and the test statistic for the change regression (1st differences) is $\chi^2(8) = 3.3$ (significance level =0.91

[2] The share attributed to the ith ingredient is computed as: $\hat{\beta}_i \overline{P}_i / \overline{P}_{Big\ Mac}$, where $\overline{P}_i$ is the average price of the ith input.

Table 4: Persistence Estimates

Tradables	$\hat{\beta}$	Half-life[1]	# obs	$\overline{R}^2$	$H_0^1 : \lambda_i = 0$	$H_0^2 : \lambda_i = 0$, $\theta_t = 0$
Beef	-0.431	1.2	256	0.17	2.357	2.100
	(0.056)				(0.001)	(0.001)
Cheese	-0.451	1.2	252	0.22	3.074	2.667
	(0.055)				(0.001)	(0.000)
Lettuce	-0.358	1.6	246	0.13	1.197	1.449
	(0.055)				(0.252)	(0.064)
Onions	-0.609	0.7	256	0.27	4.678	3.576
	(0.060)				(0.000)	(0.000)
Bread	-0.252	2.4	256	0.08	1.980	1.734
	(0.049)				(0.007)	(0.011)
Median		1.2				
Non-Tradables						
Labor	-0.250	2.4	227	0.09	1.689	1.746
	(0.052)				(0.038)	(0.013)
Rent	-0.157	4.1	253	0.03	1.329	1.318
	(0.040)				(0.154)	(0.127)
Electricity	-0.177	3.6	256	0.16	2.780	2.332
	(0.035)				(0.000)	(0.000)
Median		3.6				
Big Mac	-0.326	1.8	203	0.12	1.528	1.696
	(0.061)				(0.070)	(0.016)
Country fixed effects			yes			
Time fixed effects			yes			

This table reports the results of estimating the following equation, using a fixed-effects estimator: $\Delta q_{i,t} = \beta q_{i,t-1} + \sum \lambda_i \text{country}_i + \sum \theta_t \text{time}_t + \varepsilon_{i,t}$

[1] The half-life is computed assuming a zero intercept.

The columns labeled H_0^1 and H_0^2, report the calculated F-test statistic for the indicated test, with the associated significance levels in parenthesis.

Table 5: Persistence Estimates Compared (TAR specifications)

	EQ-TAR				Band-TAR			
Tradables	$\hat{\beta}$	Threshold	Half-life	# obs	$\hat{\beta}$	Threshold	Half-life	# obs
Beef	-0.447 (0.058)	0.024	1.17	237	-0.462 (0.061)	0.024	1.12	237
Cheese	-0.500 (0.057)	0.018	1.00	230	-0.488 (0.058)	0.024	1.04	226
Lettuce	-0.393 (0.062)	0.051	1.39	207	-0.427 (0.065)	0.063	1.25	198
Onions	-0.666 (0.064)	0.030	0.63	237	-0.680 (0.065)	0.030	0.61	237
Bread	-0.277 (0.052)	0.018	2.14	233	-0.280 (0.053)	0.018	2.11	233
Median			1.17				1.12	
Non-Tradables								
Labor	-0.261 (0.054)	0.009	2.29	214	-0.265 (0.057)	0.009	2.25	214
Rent	-0.168 (0.043)	0.018	3.76	228	-0.200 (0.049)	0.049	3.10	208
Electricity	-0.188 (0.035)	0.006	3.32	241	-0.183 (0.037)	0.021	3.42	233
Median			3.32				3.10	
Big Mac	-0.365 (0.065)	0.015	1.53	181	-0.407 (0.072)	0.018	1.33	176

This table reports estimates of equation (4) in the text.

