The Economic Journal, **117** (*October*), 1336–1356. © The Author(s). Journal compilation © Royal Economic Society 2007. Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

A PRISM INTO THE PPP PUZZLES: THE MICRO-FOUNDATIONS OF BIG MAC REAL EXCHANGE RATES*

David C. Parsley and Shang-Jin Wei

We match Big Mac prices with prices of its ingredients as a unique prism to study real exchange rates (RERs). This approach has several advantages. First, the *levels* of the Big Mac RER can be measured meaningfully. Second, as the *exact* composition of a Big Mac is known, the contributions of its tradable and non-tradable components can be estimated relatively precisely. Third, the dynamics of the RER can be studied in a setting free of several biases inherent in CPI-based RERs. Finally, a large cross-country dimension allows us to overturn the Engel result on what drives RERs.

> 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 recognised; see, for example, Friedman (1953) and, more recently, Obstfeld and Rogoff (2000). According to a recent study by Sarno and Taylor (2002), the concept of Purchasing Power Parity (PPP) has been propounded as a theory of real exchange rate behaviour at least since the 1500s. Indeed, Frankel and 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, Rogoff (1996) observed a 'remarkable consensus' on the estimated half-life of deviations from PPP 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 *et al.*, 2002).

Second, the role of the relative price of non-tradable goods (e.g., through the Balassa-Samuelson effect) in accounting for movements in real exchange rates seems too small.[1] A much-cited paper by Engel (1999) seriously undermines the conventional view; he

* This is a substantially revised version of NBER Working Paper no.10074. A longer version is available on the authors' websites. We thank Robert Cumby, Peter Pedroni, Bruce Reynolds, Kenneth Rogoff, Lucio Sarno, Moto Shintani, the Editor, two anonymous referees, and seminar participants at the University of Chicago, the University of Virginia, UC San Diego, Vanderbilt, Louisiana State University, the IMF, the ASSA meetings, the Bank of France, and the European Central Bank for useful comments. The views presented are those of the authors and do not necessarily reflect those of any institution with which they are affiliated.

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

reports 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. First, the CPI baskets across countries are not identical and 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, commonly used linear estimation may be misspecified. Several authors (Obstfeld and Taylor, 1997; Taylor, 2001; Sarno and Taylor, 2002; O'Connell and Wei, 2002) argue that arbitrage costs dictate a nonlinear 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 nonlinearity is taken into account, the estimated half-life is usually shorter (with 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 may be overestimated. 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 an 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.

We adopt an alternative approach to study real exchange rates by using information in a panel of Big Macs. Like the aggregate CPI, the Big Mac is a 'basket' comprised of both traded and non-traded components. We demonstrate that the 'Big Mac' and CPI-based real exchange rates have similar time-series properties. In particular, 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' corporate pricing policies are not of first-order importance for our inferences.

One of our key innovations is to match Big Mac prices to the prices of individual ingredients, e.g., ground beef, bread, lettuce, labour, 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 across time periods. In fact, due to McDonalds' global advertising

[2] As an example, the French basket may have some weight on cheese, which the Chinese may not care much about; while the Chinese basket may contain lots of tofu, which may be unimportant for the French. It is not particularly meaningful to speak of arbitrage between cheese prices in France and tofu prices in China.

[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 (Taylor, 2001).

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, in studying the degree of persistence, we avoid the assumption that parity is 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 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 cross-sectional or the temporal aggregation biases. However, we are not able to quantify the relative importance of various biases.

Finally, we use 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 provides us with an opportunity to explore whether departures from his result can be systematically explained as a function of country-pair characteristics.

To summarise, this article uses Big Mac real exchange rates to study these two well-known puzzles. This allows us to avoid several 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 from relying on 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, 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, these papers do not match Big Mac prices with the prices of its underlying ingredients and therefore do not address a range of questions that this article does.

The rest of the article is organised as follows. The next Section provides a 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. The final Section offers some concluding remarks.

2. Data: Sources and Ingredients

The primary data we employ are price observations for the Big Mac and its ingredients, in 34 countries over 13 years (1990–2002). The local currency data for Big Mac prices was obtained from various 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.

[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.

