Real Exchange Rates and Primary Commodity Prices*

Joao Ayres
Inter-American Development Bank

Constantino Hevia
Universidad Torcuato Di Tella

Juan Pablo Nicolini
Federal Reserve Bank of Minneapolis and Universidad Torcuato Di Tella

September 26, 2019

Abstract

In this paper, we show that there is substantial comovement between prices of primary commodities such as oil, aluminum, maize, or copper and real exchange rates between developed economies such as Germany, Japan, and the United Kingdom against the US dollar. We therefore explicitly consider the production of commodities in a two-country model of trade with productivity shocks and shocks to the supply of commodities. We calibrate the model so as to reproduce, among other things, the volatility and persistence of primary commodity prices and show that it delivers equilibrium real exchange rates that are as volatile and persistent as they are in the data. The model rationalizes an empirical strategy to identify the fraction of the variance in real exchange rates that can be accounted for by the underlying shocks, even if those shocks are not observable. We use this strategy to argue that shocks that move primary commodity prices account for a large fraction of the volatility of real exchange rates in the data. Our analysis implies that existing models used to analyze real exchange rates between large economies that mostly focus on trade between differentiated final goods could benefit, in terms of matching the behavior of real exchange rates, by also considering trade in primary commodities.

Keywords: primary commodity prices, real exchange rate disconnect puzzle.
JEL classification: F31, F41.

*We thank the coeditor Martin Uribe and three referees, as well as Manuel Amador, Roberto Chang, Juan Carlos Conesa, Deepa Datta, Jonathan Heathcote, Patrick Kehoe, Tim Kehoe, Albert Marcet, Enrique Mendoza, Martin Sola, and Esteban Rossi-Hansberg for comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the Inter-American Development Bank.

1 Introduction

This paper provides empirical evidence that points toward common factors that move a handful of primary commodity prices and real exchange rates between developed economies. Specifically, we study the behavior of the bilateral real exchange rate (RER) of Germany, Japan, and the United Kingdom with the United States for the period 1960-2014. A rough summary of the results is that shocks that move just four primary commodity prices (PCP) can account for between one-third and one-half of the volatility of the RER between the United States and those three countries.

The relevance of these results is highlighted by the so-called *exchange rate disconnect puzzle*: the fact that real exchange rates across developed economies are very volatile, very persistent, and very hard to relate to fundamentals.[1] This difficulty opened the door for theoretical explorations of models with nominal rigidities as the source of RER movements, as in, for example, Chari, Kehoe and McGrattan (2002). We will ignore nominal rigidities in our analysis and explore how far one can go with shocks that affect the relative prices of the main primary commodities.

The disconnect puzzle is not present in small open economies in which exports of a few primary commodities are a sizable share of total exports.[2] For countries such as Australia, Chile, and Norway, changes in the international prices of the commodities each country exports are highly correlated with their real exchange rates. As we show, a very similar idea can go a long way toward explaining movements in the RER among developed economies.

The idea that we exploit in the paper is very simple: fluctuations in the prices of commodities affect manufacturing costs and, thus, manufacturing prices, which in turn induce changes in the costs of final goods. These cost fluctuations translate into price fluctuations at the country level. If changes in commodity prices have differential effects on the domestic costs of any two countries, primary commodity price changes will affect the real exchange rate between those two countries.

We must make a caveat explicit at this point. In the small open economy literature, it has been customary to assume that the movements in PCP are exogenous to the country. This assumption is reasonable for small countries and simplifies the analysis considerably. When dealing with large economies, such as the ones we analyze in this paper, PCP and RERs are determined endogenously as functions of the state of the economy. This approach complicates the analysis somewhat, but, we argue, the problem is still solvable. We address

[1] See, for example, Meese and Rogoff (1983), Engel (1999), Obstfeld and Rogoff (2001), and Betts and Kehoe (2004).

[2] See Chen and Rogoff (2003) and Hevia and Nicolini (2013).

this issue in two different ways.

We first consider a very simple model in which the exogenous driving forces are productivity shocks and shocks to the supply of commodities. We calibrate the production functions so that the shares of commodities value added in the model are as low as they are in the data, and we calibrate the shocks to the supply of primary commodities so as to match the volatility and persistence of the PCP in the data. We then show that the model can go a long way toward reproducing the volatility and persistence of the RER. We also show that, as we make the shares of commodities in the production function very small, the volatility and persistence of the RER drop dramatically.

We then use the model as an example of how the coefficient of determination in log-linear regressions can be informative regarding the common shocks that move both PCP and RER, even if those shocks are unobservable. In the model, we take a stand on what these shocks are. In the empirical exploration, we are agnostic with respect to which are the underlying shocks and we treat them as unobservable.

Relating PCP changes to RER changes is a promising avenue to explore for several reasons. First, PCP are highly volatile (even more volatile than real exchange rates, as we show below) and very persistent, a feature that, as we mentioned, real exchange rates also exhibit. Second, the share of trade in primary commodities in total world trade is far from trivial: total trade in a few commodities (10) accounts for between 12% to 18% of total world trade in goods, depending on the year chosen.[3] This number clearly underestimates the true share of commodities, since trade shares are not value-added measures. Thus, when steel is exported, it is fully counted as a manufactured good, even though an important component of its cost depends on iron. The same happens when a car is exported. Third, primary commodities are at the bottom of the production chain, so they directly affect final goods prices.[4] In addition, they may directly affect the prices of other domestic inputs—such as some types of labor and services in general—that are used jointly with primary commodities in the production of intermediate goods, and thus they may indirectly affect the costs of final goods. Because just a few commodities make up a high share of total trade, we need to focus on only a handful of prices. Finally, it is well known that the law of one price on those primary commodities holds, so no ambiguity with respect to their tradability exists. The literature on RER has struggled to separate the set of final goods into two categories: the ones that are traded, for which the law of one price is assumed to hold, and the ones

[3]Total trade is close to 12% in 1990 and 18% in 2012. The main difference is that the first year is one of particularly low primary commodity prices, and the second is one of particularly high prices.

[4]This direct effect is substantial enough for monetary authorities in developed countries to focus attention on measures of "core" inflation, which abstract from the "volatile" effect of primary commodity prices (food and energy).

that are not. We only need to assume that for the few commodities we analyze, the law of one price holds, and it is precisely for these prices that independent evidence that the law of one price holds is the strongest.

The exchange rate disconnect puzzle has been widely studied in the literature. Two recent attempts at quantitatively explaining several facts related to the puzzle are Itskhoki and Mukhin (2017) and Eaton, Kortum and Neiman (2016). They provide very good descriptions of the state of the literature. The connection between RER and PCP has recently been ignored for the countries we focus on, and not only on the empirical side.[5] Ever since the seminal contribution of Obstfeld and Rogoff (1995), the theoretical literature developed to study RER between the countries we consider in this paper has focused exclusively on the production and trade of final goods.[6] This tradition eventually was also followed even when studying small open economies, following Galí and Monacelli (2005).

That direction taken by the literature is in contrast with earlier analysis of real exchange determinants. Côté (1987) provides a discussion of a mechanism by which real exchange rates and primary commodity prices jointly respond to shocks; that mechanism is very close to the workings of the model we describe in Section 2. On the empirical side, Sachs (1985) and Dornbusch (1985*a*,*b*, 1987) are important contributions. These early attempts, however, have heretofore been ignored in the literature. Our evidence supports the insights of that early literature and suggests that theoretical models of RER among developed economies that ignore primary commodity markets may fall short of providing a comprehensive explanation of RER movements.

In Section 2, we describe the model and derive an equilibrium relationship between RER and PCP. We highlight the fact that in equilibrium, both are jointly determined, so there is no sense in which one can cause the other. In Section 3, using the model as an example, we describe how the R^2 of a regression can help us to identify the fraction of the volatility of the RER that can be accounted for by shocks that also move PCP. In Section 4, we describe the data and show some of its moments, and in Section 5, we present the empirical results. The calibration and simulation of the model are presented in Section 6, and Section 7 concludes.

2 Model

We develop a simple general equilibrium model with the aim of illustrating a mechanism through which *equilibrium* fluctuations in primary commodity prices are associated with

[5] A recent empirical contribution is Chen, Rogoff and Rossi (2010).

[6] A notable exception is the work of Backus and Crucini (2000), who document a close connection between oil prices and the terms of trade among developed economies.

significant movements in the bilateral real exchange rates between two large economies. The model features two large countries and the rest of the world and three sectors of production: nontradable final goods, tradable intermediate goods, and tradable primary commodities. We depart from the traditional two-country setting by including a third economy (the rest of the world), which is used as a device to generate excess demands for primary commodities that can be shocked to generate volatile and persistent commodity prices, as observed in the data. When calibrated to mimic certain moments of commodity prices and aggregate, sectoral, and trade data for the United States and Japan, the model goes a long way toward explaining the volatility of the US-Japan bilateral real exchange rate.

All the technologies in the model are Cobb-Douglas. The Cobb-Douglas production functions allow us to obtain a simple log-linear relation with constant coefficients between the bilateral real exchange rate and primary commodity prices that must hold in equilibrium. Although counterfactual for issues such as structural transformations, those simplifying assumptions are good enough for our purposes.[7] In addition, and in contrast with the small open economy literature which takes commodity prices as given, the general equilibrium approach makes transparent the endogeneity of both the RER and PCP and how they relate to each other.

Time is discrete and denoted by $t = 0, 1, 2, ...$ Households consume a single nontradable final good. Households in countries 1 and 2 inelastically supply their endowments of labor and commodity-specific fixed endowments of natural resources, whereas households in country 3 (the rest of the world) inelastically supply their fixed endowment of labor and stochastic endowments of primary commodities.

The model abstracts from capital accumulation. The main reason for doing so is to highlight that the mechanism that we exploit is static, as we explained in the introduction, and this mechanism would persist if one were to allow for capital accumulation in the model. For the same reason, we impose financial autarky, which implies that there are no intertemporal decisions to be made and that trade is balanced on a period-by-period basis. It therefore follows that equilibrium prices and quantities are determined by the production side of the economy independently of households' preferences. Since we are not concerned with normative issues, we ignore the preferences side of the model.

Throughout the paper, a superscript in a variable is used to denote a given country and a subscript refers to a particular good. For example, $x^1_{3,t}$ is the demand for commodity 3 by country 1 at time t.

[7] In Ayres, Hevia and Nicolini (2017), we show how our results generalize to a more general multi-country model with arbitrary constant returns to scale technologies, multiple types of labor inputs, intermediate goods, and primary commodities.