Table 6: Contribution of Traded Good Deviations to Big Mac Real Exchange Rate Movements (1990-2002) (Share in variance: levels of real exchange rate)

	"pure" measure		Over attribution to "x"		Over attribution to "y"	
Exchange Rate Volatility	1.429 (0.267)	1.408 (0.282)	1.523 (0.256)	1.293 (0.268	1.512 (0.237)	0.773 (0.254)
$ Peg	-0.415 (0.127)	-0.440 (0.132)	-0.250 (0.098)	-0.279 (0.096)	-0.134 (0.126)	-0.162 (0.129)
Euro	-0.130 (0.080)	-0.128 (0.082)	-0.065 (0.049)	-0.065 (0.049)	-0.014 (0.181)	0.032 (0.180)
Distance		0.038 (0.006)		0.041 (0.005)		0.069 (0.005)
Sum Tariffs		-0.008 (0.003)		-0.008 (0.002)		-0.010 (0.002)
Common Language		-0.047 (0.027)		0.000 (0.021)		0.012 (0.022)
European Union		-0.012 (0.041)		0.040 (0.032)		-0.041 (0.037)
Mercosur		0.245 (0.065)		0.199 (0.041)		0.420 (0.057)
Apec		0.119 (0.033)		0.011 (0.029)		0.076 (0.027)
Asean		0.183 (0.089)		0.164 (0.070)		0.187 (0.102)
Nafta		0.000 (0.000)		0.000 (0.000)		-0.071 (0.069)
Observations	2304	2115	2404	2214	2948	2742
Adjusted R-squared	0.304	0.312	0.110	0.130	0.027	0.087
Time Dummies	yes	yes	yes	yes	yes	yes
Country Dummies	yes	yes	yes	yes	yes	yes

This table presents results using the definition of x-share given in equation 7a in the text.

Figure 1

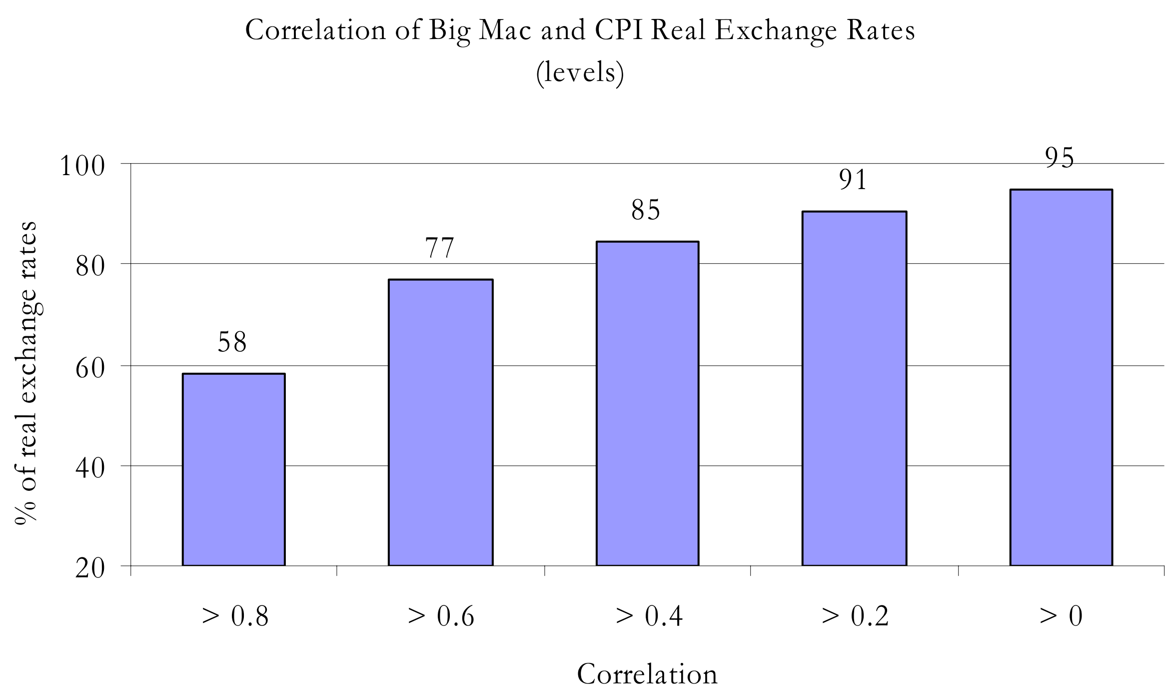
Correlation of Big Mac and CPI Real Exchange Rates
(levels)
% of real exchange rates
100
80
60
40
20
58
77
85
91
95
> 0.8
> 0.6
> 0.4
> 0.2
> 0
Correlation