Table 1

Geographic Coverage

Countries and Regions		
Europe	*Western Hemisphere*	*Asia, Pacifi 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)*	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)		

*To correspond with the *Economist*'s Big Mac Index, data for the US are an average of Atlanta, Chicago, San Francisco and Washington, D.C.

The second data set covers city specific local currency prices of various ingredients of the Big Mac – ground beef, bread, labour cost etc. – 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 minimise 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. The *EIU* data comes from the *Worldwide Cost of Living Survey*, which is designed for use by human resource managers in implementing compensation policies (see http://eiu.e-numerate.com/asp/wcol_HelpWhatIsWCOL.asp). We selected local currency prices on the following five traded inputs: ground beef, cheese, lettuce, onions and bread. When there was a choice between a supermarket and an upper-end store, we selected the supermarket price. We also include three non-traded inputs: hourly labour 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 (1990–1), Brazil (1990–4), Mexico (1990–2), and Poland (1990–4). Second, we visually checked the data for possible coding errors via scatter plots. More concretely, we 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 value of the adjacent periods ($t-1$, $t+1$) 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.

In addition, we use data on mean tariff rates from the World Bank's *World Development Indicators 2001*. For each country the tariff data are available twice – one in the early 1990s and another in the late 1990s. We use the first reported value for the years 1990–5, and the most recent value for the years 1996–2002. We also compute the distance between countries, membership in trade and currency unions, bilateral exchange rate volatility, and language, using data from the IMF's *International Financial Statistics*, and from the online CIA *World Factbook*.

3. Digesting the Big Mac

This Section contains the core of our statistical analysis. We proceed in four steps. First, we describe the connection between CPI-based and the Big Mac-based real exchange rates. We conclude that the two 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, employing both nonlinear, and linear specifications. Fourth, using the large cross-sectional dimension of bilateral real exchange rates, we examine the generality of the Engel (1999) puzzle.

3.1. *The Big Mac Versus CPI-based Real Exchange Rates*

In this sub-section, we briefly report several statistical comparisons between Big Mac and CPI-based real exchange rates. The idea is to see if Big Mac real exchange rates are informative about CPI-based real exchange rates. Specific details are available in the online version of the article: http://www.owen.vanderbilt.edu/david.parsley, or http://www.nber.org/~wei. In general, we find that Big Mac real exchange rates are typically highly positively correlated with CPI-based real exchange rates – both in levels, and in first differences. We have checked to make sure that the high correlations are not driven by high-inflationary episodes/countries. We have also implemented seven cointegration tests discussed by Pedroni (1999). All the tests allow individual heterogeneity in short-run dynamics, and all of them reject the null hypothesis of no cointegration. We conclude that there is a stable relationship between the Big Mac and CPI series that we examine. Nonetheless, we make an effort to err on the conservative side and hence restrict our attention to only those bilateral cases where *both* 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 first differences).

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, conditional on currency-specific constants, a 10% undervaluation according the Big Mac real exchange rate in one year is associated with a 3.8% appreciation over the following year.

Taken together, these pieces of information suggest that the behaviour 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.

3.2. *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:

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

Let $P^{\text{BigMac}}_{k,t}$ 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,

$$P^{\text{BigMac}}_{k,t} = \sum_j P_{k,j,t} x_j. \tag{2}$$

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, (2) is an identity.

Suppose we observe $P^{\text{BigMac}}_{k,t}$ 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 $\geq n$), then it is a matter of simple algebra to solve for all x_i, $i = 1,2, \ldots, n$. In fact, a convenient way to solve for $\{x_1, x_2, \ldots, x_n\}$ would simply be a linear regression of $P^{\text{BigMac}}_{k,t}$ on $\{P_{k,j,t}\}$. The regression in this case is not a statistical tool, but an algebraic one. Since (2) is an identity, the $R^2 = 1$. 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'. 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 labour, 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 checks will be explained later.