Country $i = 1, 2$ produces a final nontradable good, Y_t^i, using labor and two intermediate tradable goods, one of which is produced in country 1, Q_t^1, and the other is produced in country 2, Q_t^2. The technology to produce the intermediate good Q_t^i uses labor and three tradable primary commodities. In addition, country i is able to produce two of the three primary commodities (commodities i and 3) using labor and a commodity-specific fixed endowment of natural resources. The production functions are given by

$$
\begin{aligned}
Y_t^i &= Z_t^i \left(q_{1,t}^i\right)^{\alpha_1^i} \left(q_{2,t}^i\right)^{\alpha_2^i} \left(n_{y,t}^i\right)^{\alpha_3^i}, \\
Q_t^i &= Z_t^i \left(x_{1,t}^i\right)^{\beta_1^i} \left(x_{2,t}^i\right)^{\beta_2^i} \left(x_{3,t}^i\right)^{\beta_3^i} \left(n_{q,t}^i\right)^{\beta_4^i}, \\
X_{i,t}^i &= Z_t^i \left(e_{i,t}^i\right)^{1-\phi_i^i} \left(n_{x_i,t}^i\right)^{\phi_i^i}, \\
X_{3,t}^i &= Z_t^i \left(e_{3,t}^i\right)^{1-\phi_3^i} \left(n_{x_3,t}^i\right)^{\phi_3^i},
\end{aligned}
$$

where Z_t^i denotes total factor productivity (TFP) in country i, common across sectors; $q_{1,t}^i$ and $q_{2,t}^i$ are the inputs of intermediate goods used to produce final goods; $x_{j,t}^i$ for $j = 1, 2, 3$ are the inputs of primary commodities used to produce the intermediate good; $n_{y,t}^i$, $n_{q,t}^i$, $n_{x_i,t}^i$, and $n_{x_3,t}^i$ are the labor inputs allocated to each sector; $X_{i,t}^i$ and $X_{3,t}^i$ denote the production of commodities i and 3; and $e_{1,t}^i$ and $e_{3,t}^i$ are the commodity-specific fixed endowments of natural resources. All the production functions are constant returns to scale.

Country 3 (the rest of the world) is different. As noted earlier, country 3 is a device to generate volatile and persistent commodity prices. As such, the production and consumption structure of this country is simpler than that of countries 1 and 2. Country 3 receives a stochastic endowment of the three primary commodities, $X_{1,t}^3$, $X_{2,t}^3$, $X_{3,t}^3$, and produces and consumes a nontradable final good made of commodities with a technology represented by

$$
Y_t^3 = \left(x_{1,t}^3\right)^{\pi_1} \left(x_{2,t}^3\right)^{\pi_2} \left(x_{3,t}^3\right)^{\pi_3} \left(n_{y,t}^3\right)^{\pi_4},
$$

where $\pi_i > 0$ for $i = 1, 2, 3, 4$, $\pi_1 + \pi_2 + \pi_3 + \pi_4 = 1$, and $n_{y,t}^3$ denotes the labor input.

We interpret the stochastic nature of the endowment of commodities of country 3 as capturing all world contingencies that affect primary commodity markets, such as the weather, natural disasters, monopolistic behavior by the OPEC, and so on.

The source of uncertainty in the model is represented by the two aggregate productivity shocks in countries 1 and 2 and by the stochastic endowment of the three primary commodities of country 3. We assume the following AR(1) processes:

$$\begin{aligned}
\ln\left(Z_t^1\right) &= \left(1-\rho^{z_1}\right)\ln\left(Z^1\right)+\rho^{z_1}\ln\left(Z_{t-1}^1\right)+\varepsilon_t^{z_1},\\
\ln\left(Z_t^2\right) &= \left(1-\rho^{z_2}\right)\ln\left(Z^2\right)+\rho^{z_2}\ln\left(Z_{t-1}^2\right)+\varepsilon_t^{z_2},\\
\ln\left(X_{1,t}^3\right) &= \left(1-\rho^{x_1^3}\right)\ln\left(X_1^3\right)+\rho^{x_1^3}\ln\left(X_{1,t-1}^3\right)+\varepsilon_t^{x_1},\\
\ln\left(X_{2,t}^3\right) &= \left(1-\rho^{x_2^3}\right)\ln\left(X_2^3\right)+\rho^{x_2^3}\ln\left(X_{2,t-1}^3\right)+\varepsilon_t^{x_2},\\
\ln\left(X_{3,t}^3\right) &= \left(1-\rho^{x_3^3}\right)\ln\left(X_3^3\right)+\rho^{x_3^3}\ln\left(X_{3,t-1}^3\right)+\varepsilon_t^{x_3},
\end{aligned}$$

where the vector of innovations $[\varepsilon_t^{z_1},\varepsilon_t^{z_2},\varepsilon_t^{x_1},\varepsilon_t^{x_2},\varepsilon_t^{x_3}]$ is normally distributed with a zero mean and an arbitrary covariance matrix, and variables without time subscripts are the respective means. Note that we assume that productivity is the same across sectors within each country. This assumption is without loss of generality, as shocks to TFP will have negligibly small effects on the bilateral real exchange rate in this model. We leave the complete description of the model and the details of the computation of the equilibrium to Appendix B.

We turn next to discuss the determination of the bilateral real exchange rate and its equilibrium relationship to commodity prices and the other shocks in the model. The bilateral real exchange rate is defined as $\left(P_t^{y_1}E_t^{1,2}\right)/P_t^{y_2}$, where $P_t^{y_i}$ is the final good price index in country $i=1,2$, and $E_t^{1,2}$ transforms units of account in country 1 into units of account in country 2. We let ξ_t be the logarithm of the real exchange rate between countries 1 and 2, which is then given by

$$\xi_t \equiv \ln\left(\frac{P_t^{y_1}E_t^{1,2}}{P_t^{y_2}}\right)=\left(p_t^{y_1}+e_t^{12}\right)-p_t^{y_2},$$

where $x=\ln(X)$.

Perfect competition implies that prices equal marginal costs in all markets. Since all technologies are Cobb-Douglas, all marginal costs are Cobb-Douglas functions of their input prices as well. Thus, for example, for country 1, we obtain

$$P_t^{y_1}=\frac{1}{Z_t^1}\left(\frac{P^{q_t^1}}{\alpha_1^1}\right)^{\alpha_1^1}\left(\frac{P^{q_t^2}}{\alpha_2^1}\right)^{\alpha_2^1}\left(\frac{W_t^1}{\alpha_3^1}\right)^{\alpha_3^1}$$

In a similar fashion, intermediate good prices $P^{q_t^1}$ and $P^{q_t^2}$ are also Cobb-Douglas functions of commodity prices and wages. But if we use the cost minimization conditions of the commodity-producing sectors, wages can be written as exponential functions of the commodity prices and the prices of the endowments. Using the law of one price for the commodities, we can write the term $(p_t^{y_1}+e_t^{12})$ as a linear function of the logarithms of the commodity prices and the prices of the endowments, all measured in units of account in country 2.[8]

By applying a similar logic to the price level in country 2, it is possible to write

$$\xi_t = \gamma_{z_1} z_t^1 + \gamma_{z_2} z_t^2 + \gamma_{e_1} p_t^{e_1} + \gamma_{e_2} p_t^{e_2} + \delta_{x_1} p_t^{x_1} + \delta_{x_2} p_t^{x_2} + \delta_{x_3} p_t^{x_3}, \tag{1}$$

where z_t^1 and z_t^2 denote the log of aggregate productivity in countries 1 and 2; $p_t^{x_1}$, $p_t^{x_2}$, and $p_t^{x_3}$ are the log prices of the primary commodities; $p_t^{e_1}$ and $p_t^{e_2}$ are the log prices of the natural resources in countries 1 and 2; and the coefficients multiplying commodity prices are given by

$$\begin{aligned}
\delta_{x_1} &= \left(\alpha_1^1 - \alpha_1^2\right)\beta_1^1 + \left(\alpha_2^1 - \alpha_2^2\right)\beta_1^2 + \frac{\left(\alpha_1^1 - \alpha_1^2\right)\beta_4^1 + \alpha_3^1}{\phi_1^1}, \\
\delta_{x_2} &= \left(\alpha_1^1 - \alpha_1^2\right)\beta_2^1 + \left(\alpha_2^1 - \alpha_2^2\right)\beta_2^2 + \frac{\left(\alpha_2^1 - \alpha_2^2\right)\beta_4^2 - \alpha_3^2}{\phi_2^2}, \\
\delta_{x_3} &= \left(\alpha_1^1 - \alpha_1^2\right)\beta_3^1 + \left(\alpha_2^1 - \alpha_2^2\right)\beta_3^2.
\end{aligned}$$

As we made clear before, in equilibrium, all the prices on the right-hand side and the real exchange rate move together, so the coefficients just described do not measure the correlation between any given commodity price and the real exchange rate. But this equation makes explicit that the real exchange rate will be correlated with prices of traded goods such as $p_t^{x_1}, p_t^{x_2}$, and $p_t^{x_3}$ but will also be correlated with prices of nontraded goods, such as $p_t^{e_1}$ and $p_t^{e_2}$. This reflects the mechanism discussed in the introduction: final good prices have a traded component and a nontraded component. The relationship between the traded components and the real exchange rate depend on the asymmetries across countries. Notice that if the countries are the same, in the sense that $\alpha_1^1 = \alpha_1^2$ and $\alpha_2^1 = \alpha_2^2$, several terms in those coefficients do become zero. Thus, these coefficients are further away from zero when the two countries have different production structures. When we calibrate the model to the United States and Japan in Section 6, however, we indeed find large asymmetries in factor shares between the two countries. Even if the countries had some symmetries, however, the real exchange rate will still be correlated with the prices of nontradable goods, $p_t^{e_1}$and $p_t^{e_2}$, which themselves will be correlated with PCP.

A similar decomposition has been used by Crucini and Landry (2019) to analyze a large micro dataset for OECD countries. They show that the nontraded component of final good prices accounts for a large fraction of the movements of the real exchange rates. That evidence is fully consistent with the static mechanism that operates in our model. For

[8] Appendix B contains the derivation of equation (1) and shows the formulas for the coefficients γ_{z_1}, γ_{z_2}, γ_{e_1}, and γ_{e_2}.

simplicity, they assume a certain symmetry across countries, which in our context would be similar to imposing $\alpha_1^1 = \alpha_1^2$ and $\alpha_2^1 = \alpha_2^2$, a feature that shuts down a direct channel that generates the correlation between the RER and the PCP.

Equation (1) is one of many possible representations of the real exchange rates that must hold in equilibrium. Other representations, which must also hold in equilibrium, can be derived by substituting different equilibrium conditions into the definition of the real exchange rate. For example, we could have written the real exchange rate in terms of commodity prices and labor allocations rather than natural resource prices (see Ayres, Hevia and Nicolini (2017) for other representations of the real exchange rate). In addition, note that, except for the aggregate productivity shocks z_t^1 and z_t^2, all the other variables in equation (1) are endogenous and, therefore, correlated with the productivity shocks and the other variables. Hence, there is no hope of obtaining consistent estimates of the δ parameters by running a regression of real exchange rates on commodity prices. We now proceed to discuss this issue in detail and explain our empirical methodology.

3 Empirical methodology

In the model outlined above, we took a stand on the set of structural shocks that drive fluctuations in prices and quantities. The model implies a relationship between the real exchange rate and a set of primary commodity prices and unobserved variables (such as productivity shocks and prices of natural resources) that holds in equilibrium. We calibrate and analyze the quantitative implications of this model in Section 6.

In contrast, in this section we describe an empirical procedure to capture such a relationship but are agnostic about the particular model and structural shocks that drive fluctuations in real exchange rates and commodity prices in the data. The procedure is then used in the next section to measure how much of the variability in real exchange rates can be explained by shocks that also affect the prices of primary commodities.