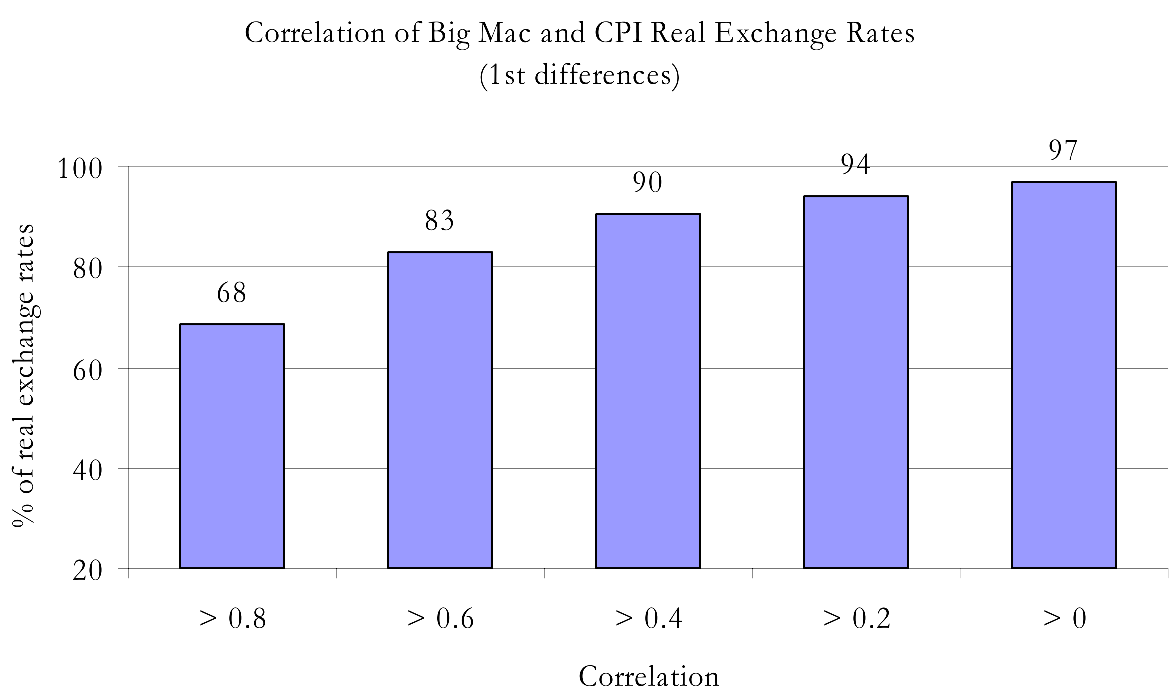
Correlation of Big Mac and CPI Real Exchange Rates
(1st differences)
% of real exchange rates
100
80
60
40
20
68
83
90
94
97
> 0.8
> 0.6
> 0.4
> 0.2
> 0
Correlation

Figure 2

Correlation of Big Mac and CPI Real Exchange Rates
(Low nominal exchange rate variability sub-sample of highly correlated RERs)