With these points in mind, we regress the price of a Big Mac on the prices of the eight main inputs, and report the results in Table 2. 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 2 is the importance of non-traded inputs – especially labour – for the price of Big Macs. According to the Table, the total non-traded goods share is at least 55%, i.e., $\alpha = 0.456 + 0.046 + 0.051 \approx 0.55$. Alternatively, if we normalise the non-tradable share by

Table 2

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

	Regression in Levels		Change Regression
	Coefficient Estimates‡	Implied Cost Share (%)†	Coefficient Estimates‡
Ingredient 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)
Non-traded:			
Labour	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%	
No. of observations	318		284
Adjusted R-squared	0.95		0.66

†The share attributed to the *i*th ingredient is computed as: $\hat{\beta}_i \bar{P}_i / \bar{P}_{BigMac}$, where $\bar{P}_i$ is the average price of the *i*th input.

‡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 (first differences) is $\chi^2(8) = 3.3$ (significance level = 0.91)

the total amount explained by all observed inputs, then, non-tradables collectively explain 64% 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.

3.3. *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 'σ-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 σ-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 β-convergence in economic growth empirics. Unlike studies of β-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 1*a*, 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 1*b* presents the same information for its non-traded inputs. The first thing to notice is that the 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. In our sample, the average dispersion in prices of traded inputs declined by 8%, while the dispersion in non-traded input prices rose by 10%. Big Mac price dispersion also rose by 6%. We will return to these results as we interpret our tests of β-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 US 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

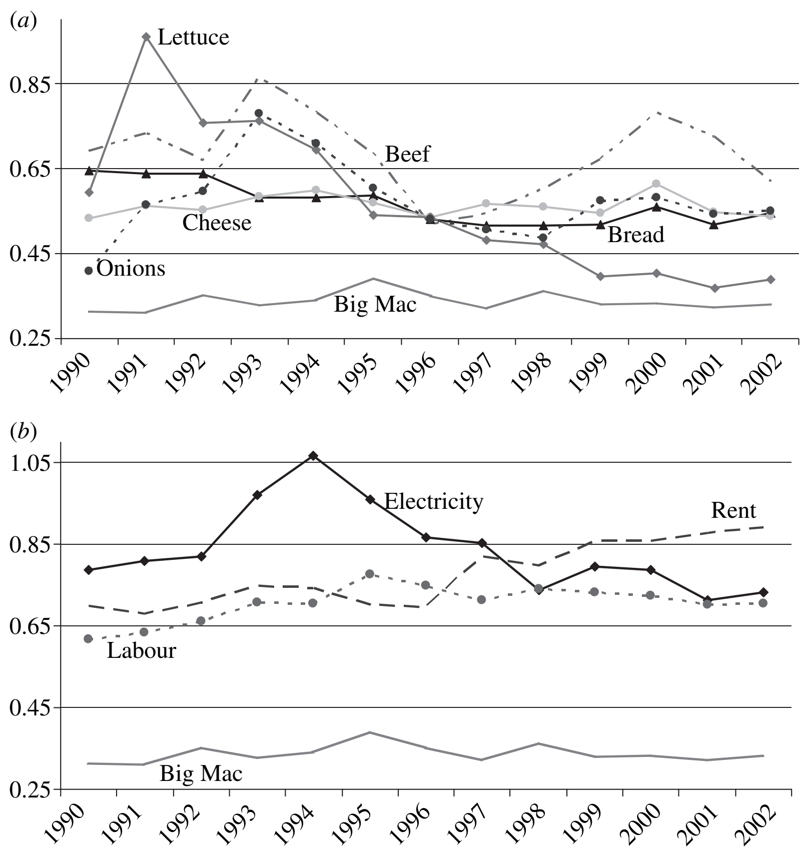

Fig. 1. *a: Price Dispersion in Traded Inputs b: Price Dispersion in Non-traded Inputs*

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 3 by providing OLS estimates of $\hat{\beta}$ from (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 fixed effects are included, and robust standard errors are reported in parenthesis beneath the estimates of $\hat{\beta}$ for each equation. Column (*b*) constrains the country fixed effects to be equal, which as argued by Crucini and Shintani (2004) results in upwardly biased estimates of persistence. Our results support this view. Column (*c*) reports the common correlated effects (CCE) estimator, which as argued by Imbs *et al.* (2005), controls for aggregation bias induced by heterogeneity. Reidel and Szilagyi (2006) however argue that in samples of less than 40 years, this bias is swamped by small sample bias. A comparison of columns (*b*) and (*c*) suggests the (cross-sectional) aggregation bias is not large. In the online version of the article we report several additional specifications (discussed below) which reinforce this assessment. The