To be specific, let us consider a regression of the bilateral real exchange rate between the United States and United Kingdom on a set of primary commodity prices:

$$\xi_t^{USA,UK} = \eta' \mathbf{p}_t^{X,USA} + v_t. \tag{2}$$

The left-hand side of equation (2) is the log of the bilateral real exchange rate, $\xi_t^{USA,UK} = \ln P_t^{USA} - \ln P_t^{UK} + \ln S_t$, where P_t^{USA} denotes the price level in the United States, P_t^{UK} denotes the price level in the United Kingdom, and $E_t^{USA,UK}$ denotes the nominal exchange rate between US dollars and British pounds. On the right-hand of the equation, $\mathbf{p}_t^{X,USA}$ is a vector of (log) primary commodity prices normalized by the US price level, η is a vector of coefficients, and v_t is an error term capturing all of the unobserved variables.[9] The model in Section 2 delivers an expression for the real exchange rate that is equivalent to equation (2).

We next argue that, even though in any general equilibrium model the real exchange rate and primary commodity prices in equation (2) are determined simultaneously, the R^2 of the regression still contains valuable information.

Suppose that, in a particular model, there are m commodity prices, $\mathbf{p}_t^{X,USA} \in \mathbf{R}^m$, and the state of the economy is represented by a vector $\omega_t \in \mathbf{R}^n$. The state vector may include endogenous state variables, such as stocks of capital, and exogenous state variables, such as productivity, endowment, and policy shocks. In such a model, the equilibrium values for the RER and the PCP are (possibly nonlinear) functions of the state variables.[10] A linear approximation to those functions implies

$$\begin{aligned} \xi_t^{USA,UK} &= \theta' \omega_t, \\ \mathbf{p}_t^{X,USA} &= \Omega \omega_t, \end{aligned} \tag{3}$$

where $\theta \in R^n$, Ω is an $m \times n$ matrix, and variables are measured as deviations from their long-run means. We treat ω_t as unobserved, so we can interpret the state variables as orthogonal with an identity covariance matrix without loss of generality.[11]

Consider projecting the real exchange rate onto the commodity prices,

$$\text{Proj}(\xi_t^{USA,UK} | \mathbf{p}_t^{X,USA}) = \eta' \mathbf{p}_t^{X,USA}.$$

Equation (3) and the orthogonality principle imply $\eta' = (\theta'\Omega')(\Omega\Omega')^{-1}$. It then follows that projecting $\xi_t^{USA,UK}$ onto $\mathbf{p}_t^{X,USA}$ is equivalent to decomposing the real exchange rate into two orthogonal components:

$$\xi_t^{USA,UK} = \eta'\Omega\omega_t + (\theta' - \eta'\Omega)\omega_t. \tag{4}$$

The first term of the projection measures how much of the variability in the real exchange rate can be accounted for by fundamental shocks that affect primary commodity prices. The second term of the projection is orthogonal to the first and measures how much of the variability in the real exchange rate is accounted for by fundamental shocks that do not manifest themselves as fluctuations in commodity prices correlated with the real exchange

[9]Since we use PCP in constant dollars, one might be worried that the US price level enters both sides of the equation. The results, however, do not depend on that normalization (Ayres, Hevia and Nicolini, 2017).

[10]In the model described in Section 2, those are the aggregate productivity shocks and the stochastic endowment of commodities in country 3.

[11]If the shocks ω_t have a non-diagonal covariance matrix $E(\omega_t\omega_t') = \Sigma$, we have an observationally equivalent system with orthogonal state variables by letting $\tilde{\omega}_t = \Sigma^{-1/2}\omega_t$, $\tilde{\theta}' = \theta'\Sigma^{1/2}$, and $\tilde{\Omega} = \Omega\Sigma^{1/2}$.

rate. In terms of this decomposition, the R^2 of the regression (2) can be written as

$$R^2 = \frac{E[\eta'\Omega\omega_t\omega_t'\Omega'\eta]}{E[(\theta'\omega_t\omega_t'\theta)]} = \frac{\eta'\Omega\Omega'\eta}{\theta'\theta}. \tag{5}$$

The underlying (implicit) assumption in much of the literature on bilateral real exchange rates between developed countries is that the component associated with commodity prices can be safely ignored. We can express this no-relevance-of-commodities assumption as the requirement that the R^2 of the regression of real exchange rates on commodity prices is zero, which is true whenever $\eta'\Omega = 0$.

Let us split the state variables as $\omega_t = [\omega_{1t}' \ \omega_{2t}']'$, so that $\xi_t^{USA,UK} = \theta_1'\omega_{1t} + \theta_2'\omega_{2t}$ and $\mathbf{p}_t^{X,USA} = \Omega_1\omega_{1t} + \Omega_2\omega_{2t}$. It then follows that $\eta'\Omega = \theta_1'\Omega_1 + \theta_2'\Omega_2$. A necessary and sufficient condition for the R^2 of the regression (2) to be zero is thus $\theta_1'\Omega_1 = -\theta_2'\Omega_2$. This equality holds, for example, when $\theta_1 = 0$ and $\Omega_2 = 0$. This implies a block-recursive structure in which the set of state variables that determine the real exchange rate are different from (and orthogonal to) those that determine the primary commodity prices. If these conditions do not hold, then commodity prices will be (generically) correlated with the real exchange rate.

Before proceeding, we would like to be very clear regarding the interpretation of our results, as will become evident from the previous discussion. In no sense do we mean that primary commodity prices cause real exchange rates. As our model makes clear and the previous discussion re-emphasizes, both PCP and RER are endogenous variables that respond to a vector of underlying structural shocks. We envision that vector to include weather shocks, wars that disrupt the production of oil and metals, world recessions, disagreement among OPEC members, discoveries of mineral and oil fields, and many others. The empirical evidence we provide below implies that the total variance of the shocks that are common to both RER and PCP is high, relative to the total variance of both RER and PCP.

3.1 Higher-order terms and time-varying coefficients

Linear approximations work well when the shocks are small. The large and persistent movements in RER and PCP suggest that the approximation error may be large. One way out of this condition could be to add nonlinear terms in the regression described in equation (2). We opted for simplicity and consider only the linear terms. Since adding variables can only increase the R^2, we can interpret our results as a lower bound on the fraction of the volatility of the RER that can be accounted for by shocks that also move PCP.

Another concern is that the coefficients in the linearized system (3) are evaluated at the equilibrium around which the linearization is made. We will apply our empirical procedure to time series that are over half a century long. The economies we consider have all experienced major transformations in both their production structures and the trade patterns during that period. As a first pass, we ignore this important issue and pretend that the coefficients do indeed remain fixed during the whole period. But we also redo the analysis dividing the sample into subperiods of about 12 years to capture possible time variation in the regression coefficients η. We also use this idea to motivate the exercises on the out-of-sample fit that we do in Subsection 5.1.

4 Data

We collected monthly data on consumer price levels and nominal exchange rates for the United States, Germany, the United Kingdom, and Japan. The bilateral RERs are defined as the nominal exchange rates between Germany (DEU), Japan (JPN), and the United Kingdom (UK) against the US dollar multiplied by the ratio of CPIs. In the case of Germany, we use the mark until 2000 and the euro thereafter. We also collected price data for the 10 primary commodities with the largest shares in world trade in 1990 and for which monthly prices are available from January 1960 to December 2014. Appendix A describes the data. Throughout the paper, we show results for the data in four-year differences but show in the Online Appendix that all results hold using the data in levels.[12]

Table 1 shows the volatility (standard deviation) of the monthly data in four-year differences on the US bilateral (log) real exchange rates against the United Kingdom, Germany, and Japan between 1960 and 2014, as well as for four subperiods.[13] We also report the average volatility (simple and trade weighted) of the (log) prices of the commodities listed in Table A.1. As can be seen, the volatility of PCP is substantially higher than that of RER.[14]

In addition, it is apparent that in the subperiods in which the volatility of PCP is high, so is the volatility of the RER. We find this issue particularly interesting, since the substantial increase in the volatility of real exchange rates after the breakdown of the Bretton Woods system of fixed exchange rates is accompanied by an equally substantial increase in the volatility of commodity prices. The conventional interpretation has been that the increase

[12]We performed unit root tests for the data in (log) levels and in three-, four-, and five-year differences. There is evidence of unit roots for the raw data that vanishes for the data in four-year differences. Still, the real exchange rates and commodity prices remain very persistent even for the data in four-year differences. See Online Appendix C.

[13]When specifying the subperiods, we opted for isolating 1960–1972, the period during which the Bretton Woods system was active. Then we chose the next subperiods so that they would have similar lengths.

[14]This is also the case for small open economies, where the ratio of the volatility of the relevant PCP is between 2.5 and 3.5 times the volatility of the RER. These values are similar to the ones that can be obtained from Table 1.

in volatility after 1972 was the result of the regime change from a fixed to a flexible exchange rate regime (Mussa, 1986). An alternative interpretation is that the fundamentals that make real exchange rates and commodity prices comove were more volatile after 1973 than before.

Table 1: Volatilities of real exchange rates and primary commodity prices

	1960–2014	1960–1972	1973–1985	1986–1998	1999–2014
Real exchange rates					
US-UK	0.18	0.11	0.26	0.17	0.14
US-DEU	0.22	0.07	0.31	0.16	0.21
US-JPN	0.23	0.07	0.27	0.26	0.17
Average across commodities					
Simple	0.38	0.17	0.44	0.30	0.36
Trade weighted	0.46	0.17	0.55	0.35	0.36

Notes: Variables are in logs, and commodity prices are normalized by US CPI. Weights are based on the share of total trade in 1990. The set of 10 primary commodities is oil, fish, meat, aluminum, copper, gold, wheat, maize, timber, and cotton.

As a complementary piece of evidence, Figure 1 shows rolling volatilities computed using windows of 10 years of data for the real exchange rates and for the average of the 10 primary commodity prices. The positive association between the volatilities of the real exchange rates and commodity prices reinforces, in our view, the interest in associating RER with PCP.[15]

5 Empirical results

We start our analysis by reporting the R^2 of the OLS regression of equation (2) using the primary commodity prices listed in Table A.1. We run the regression for the whole period and also for four subperiods. The results are reported in Table 2.

Panel (a) of Table 2 shows regression results for the data in four-year differences using 10 commodities as regressors. The R^2 are 0.48, 0.63, and 0.57 for the United Kingdom, Germany, and Japan, respectively. The R^2s are larger when we consider the subperiods, although as we argue below, they are largely the effect of smaller samples.

As we show in Table C.4, the prices of the commodities that we are using are highly correlated. One could then guess that it is possible to account for a large fraction of the real exchange rate volatility even if we considerably reduce the number of PCP. To explore this possibility, we pick the 4 commodities (out of the 10) with the highest t-statistics and rerun

[15]The correlations are 0.40, 0.54, and 0.39 for the United Kingdom, Germany, and Japan, respectively.