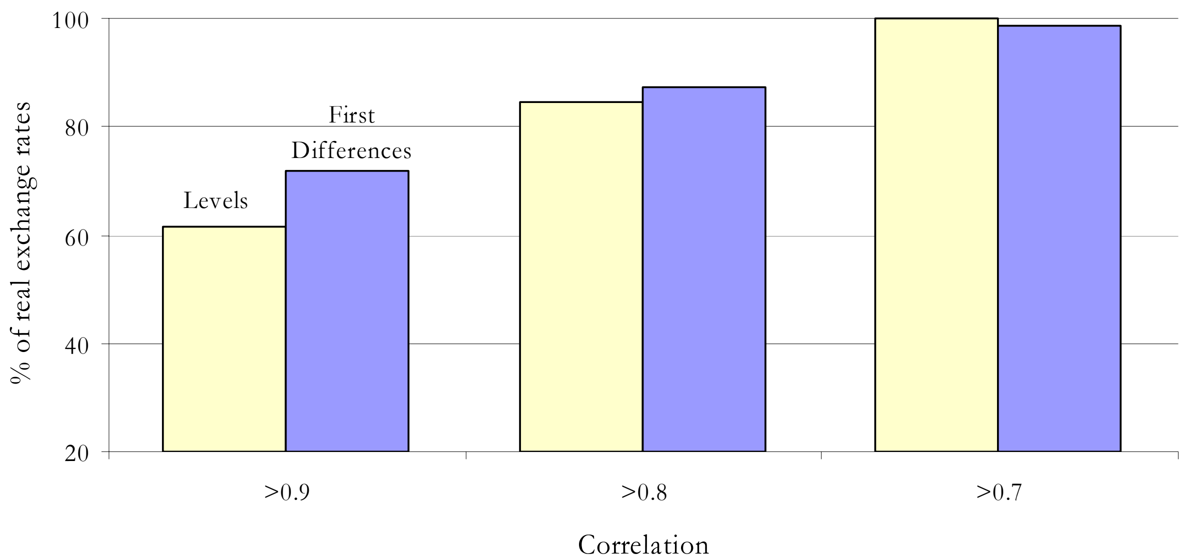

Figure 3a

Price Dispersion in Traded Inputs

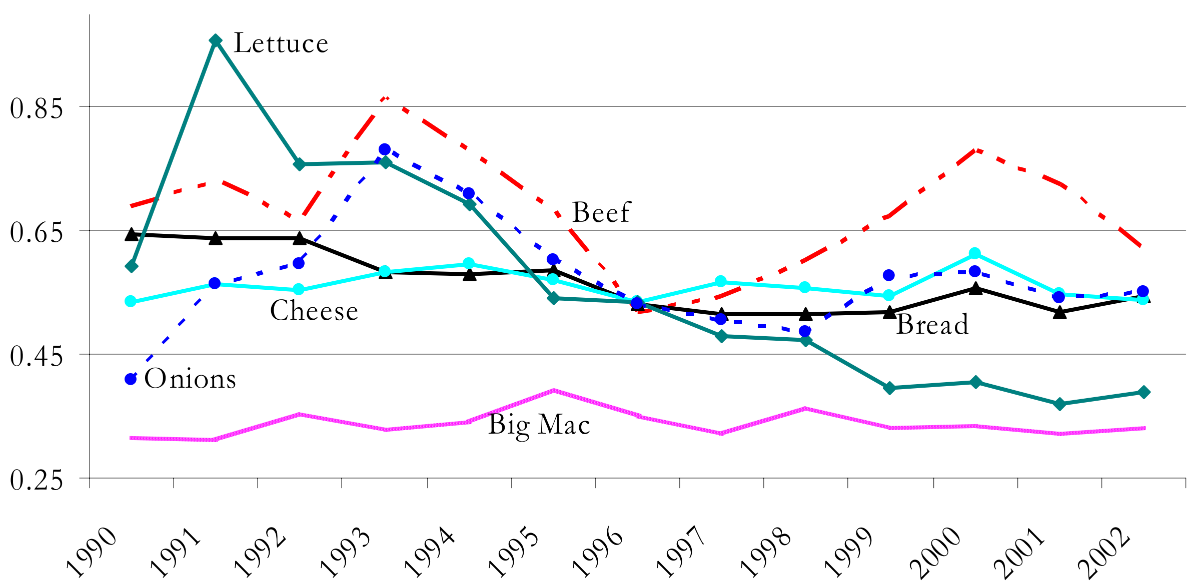

Figure 3b

Price Dispersion in Non-traded Inputs

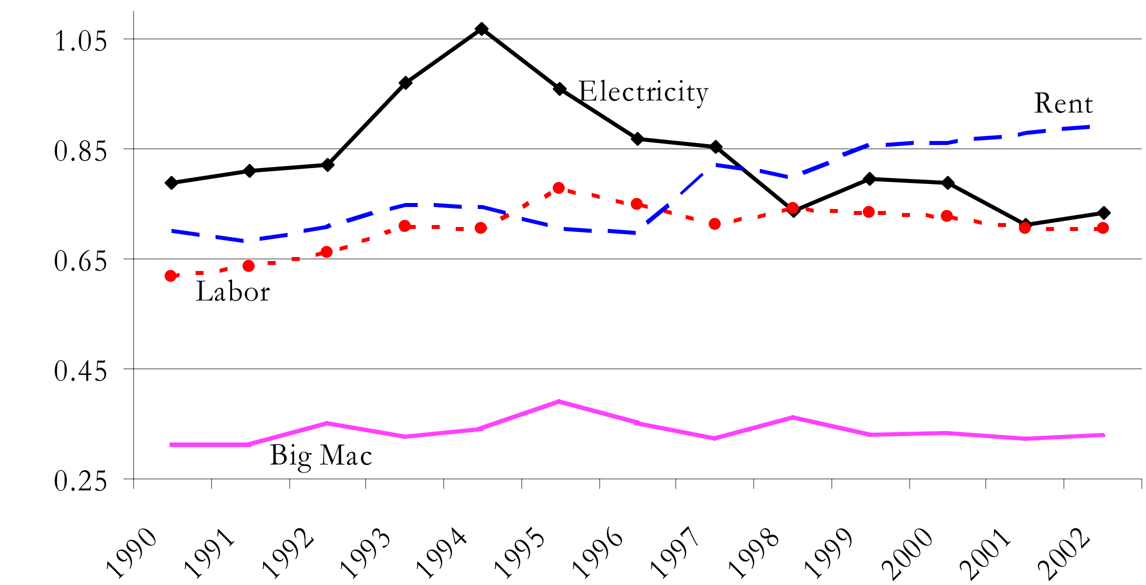

Figure 3c