Table 3

Persistence Estimates

	(*a*)				(*b*)	(*c*)	(*d*)	(*e*)
	$\hat{\beta}_1$	Half-life	no. obs	$\bar{R}^2$	$\hat{\beta}_1$	$\hat{\beta}_1$	$H_0^a: \lambda_i = 0$	$H_0^b: \lambda_i = 0, \theta_t = 0$
Tradables								
Beef	−0.431	1.2	256	0.17	−0.092	−0.094	2.357	2.100
	(0.001)				(0.075)	(0.025)		
Cheese	−0.451	1.2	252	0.22	−0.111	−0.116	3.074	2.667
	(0.000)				(0.030)	(0.031)		
Lettuce	−0.358	1.6	246	0.13	−0.172	−0.160	1.197	1.449
	(0.064)				(0.037)	(0.036)		
Onions	−0.609	0.7	256	0.27	−0.092	−0.093	4.678	3.576
	(0.000)				(0.035)	(0.034)		
Bread	−0.252	2.4	256	0.08	−0.015	−0.014	1.980	1.734
	(0.011)				(0.018)	(0.018)		
Median		1.2						
Non-Tradables								
Labour	−0.250	2.4	227	0.09	0.004	0.004	1.689	1.746
	(0.013)				(0.007)	(0.007)		
Rent	−0.157	4.1	253	0.03	−0.015	−0.016	1.329	1.318
	(0.127)				(0.021)	(0.021)		
Electricity	−0.177	3.6	256	0.16	−0.029	−0.029	2.780	2.332
	(0.000)				(0.016)	(0.015)		
Median		3.6						
Big Mac	−0.326	1.8	203	0.12	−0.073	−0.079	1.528	1.696
	(0.016)				(0.028)	(0.028)		
Country fixed effects	yes				no	no		
Time fixed effects	yes				yes	no		

Specification (*a*) reports the estimates of $\hat{\beta}_1$ from the following equation: $\Delta q_{i,t} = \beta_1 q_{i,t-1} + \sum \lambda_i \text{country}_i + \sum \theta_t \text{time}_t + \varepsilon_{i,t}$. Specification (*b*) eliminates country fixed effects, and column (*c*) reports the common correlated effects (CCE) estimator, i.e., $\Delta q_{i,t} = \beta_0 + \beta_1 q_{i,t-1} + \beta_2 \bar{q}_t + \beta_3 \bar{q}_{t-1} + \varepsilon_{i,t}$. Half-lives are computed assuming a zero intercept, and the columns labelled H_0^a and H_0^b, report the calculated F-test statistic for the indicated test, with the associated significance levels in parenthesis.

final two columns of the Table report F-tests (with p-values) whether the country fixed effects (column *d*), and whether both country and time fixed effects (column *e*) 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 1*a–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 1*b* 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.[5] Indeed, the median half-life for *Non-tradables* (3.6 years) is three times that for *Tradables* (1.2 years) and the half-life of Big Mac deviations (1.8 years) lies 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 equalised in the long run). While it is reassuring that estimated mean reversion for *Non-tradables* is slower than that for *Tradables*, the magnitude of the half-life is still somewhat surprising. Recall that the analysis of σ-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 3 holds in all of our subsequent regressions.

As a variation of the statistical method, we also estimated persistence via a 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). Again, the same general pattern remains (see the online version for details).[6] Specifically, the half-life of Big Mac deviations is bounded by that of *Tradables* from below, and of *Non-tradables* from above.

To gauge the sensitivity of the results to outliers and alternative specifications, the analysis was repeated in a number of different ways. These robustness checks and extensions are available in greater detail in the online version of the article. We briefly describe them here. First, we excluded observations associated with the largest 5% of the residuals from the corresponding regression in Table 3, and found that nearly all the half-lives rise. Next we estimated persistence using the GMM method (with and without time fixed effects) proposed by Arellano and Bond (1991) that corrects for the small-sample bias inherent in estimation of (3) by OLS. Then, to facilitate an explicit comparison with Cumby (1997), we restricted the sample to only those countries in Cumby's sample and we also restricted the time period to be closer to his by dropping the final three years from our estimation. Fourth, we estimated a regression including one lag of the dependent variable. 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 again, our conclusions about relative convergence speeds are unaffected.