Figure 1: Rolling volatilities of real exchange rates and commodity prices (10-year windows)

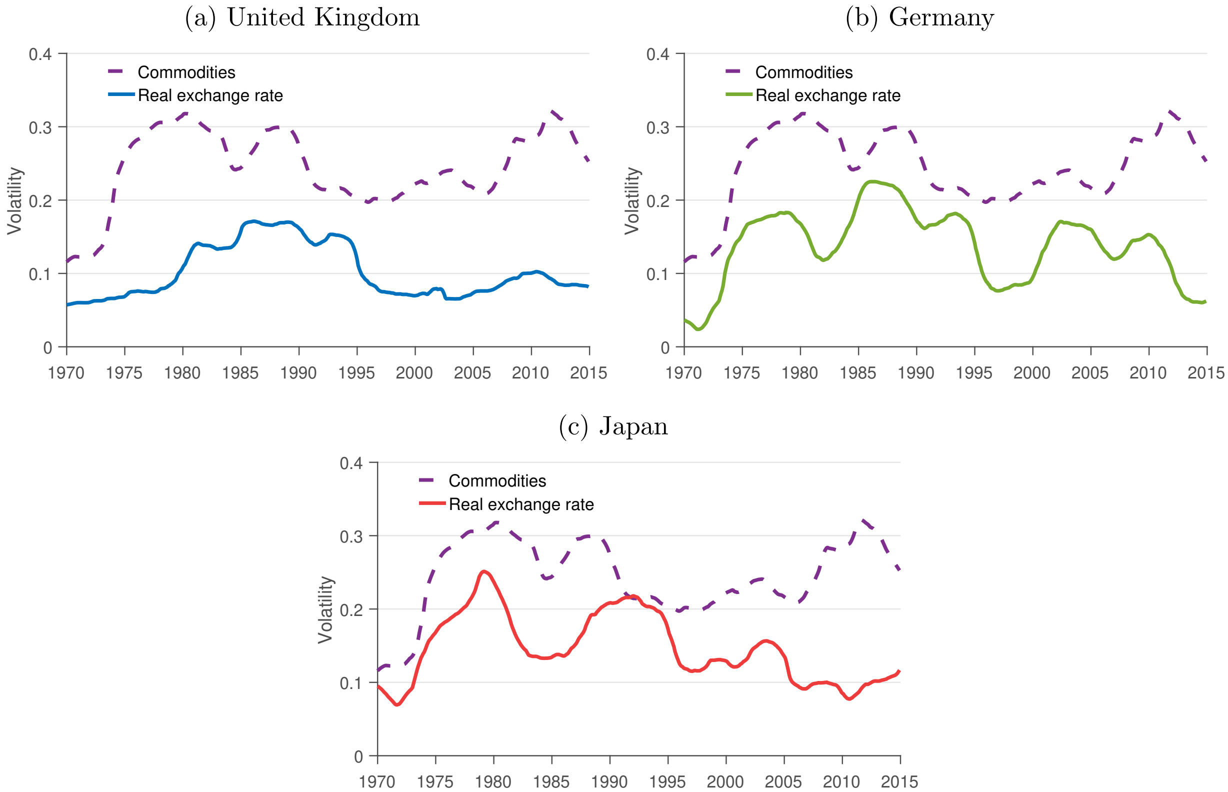

Table 2: Coefficients of determination R^2

	1960–2014	1960–1972	1973–1985	1986–1998	1999–2014
(a) 10 commodities, 4-year differences					
United Kingdom	0.48	0.90	0.90	0.81	0.60
Germany	0.63	0.95	0.87	0.83	0.75
Japan	0.57	0.92	0.84	0.92	0.82
(b) 4 commodities (best fit), 4-year differences					
United Kingdom	0.33	0.72	0.82	0.63	0.58
Germany	0.56	0.84	0.87	0.81	0.74
Japan	0.48	0.88	0.76	0.86	0.80

the regressions.[16] The results are reported in panel (b) of Table 2.[17]

[16]Throughout the paper, we compute t-statistics using the Newey-West heteroskedasticity-and-autocorrelation-consistent standard errors.

[17]Tables F.1–F.3 in Online Appendix F report the coefficients of the regressions in levels and in 4-year differences. We also show the results for the case in which we choose only three commodities.

By selecting only four commodity price series, the R^2s are between 33% and 56% for the regression using four-year differences. It is important to emphasize that they are much larger in all subperiods we consider, and, in particular, there are no systematic differences in the relationship between PCP and real exchange rates before and after the Bretton Woods system. This is in line with the alternative hypothesis about the increase in real exchange rate volatility following 1972: that it coincided with an increase in the volatility of fundamentals.

Figure 2 plots the data versus the respective fitted values for the regressions in four-year differences for the cases of both 10 and 4 PCP and also reports the respective correlation between the data and fitted values (equivalent to the square root of the R^2).[18] As can be seen, the match is very good in all cases.

The Online Appendix shows that these results are robust to several variations in the specification of the regressions: using data in log-levels considering the 10 commodities and selecting the best 4 commodities as above; regressions without subtracting the log of the US CPI from either side of equation (2); regressions using non-US pairs of bilateral real exchange rates, such as the bilateral real exchange rate between the United Kingdom and Germany; and regressions selecting commodities based on US trade data rather than by a statistical criterion. This last procedure can be rationalized by the theory outlined in Section 2. As shown in equation (1), the theory predicts that real exchange rates and commodity prices are related as long as the commodities are non-trivial inputs in the production structure of one of two economies in the country pair.

5.1 Out-of-sample fit

In the previous regressions, we chose the four primary commodities that obtain a good fit with the real exchange rate, so the regressors have been chosen precisely to match the data. To check the robustness of our results to the in-sample selection, we adopt the following procedure. We start by running a regression using data in four-year differences over the period 1960-1972. We drop the six commodities with the lowest t-statistics and rerun the regression. Based on the four commodities selected by this procedure and their estimated coefficients, we use observed commodity prices over the following h periods to fit the real exchange rate and store the h fitted values. We next add one observation to the sample and repeat previous regressions to fit the real exchange rates over the following h periods. Repeating this procedure until the end of the sample, we construct time series of out-of-sample fitted real exchange rates over the following $h = 1, 2, ..., 60$ periods.

The logic behind this exercise is related to the discussion in Section 3. We interpret

[18] For the case of the data in levels, see Figure D.1 in Online Appendix D. The results are very similar, except for Japan, which is the country for which the unit root evidence is very high.

Figure 2: Real exchange rates and fitted values, four-year differences.

(a) United Kingdom

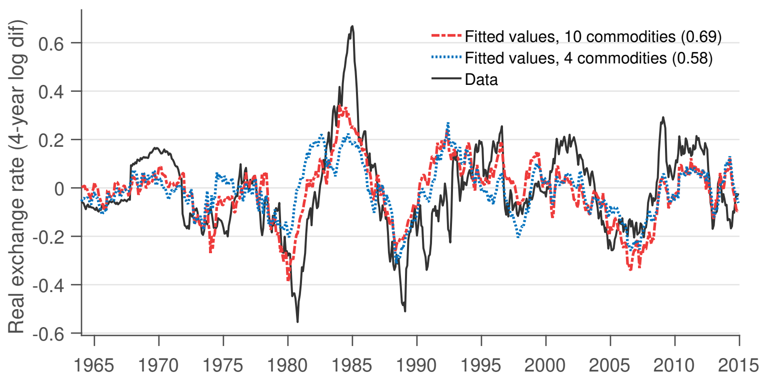

(b) Germany

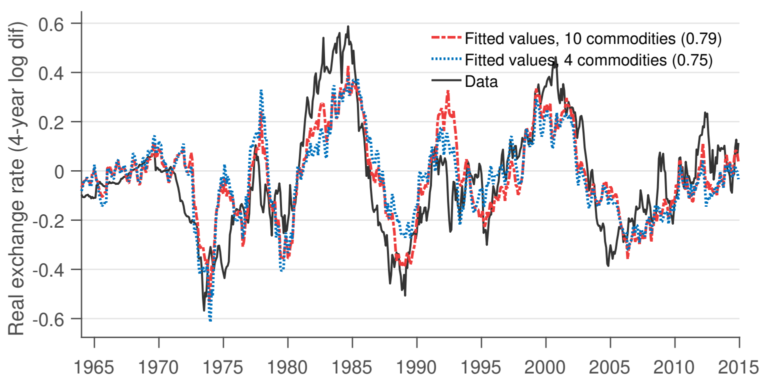

(c) Japan

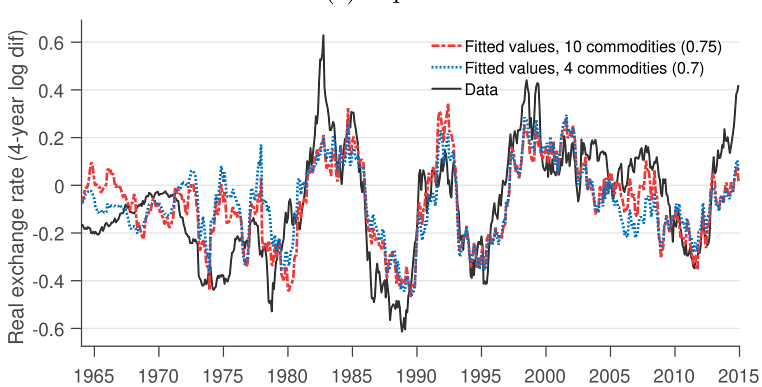

the linear regression as a linear approximation of the solution of a model in which the RER and the PCP are jointly determined, as described in equation (3). The parameters on those equations are evaluated at the equilibrium point at which the linearization is made. The maintained assumption in this exercise is that those values will not change much in a relatively short period of time, so that the reduced-form estimates could work reasonably well for an interval of time that is not too long, particularly if no major changes occurred.

Figure 3: Out-of-sample fit six months ahead with four commodities (best fit)

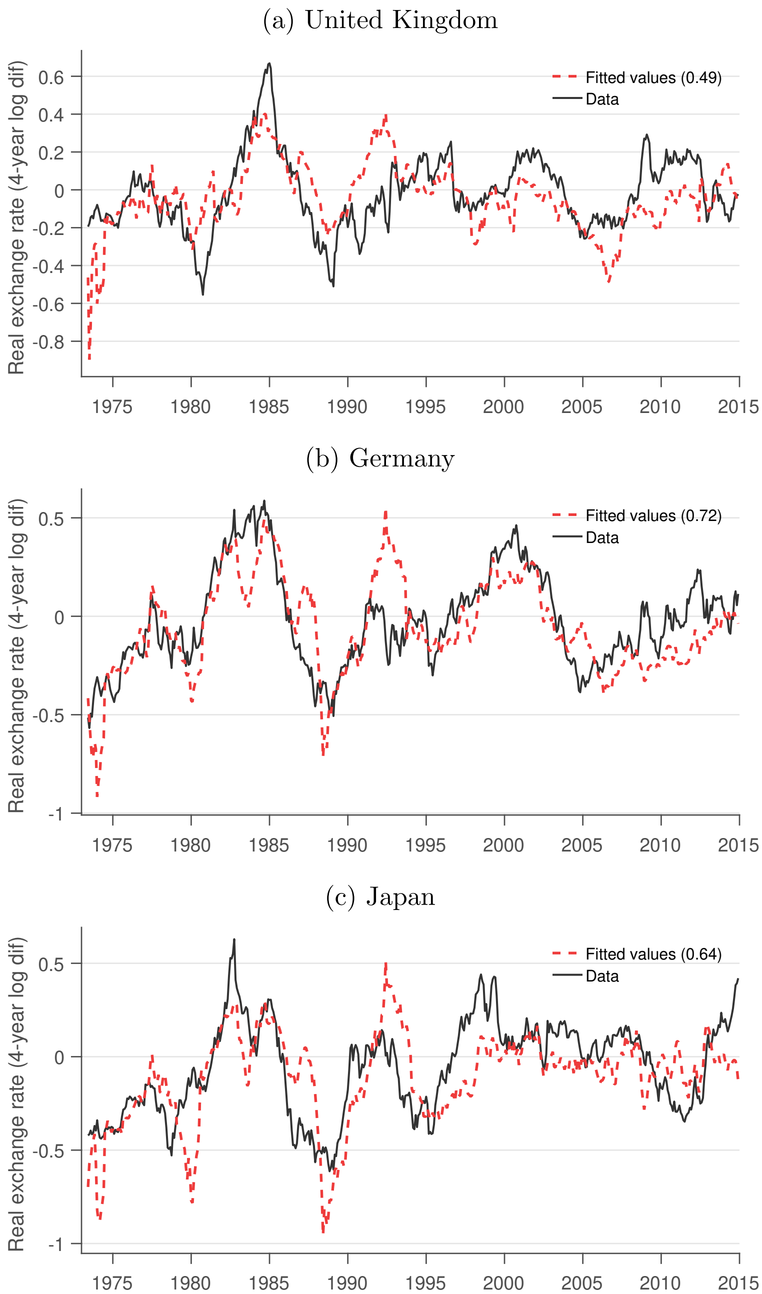

Figure 3 shows the actual and fitted real exchange rates for the case $h = 6$ months ahead. The out-of-sample fit is remarkable, with a correlation between the fitted and actual values of 0.45 for the United Kingdom, 0.73 for Germany, and 0.64 for Japan.