1990-2002 Change in Absolute Price Dispersion
(average percent change in coefficient of variation of prices)

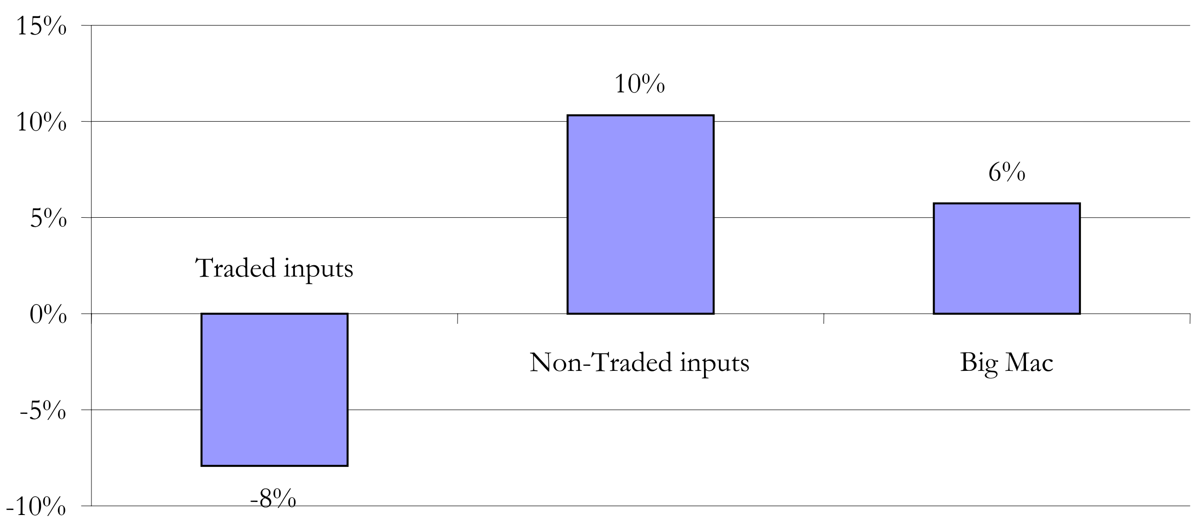

Figure 4

Share of Traded Goods Price Deviations in
Big Mac Real Exchange Rates
(343 real exchange rates, all years)