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

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

Finally, we considered the effects of taxes. It is well known that taxes and other transaction costs can create a wedge. Moreover, time variation in these transaction costs can, in effect, present a 'moving target' for mean reversion estimates. As arbitrage may be gauged on either a pre-tax, or tax-inclusive basis, the regressions presented in Table 3 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.[7] Our results are again similar to those in Table 3; the adjustment for VAT and sales taxes seem to matter little for estimated convergence rates.

We now explicitly consider the possibility of nonlinear convergence. Several authors, e.g., O'Connell (1998), Obstfeld and Taylor (1997), Taylor (2001), Sarno and Taylor (2002) and O'Connell and Wei (2002) argue that estimates of 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 the 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 behaviour does not conform to the TAR specification. We consider two such models of nonlinear price adjustment – an 'equilibrium threshold autoregressive model' (or Eq-TAR for short) and a Band-TAR. Both can be represented by the following equation with different parameter restrictions.

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

As we can reject the hypothesis that country fixed effects are zero in the linear specification, 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 centre 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 towards 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 centre 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 characterise behaviour 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

[7] 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 recognise 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 that study, taxes have virtually no effect on their persistence estimates.

model would be appropriate. Currently, there is no consensus as to which model is uniformly 'better' and there is no 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 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 estimates per good. The model with the minimum residual sum of squares is reported in Table 4. For comparison, we present the Eq-Tar and Band-Tar results in two sets of columns. Overall, the estimates of convergence are faster in these nonlinear 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. Obstfeld and Taylor (1997) report thresholds of between 8 and 10% – 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 real exchange rate 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 emphasise that

Table 4

Persistence Estimates Compared (TAR specifications)

	EQ-TAR				Band-TAR			
	$\hat{\beta}$	Threshold	Half-life	no. obs	$\hat{\beta}$	Threshold	Half-life	no. obs
Tradables								
Beef	−0.447	0.024	1.17	237	−0.462	0.024	1.12	237
	(0.058)				(0.061)			
Cheese	−0.500	0.018	1.00	230	−0.488	0.024	1.04	226
	(0.057)				(0.058)			
Lettuce	−0.393	0.051	1.39	207	−0.427	0.063	1.25	198
	(0.062)				(0.065)			
Onions	−0.666	0.030	0.63	237	−0.680	0.030	0.61	237
	(0.065)							
Bread	−0.277	0.018	2.14	233	−0.280	0.018	2.11	233
	(0.052)				(0.053)			
Median			1.17				1.12	
Non-Tradables								
Labour	−0.261	0.009	2.29	214	−0.265	0.009	2.25	214
	(0.054)				(0.057)			
Rent	−0.168	0.018	3.76	228	−0.200	0.049	3.10	208
	(0.043)				(0.049)			
Electricity	−0.188	0.006	3.32	241	−0.183	0.021	3.42	233
	(0.035)				(0.037)			
Median			3.32				3.10	
Big Mac	−0.365	0.015	1.53	181	−0.407	0.018	1.33	176
	(0.065)				(0.072)			

This Table reports estimates of (4) in the text. Standard errors are in parenthesis.

such an inference is possible due to availability of both input and Big Mac prices. While the nature of the data is such that the estimation does not suffer from the problems of time aggregation or unknown index composition, our exercise does not allow us to determine which of these two problems is more severe. We now turn to a decomposition of real exchange rates into parts attributable to tradables and non-tradables, respectively.