We summarize the results in Figure 4, in which we show the correlation between fitted and actual real exchange rates as we vary the forward window from $h = 1$ to $h = 60$ months ahead. Although the correlations decrease as the fitting horizon increases, they decrease slowly. There is a good out-of-sample fit even using data that are several years old to select the commodities and coefficients to fit real exchange rates today.

Figure 4: Out-of-sample fit, four commodities, correlations as a function of r (months ahead)

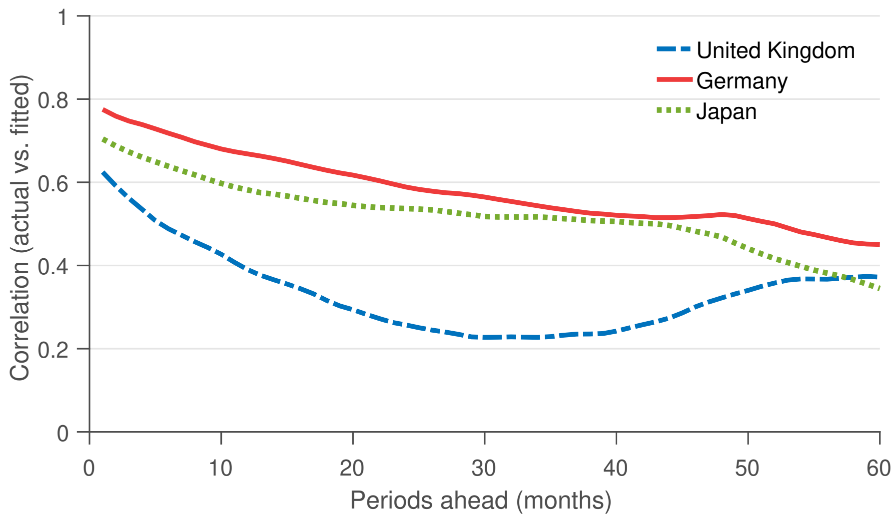

5.2 Are the results spurious?

A concern with the previous regressions is to what extent the results could be due to a problem of small sample size. It is well known that, even with stationary series, regressing two orthogonal but highly persistent series could lead to a spurious correlation for moderate sample sizes. To explore this issue, we perform small sample inference by using a parametric bootstrap procedure that generates artificial data under the null hypothesis that commodity prices are orthogonal to real exchange rates. By construction, commodity prices and real exchange rates are orthogonal, which implies that the R^2 converges to zero as the artificial sample size grows toward infinity. But for finite samples, the R^2 is positive.

Take, for example, the Germany-US real exchange rate regression with an R^2 of 0.56 in panel (b) of Table 2. The bootstrap procedure is as follows. We first estimate an autoregressive process for the Germany-US real exchange rate and an independent vector autoregression with the four commodity prices used in the regression (we use the Schwarz information criterion to select the lag lengths). To compute the small sample distribution of the R^2, we draw 10,000 samples of length 660 (the number of months between January 1960 and December 2014) by resampling from the residuals of the estimated processes and compute artificial real exchange rate and commodity price data. Next, for each artificial sample, we run a regression of the real exchange rate on the commodity prices and store the associated R^2. Finally, we compare the estimated R^2 using actual data with its small sample distribution to assess how common it is to observe an R^2 of 0.56 under the null hypothesis of orthogonality.[19]

Figure 5: Small sample distribution of the R^2 over the period 1960–2014

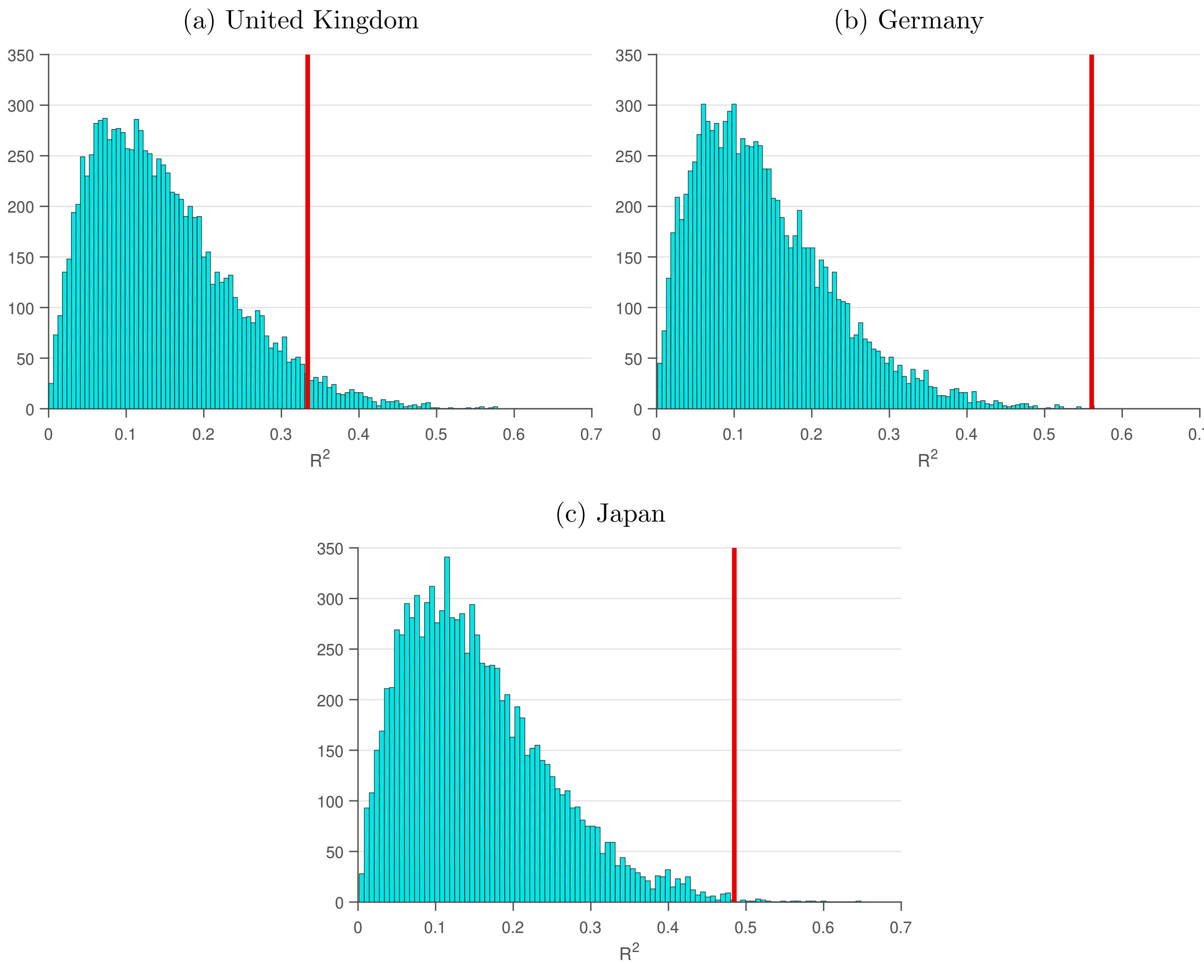

Figure 5 shows the small sample distributions of the R^2 over the entire sample period. The vertical lines are the estimated R^2 using the actual data. In all cases, the probability of obtaining an R^2 as large as that estimated in panel (b) of Table 2 is smaller than 5% and

[19]We also performed a Monte Carlo simulation in which the only difference with the bootstrap procedure is that shocks are drawn from zero-mean normal distributions with a variance and covariance matrix equal to those estimated with the real exchange rate and commodity prices data. The results obtained using the Monte Carlo simulation are virtually identical to those using the bootstrap procedure.

as low as 0% for the case of Germany. The three distributions under the null hypothesis are positively skewed with a mode of about 0.1, which is much smaller than the estimated R^2s in the table.

Table 3: Bootstrapped distributions of R^2 under the null hypothesis of orthogonality.

		Percentiles distribution of R^2				
	$\hat{R}^2$	Median	75	90	95	$\Pr(R^2 \geq \hat{R}^2)$
United Kingdom						
1960-2014	0.33	0.13	0.20	0.27	0.31	0.037
1960-1972	0.72	0.52	0.66	0.75	0.80	0.143
1973-1985	0.82	0.37	0.52	0.64	0.70	0.004
1986-1998	0.63	0.37	0.50	0.61	0.67	0.077
1999-2014	0.58	0.29	0.41	0.53	0.59	0.059
Germany						
1960-2014	0.56	0.13	0.19	0.26	0.31	0.000
1960-1972	0.84	0.56	0.69	0.79	0.83	0.032
1973-1985	0.87	0.49	0.63	0.73	0.78	0.005
1986-1998	0.81	0.40	0.54	0.65	0.71	0.007
1999-2014	0.74	0.30	0.43	0.55	0.61	0.007
Japan						
1960-2014	0.48	0.14	0.21	0.29	0.34	0.003
1960-1972	0.88	0.59	0.72	0.81	0.85	0.022
1973-1985	0.76	0.46	0.60	0.70	0.75	0.045
1986-1998	0.86	0.41	0.55	0.66	0.71	0.001
1999-2014	0.80	0.33	0.46	0.57	0.63	0.002

Table 3 shows statistics of the small sample distributions under the null of orthogonality for the three bilateral real exchange rates and for the five subperiods, together with the probability of observing an R^2 as large as that estimated with the actual data under the null of orthogonality. For comparison, the table also includes the estimated R^2s from panel (b) of Table 2. Overall, these results suggest that the estimated correlations are robust for every subperiod and bilateral real exchange rate. Of course, for some subperiods and countries, the small sample distributions are more dispersed, and it is not uncommon to observe a relatively large R^2 under the null of orthogonality, especially for smaller sample sizes. For example, although the estimated R^2 for Germany over the period 1964-1972 is 0.84, the median R^2 under the null of orthogonality is 0.56.

We also computed the small sample distributions of the out-of-sample-fit exercise of Subsection 5.1. For each country, we created 2,000 artificial correlations as a function of $h = 1, 2, ..., 60$ replicating the procedure in Figure 4 but imposing that real exchange rates are orthogonal to commodity prices, as we did before. Figure 6 displays the median correlation for each country under the null hypothesis (solid line), and the shaded areas represent the 5th and 95th percentiles of the small sample distribution of the correlation as a function of the horizon $h = 1, 2, ..., 60$. The dashed lines represent the estimated correlations from Figure 4. In most cases, we reject the null hypothesis of orthogonality.