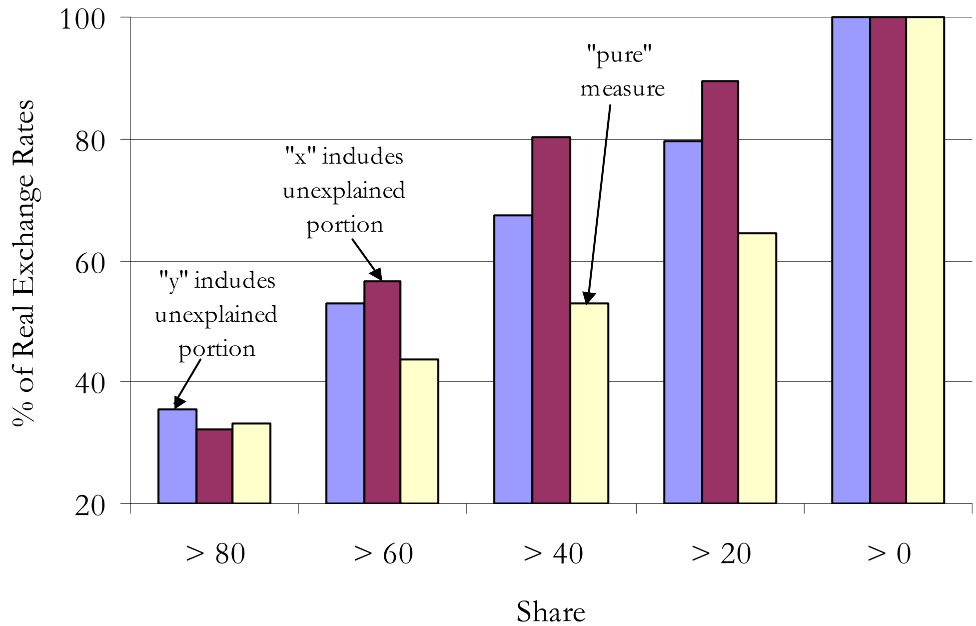

References

Arellano, Manuel, and Stephen Bond, 1991, "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations", *Review of Economics Studies*, 58, pages 277-297.

Barro, R. J. and Sala-i-Martin, X., 1995, **Economic Growth**, MIT Press, Cambridge.

Caves, Richard, J. Frankel, and R. Jones, 2002, **World Trade and Payments: An Introduction**, Ninth edition Harper Collins, New York.

Chari, V.V., Patrick. Kehoe, and Ellen. McGratten, 2002, "Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?", *Review of Economic Studies*, 69(3), pages 533-63.

Click, Reid W., 1996, "Contrarian MacParity", *Economics Letters,* 53(2)*,* November 1996, pages 209-12.

Cumby, Robert, 1997, "Forecasting Exchange Rates and Relative Prices with the Hamburger Standard: Is What You Want What You Get with McParity?", also **NBER Working Paper #5675**, July 1996.

Dornbusch, Rudiger, 1989, "Real Exchange Rates and Macroeconomics: A Selective Survey", *Scandinavian Journal of Economics*, 91(2), pages 401-32.

Engel, Charles, 1999, "Accounting for U.S. Real Exchange Rate Changes", *Journal of Political Economy*, V107, pages 507-38.

Frankel, Jeffrey, and Andrew Rose, 1996, "A Panel Project on Purchasing Power Parity: Mean Reversion Within and Between Countries", *Journal of International Economics*, 40(1-2), pages 209-224.

Franses, Philip Hans, and Dick van Dijk, 2000, **Nonlinear Time Series Models in Empirical Finance**, Cambridge University Press, Cambridge, United Kingdom.

Friedman, Milton, 1953, **Essays in Positive Economics**, University of Chicago Press, Chicago.

Froot, Kenneth, and Kenneth Rogoff, 1991, "The EMS, the EMU, and the Transition to a Common Currency", *NBER Macroeconomics Annual*, 6, p. 269-317.

Froot, Kenneth, and Kenneth Rogoff, 1995, "Perspectives on PPP and Long-Run Real Exchange Rates", in **Handbook of International Economics, Volume 3**, Grossman and Rogoff, editors, Elsevier, New York.

Greene, William., 1997. Econometric Analysis. Third edition, Prentice Hall, New Jersey.