3.4. *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.[8] Productivity growth however, is not the only source of movements in the relative price of non-tradables across countries. 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 come under assault. In an influential paper, Engel (1999) concludes that movements in relative prices of non-traded goods appear to account for essentially *none* of the movements in aggregate US-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 not general. As a hint, in a study of the Mexican peso/US dollar real exchange rate, Mendoza (2000) 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 the Mexican case 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 drawback to the existing accounting exercises is the measurement error in characterising the aggregate price index into traded and non-traded subcomponents. One indication of the scope for such measurement error is Engel's estimated weight of the non-traded goods in the six industrialised countries' CPI indexes: it varies from 24% in Italy to 46% in the US, a range that implies substantial variation in CPI baskets across countries. 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 assumes

[8] For textbook treatments, see, Caves *et al.* (2002, pp. 372–3) or Obstfeld and Rogoff (1996, pp. 210–4).

a Cobb-Douglas function without solid foundations. In contrast, for the Big Mac real exchange rates, there is very little room for substitution across inputs. Hence, the decomposition is arguably more straightforward.

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 non-traded 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 US dollar price of a Big Mac in the US. The nominal exchange rate (US 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 expression in the first square bracket is 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.,

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

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 2. 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 non-traded 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 30 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 allows. 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. Since we observe prices (and not price indexes) we construct absolute (i.e., levels) measures of x, y, and q^{BM} 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).[9] Our approach has the advantage that we can systematically

[9] With 34 countries, we have 561 ($=34 \times 33/2$) real exchange rates. However, we continue to focus only on those real exchange rates that are highly correlated with CPI real exchange rates.

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

We construct 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.,

$$x - share_t = \frac{(x_t - \bar{x})^2}{(x_t - \bar{x})^2 + (y_t - \bar{y})^2}. \tag{7a}$$

We label this as 'share in variance' since it most closely approximates Engel's variance decomposition, though (7*a*) preserves the time series dimension. As in Engel's formulation, the denominator 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 2 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 4,500 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 percentage 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

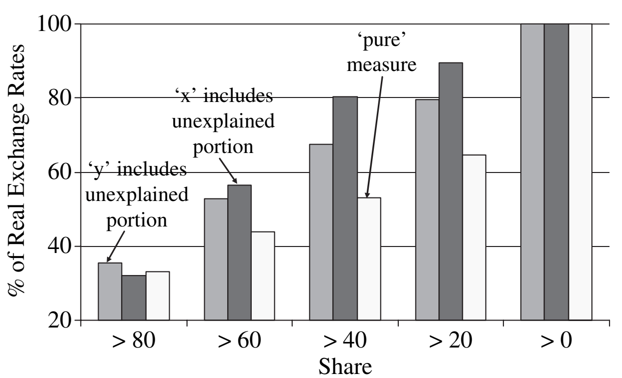

Fig. 2. *Share of Traded Goods Price Deviations in Big Mac Real Exchange Rates* (343 real exchange rates, all years)

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 (7*a*). 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 does not hold generally.

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.[10] We begin by incorporating a dummy variable for the US 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. The results are reported in Table 5. The basic specification includes the three variables

Table 5

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	1.429	1.408	1.523	1.293	1.512	0.773
Volatility	(0.267)	(0.282)	(0.256)	(0.268	(0.237)	(0.254)
$ Peg	−0.415	−0.440	−0.250	−0.279	−0.134	−0.162
	(0.127)	(0.132)	(0.098)	(0.096)	(0.126)	(0.129)
Euro	−0.130	−0.128	−0.065	−0.065	−0.014	0.032
	(0.080)	(0.082)	(0.049)	(0.049)	(0.181)	(0.180)
Distance		0.038		0.041		0.069
		(0.006)		(0.005)		(0.005)
Sum Tariffs		−0.008		−0.008		−0.010
		(0.003)		(0.002)		(0.002)
Common Language		−0.047		0.000		0.012
		(0.027)		(0.021)		(0.022)
European Union		−0.012		0.040		−0.041
		(0.041)		(0.032)		(0.037)
Mercosur		0.245		0.199		0.420
		(0.065)		(0.041)		(0.057)
APEC		0.119		0.011		0.076
		(0.033)		(0.029)		(0.027)
ASEAN		0.183		0.164		0.187
		(0.089)		(0.070)		(0.102)
Nafta		0.000		0.000		−0.071
		(0.000)		(0.000)		(0.069)
Observations	2,304	2,115	2,404	2,214	2,948	2,742
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 (7*a*) in the text. Standard errors are in parenthesis.