Figure 6: Fitted correlations and bootstrap bands under the null hypothesis of orthogonality (with four commodities, best fit)

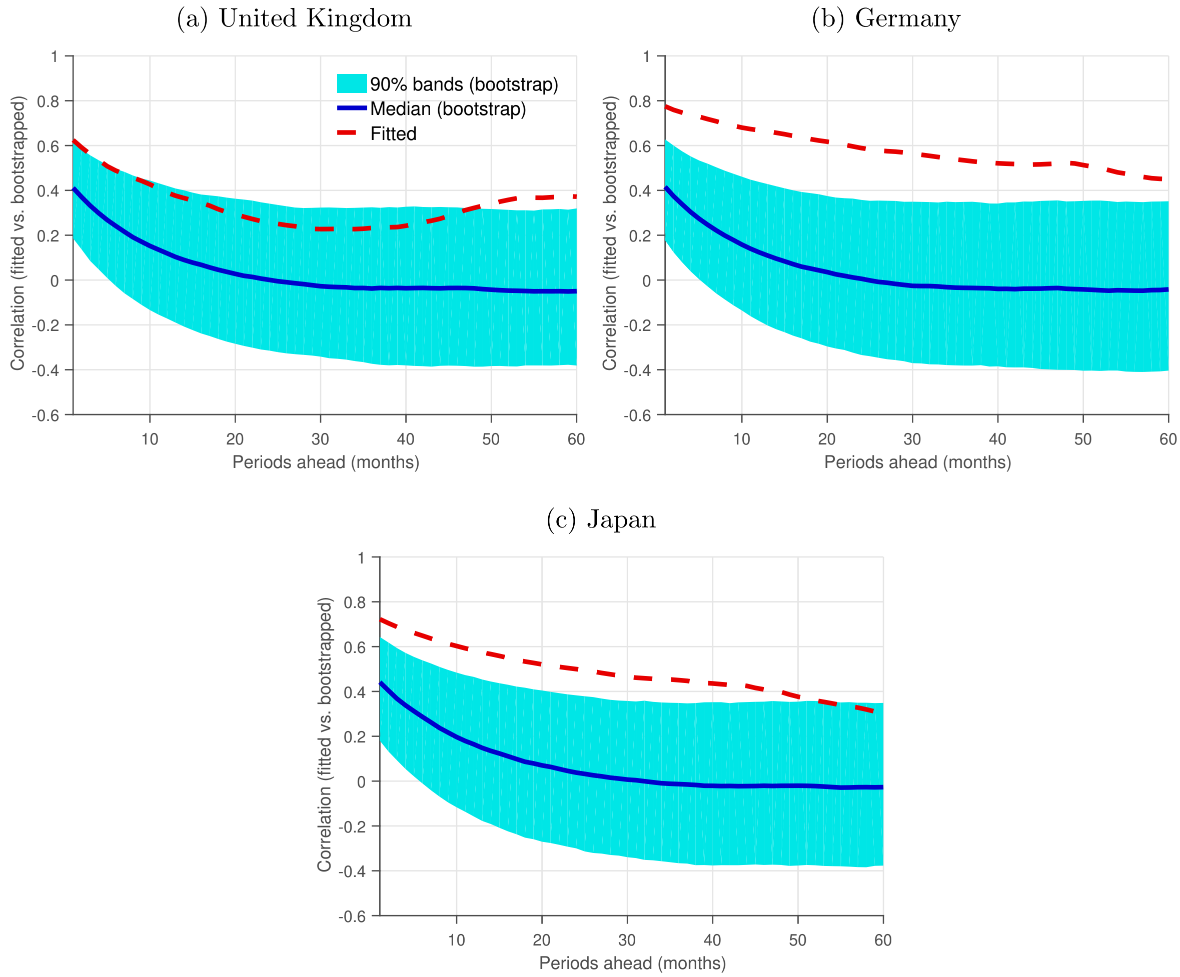

6 Calibration and simulations

In this section, we calibrate the model presented in Section 2 for the United States and Japan in order to run some numerical experiments. First, we calibrate the stochastic processes of the exogenous variables—the variance-covariance matrix and the autocorrelations of the shocks to the endowments in the rest of the world and the evolution of the total factor productivity shocks—so as to match the behavior of output in the two countries and the behavior of the primary commodity prices. We then show that, in spite of its extreme simplicity, the model is able to generate a very volatile and persistent RER. Second, we use the model to quantify the size of the bias in the ordinary least squares estimates of the coefficients of equation (2). In Section 3, we explain theoretically why the estimates of those coefficients lack any interpretation and why they could be unrelated to the structural parameters of the model. The advantage of running the same experiment with model-simulated data is that we can compare the estimates with the structural parameters in equation (1). As we show, the bias can be very large, and the values of the estimates can even change sign when one omits a regressor. These exercises reinforce our decision to completely disregard the discussion of the parameter estimates in the regressions. On the other hand, we show that the R^2 is a relevant measure in the simulated data.

To calibrate our model economy, we proceed in three steps. The first step consists of calibrating the parameters of the Cobb-Douglas production functions, which correspond to the factor shares in equilibrium. We use the 2005 Japan-US Input-Output Table published by the Ministry of Economy, Trade, and Industry (METI) of Japan. We map each of the 174 sectors in the input-output table into the three sectors considered in our model: final goods, intermediate goods, and primary commodities. The exact mapping is presented in Appendix A.2. The group of all final goods in the United States is assumed to be Y^1, and the group of all its intermediate goods is assumed to be Q^1. We do the same for Y^2 and Q^2 in the case of Japan. Regarding the primary commodities, we assume that X_1 corresponds to the sector "petroleum and natural gas" and that X_2 corresponds to the sectors "fishing" and "seafood" together. The rest of the primary commodities are grouped into X_3. These are, for each of the two economies, the commodities that have large shares in commodity production, where exports are a large share and where the two countries differ the most. Recall that the model suggests that asymmetries are important in generating fluctuations in the real exchange rate, and the main source of heterogeneity between the two countries will be in these two sectors, since we will group all other commodities in the two countries in X_3.

The input-output table contains data on the payments to each of the factors of production such as intermediate inputs, compensation of employees, and operating surplus. We compute the shares of each factor of production considered in the model to pin down the parameters of the Cobb-Douglas production functions described in Section 2. For the intermediate and final good sectors, we assume that the payments to labor input are equal to the value added in the data. In the primary commodity sector, on the other hand, the labor share is computed as the share of compensation of employees in value added.

We do not have a similar dataset for the input-output table of country 3, the rest of the world (ROW). We use data from the 10-Sector Database available from the Groningen Growth and Development Center to compute the share of the commodity sector in the ROW GDP. The commodity sector is assumed to comprise both the agriculture and mining sectors. Then we use data from Comtrade to compute the relative factor shares of each primary commodity based on their respective shares in total world trade in primary commodities in 2000. The resulting factor shares are presented in Table 4.

Table 4: Calibration: factor shares (%)

	Country 1 (USA)	**Country 2 (JPN)**
Final good		
intermediate good Q_1	$\alpha_1^1 = 20.2$	$\alpha_1^2 = 0.3$
intermediate good Q_2	$\alpha_2^1 = 0.1$	$\alpha_2^2 = 23.5$
labor n^q	$\alpha_3^1 = 79.7$	$\alpha_3^2 = 76.2$
Intermediate good		
primary commodity X_1	$\beta_1^1 = 6.5$	$\beta_1^2 = 5.9$
primary commodity X_2	$\beta_2^1 = 0.0$	$\beta_2^2 = 0.1$
primary commodity X_3	$\beta_3^1 = 4.7$	$\beta_3^2 = 11.5$
labor n^q	$\beta_4^1 = 88.8$	$\beta_4^2 = 82.5$
Primary commodity X_i		
labor $n_i^{x_i}$	$\phi_1^1 = 27.8$	$\phi_2^2 = 33.0$
natural resource e_i^i	$1 - \phi_1^1 = 72.2$	$1 - \phi_2^2 = 67.0$
Primary commodity X_3		
labor $n_i^{x_3}$	$\phi_3^1 = 50.0$	$\phi_3^2 = 28.5$
natural resource e_3^i	$1 - \phi_3^1 = 50.0$	$1 - \phi_3^2 = 71.5$
	Country 3 (ROW)	
Final good		
primary commodity X_1	$\pi_1 = 5.8$	
primary commodity X_2	$\pi_2 = 0.5$	
primary commodity X_3	$\pi_3 = 3.5$	
labor n_3	$\pi_4 = 90.2$	

The second step consists of calibrating the relative size of each economy in steady state. We normalize the TFP level in countries 1 and 2 to be equal to one, so the relative sizes of the economies are determined by their relative endowments of labor, natural resources, and primary commodities. We normalize country 1, the United States, to have size equal to one and calibrate the other parameters in order to exactly match the average relative size of each economy in terms of nominal GDP between 1960 and 2014. We use data on nominal

GDP in US dollars from the World Development Indicators (WDI) from the World Bank to decompose world GDP between the United States, Japan, and ROW in the 1960–2014 period. The calibration is presented in Table 5. The relative size of the endowments of natural resources with respect to the endowment of labor within each country is set to be one. The relative endowments of primary commodities in the ROW are chosen such that the model matches the shares of primary commodities X_1 and X_2 in the commodity sector GDP of countries 1 and 2, respectively. The labor endowment in the ROW is normalized to one. All the results presented here are robust to changes in the relative size of endowments within countries while keeping the relative sizes of their economies constant.

Table 5: Calibration: relative sizes in steady state

		Share of world GDP (%)	
	Parameters	**Data**	**Model**
Country 1 (USA)	$n^1 = e_1^1 = e_3^1 = 1$	30	30
Country 2 (Japan)	$n^2 = e_2^2 = e_3^2 = 0.33$	10	10
Country 3 (ROW)	$n^3 = 1$, $X_1^3 = 2.32$, $X_2^3 = 1.04$, $X_3^3 = 0.32$	60	60

The third step consists of calibrating the stochastic processes for the five exogenous shocks to the economy: the TFP shocks in the United States and Japan, and the endowment of the three commodities in the rest of the world. In the model, we conceptually isolate the effect of productivity shocks from the effect of shocks in the endowments by assuming that they are orthogonal. Thus, five parameters will govern the variance-covariance matrix, plus the autocorrelations of the two productivity shocks. In addition, nine parameters will govern the variance-covariance matrix, plus the autocorrelations of the three endowment shocks. These 14 parameters will govern the stochastic processes of output in the two countries and the three primary commodity prices. Therefore, these parameters will determine the invariant distribution of the variables in the model economy, and we will use 14 of its moments to calibrate the parameter values based on the same moments observed in the data.

Although 14 moments in the data are used to calibrate the 14 parameters in the model, an exact solution may not exist because of the nonlinearities of the model. Thus, to calibrate the 14 parameters, we minimize the (Euclidean) distance between 14 moments in the model and the data with respect to fluctuations in real GDP and primary commodity prices. The time period of our model is chosen to be one year. The parameters, moments, and their respective values are described in Table 6.

We use the cyclical component of the annual real GDP per capita series for the United

States and Japan using the HP filter with parameter 6.25. For the primary commodity prices, we use the price of oil as p^{x_1}, the price of fish as p^{x_2}, and the price of aluminum as a proxy for the price of the rest of the commodities p^{x_3}, all described in Section 4. We choose the price of aluminum as p^{x_3} because the group of metals is the largest group of primary commodities in terms of trade volume excluding oil. Table 6 shows that the moments implied by our model closely match their counterparts in the data. The results are based on a simulation of $6,000$ periods, in which we drop the first $1,000$ periods.