Imbs, Jean., H. Mumtaz, M. Ravn, and Helen. Rey, 2005, "PPP Strikes Back: Aggregation and the Real Exchange Rate", *Quarterly Journal of Economics*, 120(1), p.1-44.

Lutz, Matthias, 2001, "Beyond Burgernomics and MacParity: Exchange Rate Forecasts Based on the Law of one Price", University of St. Gallen working paper.

Mendoza, Enrique, 2000, "On the Instability of Variance Decompositions of the Real Exchange Rate across Exchange-Rate-Regimes: Evidence from Mexico and the United States", Duke University working paper.

O'Connell, Paul, 1998, "Market Frictions and Real Exchange Rates", *Journal of International Money and Finance,* 17(1), pages 71-95.

O'Connell Paul, and Shang-Jin Wei, 2002, "The Bigger They Are, The Harder They Fall: How Price Differences Across U.S. Cities Are Arbitraged", *Journal of International Economics*, 56(1), pages 21-53.

Obstfeld, Maurice, and Kenneth Rogoff, 1996, **Foundations of International Macroeconomics**, MIT Press, Cambridge, Mass.

Obstfeld, Maurice, and Kenneth Rogoff, 2000, "The Six Major Puzzles in Macroeconomics: Is There a Common Cause?", *NBER Macroeconomics Annual* (15).

Obstfeld, Maurice, and Alan Taylor, 1997, "Non-linear Aspects of Goods-Market Arbitrage and Adjustment: Heckscher's Commodity Point Revisited," *Journal of Japanese and International Economics,* 11, pages 441-479.

Ong, Li-Lian (1997): "Burgernomics: The Economics of the Big Mac Standard," *Journal of International Money and Finance, 16(6),* December 1997, pages 865.

Pakko, Michael, and Patricia S. Pollard, 1996 "For Here or To Go? Purchasing Power Parity and the Big Mac," *Federal Reserve Bank of St. Louis Review, 78(1), Jan.-Feb. 1996, pages 3-22.*

Parsley, David, 2001, "Accounting for Real Exchange Rate Changes in East Asia", Hong Kong Institute of Monetary Research Working Paper #6/2001.

Parsley, David, and Helen Popper, "Accounting for Real Exchange Rate Changes with Individual Goods", December 2002, Vanderbilt University working paper.

Parsley, David, and Shang-Jin Wei, 1996, "Convergence to the Law of One Price without Trade Barriers or Currency Fluctuations", *Quarterly Journal of Economics*, 111(4), pages 1211-1236.

Parsley, David, and Shang-Jin Wei, 2003, "How Big and Heterogeneous are the Effects of Currency Arrangements on Goods Market Integration: a Price Based Approach?", Vanderbilt University working Paper.

Pedroni, Peter, 1999, "Critical Values for Cointegration Tests in Heterogeneous Panels with Multiple Regressors", *Oxford Bulletin of Economics and Statistics*, 61(Special Issue), pages 653-70.

Rogoff, Kenneth, 1996, "The Purchasing Power Parity Puzzle", *Journal of Economic Literature*, 34, pages 647-668.

Sala-i-Martin, Xavier X., 1996, "Regional Cohesion: Evidence and Theories of Regional Growth and Convergence", *European Economic Review*, 40, pages 1325-1352.

Sarno, Lucio, and, Mark Taylor, 2004, "Non-Linear Dynamics in Deviations from the Law of One Price: A Broad-Based Empirical Study", *Journal of International Money and Finance, 23*(1), pages 1-25.

Sarno, Lucio, and, Mark Taylor, 2002, "Purchasing Power Parity and the Real Exchange Rate", *International Monetary Fund Staff Papers*, 49(1), pages 65-105.

Taylor, Alan, 2001, "Potential Pitfalls for the Purchasing-Power Parity Puzzle? Sampling and Specification Biases in Mean-Reversion Tests of the Law of One Price", *Econometrica*, 69, pages 473–498.

Taylor, Alan, and Mark Taylor, 2004, "The Purchasing Power Parity Debate", *Journal of Economic Perspectives*, 18(4), p. 135-58.