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

(*\$peg*, *xrvol*, and *Euro*) plus time and city dummies. In the second 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. This general specification is shown below.

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

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 (for the country-pairs involved) during each year. If both countries i and j are in the same free trade area or customs union (such as within the US, or within the European Union) the value for *tariff* is set to zero. The first two columns of Table 5 (labelled *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 entirely to x, and in the third group of columns, this variation has been attributed entirely 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., it exaggerates the importance of traded goods deviation. Second, having a peg to the US dollar lowers the contribution of deviations from the law of one price to movements in aggregate real exchange rates. 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 as Appendix Table 7 in the online version) 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 (7*b*) is approximately the share of the mean squared error (MSE) of the real exchange rate attributable to x.[11] Here, 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 examine three different measures of (7*b*) depending on how we treat potential covariation between x and y.

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

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

Results using this as the dependent variable (see online Appendix) are largely similar to those using the share in variance as the dependent variable (i.e., Table 5). Namely,

(1) higher exchange rate volatility is associated with a larger *x-share*;
(2) having a peg to the US 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.

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 ($7a$), the new dependent variable becomes:

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

Results using this specification, (and for the MSE) are reported in the online appendix. Statistical significance generally rises using this specification but other conclusions remain qualitatively the same. The only notable differences 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 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 be attributed 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 made possible due to our ability to focus on real exchange rate levels, we repeated all estimations using first-differenced versions of the dependent variables (see online version for details). It continues to be the case that exchange rate variability raises the importance of deviations from the law of one price in real exchange rate movements. However, other conclusions are less apparent in this weaker version of the decomposition. That is, the formerly robust conclusions concerning the dollar peg, distance, 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 (7*a*) and (7*b*) as more informative.

4. Thoughts at the Checkout Counter

This article 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 and CPI-based real exchange rates are generally highly correlated. Our main innovation is to match the Big Mac prices to the prices of individual ingredients (ground beef, bread, lettuce, labour cost, rent etc.) across 34 countries, 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 produce 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.

Our setting is arguably free of a number of possible biases induced by non-comparability of consumption baskets across countries (and the associated arbitrary split between traded and non-traded components), product aggregation bias and time aggregation bias. We find that the speed of convergence for tradable inputs is sufficiently fast to be compatible with economic theories (Chari *et al.*, 2002) and that convergence for the Big Mac real exchange rates is slower than the speed for its tradable inputs but faster than its non-tradable inputs.

By comparing regressions with and without fixed effects, we learn that persistence is higher when constraining all country fixed effects to be the same. Comparing non-linear TAR models with linear regressions, we see that ignoring the non-linearity tends to bias the persistence estimate upward. However, the bias is not big. Finally, we show see that cross-sectional heterogeneity also biases the persistence estimate upward. But this bias is not big either.

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 not correct generally. Reduced exchange rate volatility, exchange rate pegs and tariff barriers generally weaken the Engel effect.

Vanderbilt University
Columbia Business School, CCRF, NBER and CEPR

Submitted: 18 April 2005
Accepted: 24 April 2006

References

Arellano, M., and Bond, S. (1991). 'Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations', *Review of Economics Studies*, vol. 58 (April), pp. 277–97.

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

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

Chari, V.V., Kehoe, P. and McGratten, E. (2002). 'Can sticky price models generate volatile and persistent real exchange rates?', *Review of Economic Studies*, vol. 69 (July), pp. 533–63.

Click, R. W. (1996). 'Contrarian MacParity', *Economics Letters*, vol. 53 (November), pp. 209–12.

Crucini, M. and Shintani, M. (2004). 'Persistence in law-of-one-price deviations: evidence from micro-data', Vanderbilt University Working Paper.

Cumby, R. (1997). 'Forecasting exchange rates and relative prices with the hamburger standard: is what you want what you get with McParity?', NBER Working Paper no.5675.

Dornbusch, R. (1989). 'Real exchange rates and macroeconomics: a selective survey', *Scandinavian Journal of Economics*, vol. 91 (June), pp. 401–32.

Engel, C. (1999). 'Accounting for U.S. real exchange rate changes', *Journal of Political Economy*, vol. 107 (June), pp. 507–38.