6.1 Results

With the calibrated model, we can assess how well the model performs with respect to non-targeted moments. The results are presented in Table 7. The first column shows the data and the second column (labeled "Model") shows the results under the calibration discussed above.

To highlight the quantitative role of commodity production and consumption, we also simulate the model using the same parameterization except that we reduce the Cobb-Douglas coefficients of primary commodities in the production of intermediate goods β_1^i, β_2^i, and β_3^i for $i = 1, 2$.[20] Specifically, we divide the coefficients on the commodities in Table 4 by 10,000. We report the results of these exercises in the third column (labeled "Model II") of Table 7. As we show below, this change implies that the equilibrium value added of commodities over total GDP is negligible.

We start by analyzing the sectoral composition of GDP in each country. The main difference between the model and the data is with respect to the shares of the final and intermediate goods sectors. For example, in the case of the United States, the share of the intermediate goods sector in US GDP is 47.9% in the data, whereas it is just 18% in the model. This result is not surprising, since we have a simplified model of each economy. What is most important in the sectoral composition in Table 7 is that we are not overstating the relative size of the primary commodity sector in these countries, which is the main feature of our analysis. In the case of the United States, for example, the primary commodity sector accounts for 3.6% of GDP in the data and 2.2% in the model.

Next, we analyze the standard deviation and autocorrelation of the RER that are implied by the model. The result is striking: the model is able to deliver volatility and autocorrelation measures of the RER that are close to the data. The standard deviation of the RER is 28% in the model and 37% in the data, whereas its autocorrelation is 0.99 in the model and 0.96 in the data.

[20] We increase the coefficient on labor β_4^i accordingly.

Table 6: Calibration: stochastic processes

Moments	**Data**	**Model**
Standard deviation of US real GDP per capita (%)	1.3	1.3
Standard deviation of Japan real GDP per capita (%)	1.6	1.6
Autocorrelation of US real GDP	0.31	0.31
Autocorrelation of Japan real GDP	0.18	0.18
Correlation between US and Japan real GDP	0.41	0.41
Standard deviation of the price of oil (%)	66.7	86.7
Standard deviation of the price of fish (%)	35.5	29.4
Standard deviation of the price of aluminum (%)	31.5	31.7
Autocorrelation of the price of oil	0.92	0.99
Autocorrelation of the price of fish	0.76	0.79
Autocorrelation of the price of aluminum	0.84	0.99
Correlation between the prices of oil and fish	0.29	0.22
Correlation between the prices of oil and aluminum	-0.22	-0.72
Correlation between the prices of fish and aluminum	0.37	0.27

Parameters	**Values**
$100 \times$ standard deviation of $\varepsilon_t^{z_1}$	1.0
$100 \times$ standard deviation of $\varepsilon_t^{z_2}$	1.2
ρ^{z_1}: autocorrelation of $z_{1,t}$	0.30
ρ^{z_2}: autocorrelation of $z_{2,t}$	0.18
correlation between $\varepsilon_t^{z_1}$ and $\varepsilon_t^{z_2}$	0.41
$100 \times$ standard deviation of $\varepsilon_t^{x_1}$	6.6
$100 \times$ standard deviation of $\varepsilon_t^{x_2}$	13.8
$100 \times$ standard deviation of $\varepsilon_t^{x_3}$	4.0
ρ^{x_1}: autocorrelation of $x_{1,t}^3$	0.99
ρ^{x_2}: autocorrelation of $x_{2,t}^3$	0.00
ρ^{x_3}: autocorrelation of $x_{3,t}^3$	0.99
correlation between $\varepsilon_t^{x_1}$ and $\varepsilon_t^{x_2}$	-0.00
correlation between $\varepsilon_t^{x_1}$ and $\varepsilon_t^{x_3}$	-0.77
correlation between $\varepsilon_t^{x_2}$ and $\varepsilon_t^{x_3}$	0.08

Notes: Real GDP per capita corresponds to the cyclical component of the HP-filtered data with smoothing parameter 6.25. Primary commodity prices are normalized by US CPI in both the model and the data. We simulate the model at an annual frequency, so the statistics on primary commodity prices are based on an annual series constructed with the average values within each year. The correlations between TFP shocks and shocks to the endowments of primary commodities are set to zero.

The importance of primary commodities can be assessed from the results in Model II. When we reduce the shares of primary commodities in the production of intermediate goods, the standard deviation drops substantially. Note, also, that the share of value added accounted for by commodities is much smaller for both the United States and Japan.

The close relationship between the RER and primary commodity prices is also present

Table 7: Non-targeted moments

	Data	Model	Model II
(1) Share of country GDP (%)			
United States			
Final good Y_1	48.50	79.65	79.65
Intermediate good Q_1	47.88	18.03	20.35
Primary commodity X_1	1.18	0.77	0.00
Primary commodity X_3	2.40	1.55	0.00
Japan			
Final good Y_2	46.34	79.24	76.21
Intermediate good Q_2	49.85	20.51	23.79
Primary commodity X_2	0.38	0.03	0.00
Primary commodity X_3	3.42	0.21	0.00
(2) Standard deviation of RER (%)	37.0	27.83	1.52
(3) Autocorrelation of RER	0.96	0.99	0.22
(4) $\mathbf{R^2}$ of OLS regression	0.80	0.98	0.07

Notes: We report the standard deviation of the log of the RER. OLS regressions are based on the four-year differences of the log of RER and primary commodity prices. The latter are normalized by US CPI. The R^2 of the OLS regression in the data corresponds to the R^2 in the 1999–2014 sub-period in Table 2. In Model II we depart from the benchmark calibration by reducing the importance of primary commodities in the production of intermediate goods. That is achieved by dividing β_1^1, β_2^1, β_3^1, β_1^2, β_2^2, and β_3^2 in Table 4 by 10,000.

when we run the OLS regressions of the RER on primary commodity prices in the simulated data. In both cases, we use the regression in four-year differences. The R^2 of the OLS regression is 0.98. This result is less surprising: movements in the real exchange rate in the model are due to either productivity shocks or shocks to the endowment of commodities. It is well known that productivity shocks do not move real exchange rates very much, so most of the volatility of the real exchange rate is driven by shocks that, in the model, affect the supply of commodities. The comparable value for the R^2 in the data is the one for Japan for the last subperiod, which runs from 1999 to 2014. The value is comparable because the input-output table used in the calibration is from this period, and the values change substantially from decade to decade. The R^2 for the case of four commodities is 0.80, below the one delivered by the model. Part of the reason for this value is that the model has no other shocks besides productivity shocks; moreover, in the model, we know exactly what the relevant prices of primary commodities are. The data, however, include many primary commodities, and we are using only a few in the regression.

The importance of primary commodities in generating such results can also be seen from the comparison of the R^2s of the OLS regressions under the benchmark calibration and the calibration with low shares of primary commodities in the production of intermediate goods, Model II. In the latter case, the R^2 drops to 0.07.

Table 8: Variance decomposition

	Benchmark	**No TFP shock**	**No endowment shock**
100 x standard deviation of			
Country 1 real GDP	1.3	0.0	1.3
Country 2 real GDP	1.6	0.0	1.6
Real exchange eate (RER)	27.8	45.8	0.4
Price of X_1	86.7	89.8	0.4
Price of X_2	29.4	42.1	0.3
Price of X_3	31.7	52.1	0.2

Table 8 shows a variance decomposition exercise. The first column reproduces the volatility of output for the two countries, the RER, and the three commodity prices for the benchmark calibration. The second column shows the volatility for the same variables when we shut down the productivity shocks in the two countries. Naturally in this case, the volatility of output vanishes in the two countries, since the correct measure of output is independent of changes in the terms of trade (Kehoe and Ruhl, 2008). Interestingly, the volatility of the real exchange rate is higher than in the benchmark calibration. This finding should not necessarily be surprising given that the model is highly nonlinear and thus volatilities are not additive. The last column shows the results when we shut down the shocks to the supply of commodities in ROW. In this case, the model reproduces the volatility of output in both countries, but the volatility of the RER and the PCP collapse to very small values. One could conjecture that by assuming a common TFP shock across sectors, we undermined the ability of the model to generate a volatile RER with TFP shocks alone. However, this is not the case. A calibration in which the volatility of the TFP shock to commodity production is 10 times higher than for the final good, and which matches the volatility of output, increases the volatility of the RER to only 28.8%, a value is still very close to the benchmark of 27.8%.

Finally, we can use the OLS regressions in the simulated data to quantify the bias in estimating the coefficients on the primary commodity prices. The coefficients are reported in Table 9. The table shows that we cannot extract any relevant information from the estimated coefficients. The first reason is that we do not know what makes up the relevant set of primary commodity prices in the data. As is evident in Table 9, the coefficients change significantly when we estimate the regression without including the price of X_3 as a regressor, for example. The coefficient on p^{x_1} changes its sign from -0.007 to 0.156. The second reason

Table 9: Coefficients of the OLS regressions

	p^{x_1}	p^{x_2}	p^{x_3}
Baseline regression	-0.007	-0.008	-0.886
Regression without $p_t^{x_3}$	0.156	-0.040	
Coefficients implied by equation (1)	2.983	-2.612	-0.017

is that even when we include the correct set of primary commodity prices, which is the case in our baseline regression, the estimated coefficients are very different from the ones implied by equation (1). Table 9 shows that, based on our benchmark calibration, equation (1) would imply a coefficient equal to 2.983 on the price of primary commodity X_1, whereas the regression in the simulated data delivers a coefficient equal to -0.007.

7 Conclusions

In this paper, we provide empirical evidence that points toward a common factor that moves a handful of primary commodity prices on the one hand and real exchange rates between the United States and the United Kingdom, Germany, and Japan on the other. More specifically, we show that shocks that move just four primary commodity prices can account for between one-third to one-half of the volatility of the real exchange rates for a period that lasts more than half a century. For periods that are just one decade and a half, that fraction can go all the way up to 90%.

We also numerically solve a very simple model in which the exogenous stochastic driving processes are shocks to productivity and shocks to the supply of primary commodities. We calibrate the model so as to reproduce the volatility and persistence of primary commodity prices. We show that the model is able to replicate the high volatility and persistency of real exchange rates in the data.

Ever since the pathbreaking work of Obstfeld and Rogoff (1995), theoretical models of real exchange rates for large developed economies have almost exclusively focused on trade in final differentiated goods. A challenge for the literature has been to deliver large and persistent fluctuations in the ratio of final consumption prices measured in the same units, without large movements in quantities of final goods consumed, since quantities do not move nearly as much in the data. This first exploration in which we explicitly added the production of commodities to a model with multiple large economies went a sizable way toward solving that problem and placed a long research path ahead of us. The empirical results of this paper suggest that the path is worth pursuing.

References

Ayres, Joao Luiz, Constantino Hevia, and Juan Pablo Nicolini. 2017. "Real Exchange Rates and Primary Commodity Prices." Federal Reserve Bank of Minneapolis Working Paper 743.

Backus, David K., and Mario J. Crucini. 2000. "Oil prices and the terms of trade." *Journal of International Economics*, 50(1): 185–213.