Frankel, J. and Rose, A. (1996). 'A panel project on purchasing power parity: mean reversion within and between countries', *Journal of International Economics*, vol. 40 (February), pp. 209–24.

Franses, P. and van Dijk, D. (2000). *Nonlinear Time Series Models in Empirical Finance*, Cambridge: Cambridge University Press.

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

Froot, K. and Rogoff, K. (1991). 'The EMS, the EMU, and the transition to a common currency', *NBER Macroeconomics Annual*, vol. 6, pp. 269–317.

Froot, K. and Rogoff, K. (1995). 'Perspectives on PPP and long-run real exchange rates', in (H. Grossman and K. Rogoff, eds.), *Handbook of International Economics*, Vol. 3, pp. 1647–88, New York: Elsevier.

Greene, W. (1997). *Econometric Analysis*, Third edition, New Jersey: Prentice Hall.

Imbs, J., Mumtaz, H, Ravn, M. and Rey, H. (2005). 'PPP strikes back: aggregation and the real exchange rate', *Quarterly Journal of Economics*, vol. 120 (February), pp.1–44.

Lutz, M. (2001). 'Beyond burgernomics and MacParity: exchange rate forecasts based on the law of one price', University of St. Gallen working paper.

Mendoza, E. (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, P. (1998). 'Market frictions and real exchange rates', *Journal of International Money and Finance*, vol. 17 (February), pp. 71–95.

O'Connell, P. and Wei, S. (2002). 'The bigger they are, the harder they fall: how price differences across U.S. cities are arbitraged', *Journal of International Economics*, vol. 56 (January), pp. 21–53.

Obstfeld, M. and Rogoff, K. (1996). *Foundations of International Macroeconomics*, Cambridge, MA: MIT Press.

Obstfeld, M. and Rogoff, K. (2000). 'The six major puzzles in macroeconomics: is there a common cause?', *NBER Macroeconomics Annual*, vol. 15, pp. 339–90.

Obstfeld, M. and Taylor, A. (1997). 'Non-linear aspects of goods-market arbitrage and adjustment: Heckscher's commodity point revisited', *Journal of Japanese and International Economics*, vol. 11 (December), pp. 441–79.

Ong, L. (1997). 'Burgernomics: the economics of the Big Mac standard', *Journal of International Money and Finance*, vol. 16 (December), pp. 865–78.

Pakko, M. and Pollard, P. (1996). 'For here or to go? purchasing power parity and the Big Mac', *Federal Reserve Bank of St. Louis Review*, vol. 78 (January-February), pp. 3–22.

Parsley, D. and Wei, S. (1996). 'Convergence to the law of one price without trade barriers or currency fluctuations', *Quarterly Journal of Economics*, vol. 111(November), pp. 1211–36.

Parsley, D. and Wei, S. (2003). 'Currency arrangements and market integration: a price based approach', Vanderbilt University working paper.

Pedroni, P. (1999). 'Critical values for cointegration tests in heterogeneous panels with multiple regressors', *Oxford Bulletin of Economics and Statistics*, vol. 61(1) (November), pp. 653–70.

Reidel, D. and Szilagyi, J. (2006). 'A biased view of PPP', Harvard University working paper.

Rogoff, K. (1996). 'The purchasing power parity puzzle', *Journal of Economic Literature*, vol. 34 (June), pp. 647–68.

Sala-i-Martin, X. (1996). 'Regional cohesion: evidence and theories of regional growth and convergence', *European Economic Review*, vol. 40 (June), pp. 1325–52.

Sarno, L. and, Taylor, M. (2002). 'Purchasing power parity and the real exchange rate', *International Monetary Fund Staff Papers*, vol. 49 (April), pp. 65–105.

Sarno, L. and Taylor, M. (2004). 'Non-linear dynamics in deviations from the law of one price: a broad-based empirical study', *Journal of International Money and Finance*, vol. 23 (February), pp. 1–25.

Taylor, A. (2001). 'Potential pitfalls for the purchasing-power parity puzzle? Sampling and specification biases in mean-reversion tests of the law of one price', *Econometrica*, vol. 69 (March), pp. 473–98.