Betts, Caroline, and Timothy Kehoe. 2004. "U.S. Real Exchange Rate Fluctuations and Relative Price Fluctuations." Federal Reserve Bank of Minneapolis Staff Report 334.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan. 2002. "Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?" *Review of Economic Studies*, 69(3): 533–563.

Chen, Yu-Chin, and Kenneth Rogoff. 2003. "Commodity Currencies." *Journal of International Economics*, 60(1): 133–160.

Chen, Yu-Chin, Kenneth S. Rogoff, and Barbara Rossi. 2010. "Can Exchange Rates Forecast Commodity Prices?" *The Quarterly Journal of Economics*, 125(3): 1145–1194.

Côté, Agathe. 1987. "The Link between the U.S. Dollar Real Exchange Rate, Real Primary Commodity Prices, and LDCs' Terms of Trade." *The Review of Economics and Statistics*, 69(3): 547–551.

Crucini, Mario J., and Anthony Landry. 2019. "Accounting for real exchange rates using micro-data." *Journal of International Money and Finance*, 91: 86–100.

Dornbusch, Rudiger. 1985*a*. "Inflation, Exchange Rates and Stabilization." National Bureau of Economic Research Working Paper 1739.

Dornbusch, Rudiger. 1985*b*. "Policy and Performance Links between LDC Debtors and Industrial Nations." *Brookings Papers on Economic Activity*, 16(2): 303–368.

Dornbusch, Rudiger. 1987. "Exchange Rates and Prices." *American Economic Review*, 77(1): 93–106.

Eaton, Jonathan, Samuel Kortum, and Brent Neiman. 2016. "Obstfeld and Rogoff's International Macro Puzzles: A Quantitative Assessment." *Journal of Economic Dynamics and Control*, 72: 5–23.

Engel, Charles. 1999. "Accounting for U.S. Real Exchange Rate Changes." *Journal of Political Economy*, 107(3): 507–538.

Galí, Jordi, and Tommaso Monacelli. 2005. "Monetary Policy and Exchange Rate Volatility in a Small Open Economy." *Review of Economic Studies*, 72(3): 707–734.

Hevia, Constantino, and Juan Pablo Nicolini. 2013. "Optimal Devaluations." *IMF Economic Review*, 61(1): 22–51.

Itskhoki, Oleg, and Dimitry Mukhin. 2017. "Exchange Rate Disconnect in General Equilibrium." National Bureau of Economic Research Working Paper 23401.

Johansen, Soren. 1991. "Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models." *Econometrica*, 59(6): 1551–1580.

Kehoe, Timothy J., and Kim J. Ruhl. 2008. "Are shocks to the terms of trade shocks to productivity?" *Review of Economic Dynamics*, 11(4): 804 – 819.

Meese, Richard A., and Kenneth Rogoff. 1983. "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?" *Journal of International Economics*, 14(1-2): 3–24.

Mussa, Michael. 1986. "Nominal exchange rate regimes and the behavior of real exchange rates: Evidence and implications." *Carnegie-Rochester Conference Series on Public Policy*, 25: 117–214.

Obstfeld, Maurice, and Kenneth Rogoff. 1995. "Exchange Rate Dynamics Redux." *Journal of Political Economy*, 103(3): 624–6G0.

Obstfeld, Maurice, and Kenneth Rogoff. 2001. "The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?" In *NBER Macroeconomics Annual 2000, vol. 15, ed. Ben S. Bernanke and Kenneth Roggoff.*

Sachs, Jeffrey D. 1985. "The Dollar and the Policy Mix: 1985." *Brookings Papers on Economic Activity*, 16(1): 117–197.

Stock, James H., and Mark W. Watson. 1993. "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems." *Econometrica*, 61(4): 783–820.

A Data

We used (end-of-period) monthly series for the nominal exchange rates and the consumer price index (CPI) of each country, both obtained from Global Financial Data. The commodity price series are from the World Bank *Commodity Price Data* (Pink Sheet) and the United Nations (UNCTAD*stat*). We excluded natural gas, coal, and iron because of data availability.[21] The data sources for the price series of each commodity are as follows:

(1) Petroleum - Brent crude oil. Source: Global Financial Data, Ticker: BRT_ D.

(2) Fish - price of fish meal. Source: UNCTAD*stat*.

(3) Meat - price of beef. Source: World Bank *Commodity Price Data*.

(4) Aluminum - Source: World Bank *Commodity Price Data*.

(5) Copper: - Source: World Bank *Commodity Price Data*.

(6) Gold - Source: World Bank *Commodity Price Data*.

(7) Wheat - US, $n°1$, hard red winter. Source: World Bank *Commodity Price Data*.

(8) Maize - Source: World Bank *Commodity Price Data*.

(9) Timber - Logs, Malaysia. Source: World Bank *Commodity Price Data*.

(10) Cotton - Cotton Outlook A index. Source: World Bank *Commodity Price Data*.

Table A.1 shows the selected primary commodities and their respective definitions and shares in world trade.[22]

Table A.1: List of Primary Commodities

Commodity	(%) share of world trade in 1990	SITC (rev.3)	Commodity	(%) share of world trade in 1990	SITC (rev.3)
(1) Petroleum	7.22	33	(6) Gold	0.42	971.01
(2) Fish	1.05	03	(7) Wheat	0.35	041
(3) Meat	0.89	011/012	(8) Maize	0.28	044
(4) Aluminum	0.49	285.1/684.1	(9) Timber	0.26	24
(5) Copper	0.45	283.1/682.1	(10) Cotton	0.22	263

Note: SITC (rev3) stands for Standard International Trade Classification (revision 3).
Source: Comtrade.

[21] We also performed all the experiments in the paper using sugar instead of gold (which also serves as a store of value), and the results are virtually the same.

[22] We repeated the analysis using trade data in 2000, and the results remain the same. In this case, maize and cotton are replaced by platinum and coffee.

A.1 Trade data

Trade data were obtained from the United Nations *Comtrade* Database.[23] World trade (exports+imports) for each commodity and its total were computed as the sum of trade over all the countries in the dataset.

Table A.2: Share of imports and exports in each country (% average in 1990–1999)

	United States		**United Kingdom**		**Germany**		**Japan**	
	Imports	Exports	Imports	Exports	Imports	Exports	Imports	Exports
Petroleum	8.5	1.2	3.4	6.0	5.2	0.7	13.1	0.4
Fish	1.0	0.6	0.7	0.5	0.5	0.2	5.0	0.2
Meat	0.3	0.9	0.7	0.6	1.1	0.4	2.4	0.0
Aluminum	0.4	0.2	0.3	0.2	0.5	0.1	1.5	0.0
Copper	0.2	0.2	0.3	0.0	0.4	0.1	1.2	0.1
Gold	0.3	0.9	0.0	0.1	0.3	0.2	0.8	0.1
Wheat	0.0	0.8	0.1	0.3	0.1	0.2	0.4	0.0
Maize	0.0	1.1	0.1	0.0	0.1	0.0	0.8	0.0
Timber	0.9	1.0	0.8	0.0	0.5	0.2	3.5	0.0
Cotton	0.0	0.5	0.0	0.0	0.1	0.0	0.3	0.0
SUM	**11.8**	**7.2**	**6.5**	**7.7**	**8.7**	**2.1**	**29.0**	**0.8**

A.2 Input-output tables

Data for the input-output tables of the United States and Japan come from the 2005 Japan-US Input-Output Table published by the Ministry of Economy, Trade, and Industry (METI) of Japan.[24] We map each of the 174 sectors into three sectors: final goods, intermediate goods, and primary commodities. The mapping with sector codes is the following:

Final goods: 022, 027, 030, 132–137, 147, 152, 154, 167–171.

Primary commodities: 001–020, 024, 025, 029, 031, 032, 039–041, 043, 075–077.

Intermediate goods: 021, 023, 026, 028, 033–038, 042, 044–074, 078–131, 138–146, 148–151, 153, 155–166, 172–174.

Regarding primary commodities, we use the sector "Crude petroleum and natural gas" (code 017) as X_1 and the sectors "Fishing" and "Seafood" (codes 012 and 020) as X_2. The rest of the primary commodities are grouped in X_3.

[23] Available online at https://comtrade.un.org/data.

[24] Available online at https://www.meti.go.jp.

B The model

This appendix fills in the details of the model described in Section 2, derives the real exchange rate equation (1), and describes the computation of the equilibrium.

We begin by describing the computation of the equilibrium. Given the assumption that trade is balanced on a period-by-period basis, the equilibrium can be computed as a sequence of static problems independent of the preferences of the households, given a value of the stochastic process $\left[z_t^1, z_t^1, \ln X_{1,t}^3, \ln X_{2,t}^3, \ln X_{3,t}^3\right]$. We normalize the price of primary commodity X_3 to one, $P_t^{x_3} = 1$, and iterate on the prices of primary commodities $P_t^{x_1}$ and $P_t^{x_2}$, and on the prices of intermediate goods $P_t^{q_1}$ and $P_t^{q_2}$, such that all markets clear.

Given a guess for the vector $[P_t^{x_1}, P_t^{x_2}, P_t^{q_1}, P_t^{q_2}]$, we can compute the other prices and allocations in the economy. We start with country 1. From the cost minimization problem of the firms, perfect competition implies that the prices of the final good $P_t^{y_1}$, intermediate good $P_t^{q_1}$, and primary commodities $P_t^{x_1}$ and $P_t^{x_3}$ are equal to their respective marginal costs. With Cobb-Douglas production functions, these are given by

$$P_t^{y_1} = \frac{1}{Z_t^1}\left(\frac{P_t^{q_1}}{\alpha_1^1}\right)^{\alpha_1^1}\left(\frac{P_t^{q_2}}{\alpha_2^1}\right)^{\alpha_2^1}\left(\frac{W_t^1}{\alpha_3^1}\right)^{\alpha_3^1}, \tag{B.1}$$

$$P_t^{q_1} = \frac{1}{Z_t^1}\left(\frac{P_t^{x_1}}{\beta_1^1}\right)^{\beta_1^1}\left(\frac{P_t^{x_2}}{\beta_2^1}\right)^{\beta_2^1}\left(\frac{P_t^{x_3}}{\beta_3^1}\right)^{\beta_3^1}\left(\frac{W_t^1}{\beta_4^1}\right)^{\beta_4^1}, \tag{B.2}$$

$$P_t^{x_1} = \frac{1}{Z_t^1}\left(\frac{P_t^{e_1^1}}{1-\phi_1^1}\right)^{1-\phi_1^1}\left(\frac{W_t^1}{\phi_1^1}\right)^{\phi_1^1}, \tag{B.3}$$

$$P_t^{x_3} = \frac{1}{Z_t^1}\left(\frac{P_t^{e_3^1}}{1-\phi_3^1}\right)^{1-\phi_3^1}\left(\frac{W_t^1}{\phi_3^1}\right)^{\phi_3^1}.$$

Given $[P_t^{x_1}, P_t^{x_2}, P_t^{q_1}, P_t^{q_2}]$ and the normalization $P_t^{x_3} = 1$, the above system consists of four equations and four unknowns: the price of the final good $P_t^{y_1}$, the wage rate W_t^1, and the prices of natural resources $P_t^{e_1^1}$ and $P_t^{e_3^1}$. The system is linear in logs, so is straightforward solving for the unknowns.

Once we know the prices in country 1, we solve for the allocation. Consumption is obtained from the budget constraint of the household,

$$C_t^1 = \frac{W_t^1}{P_t^{y_1}}n^1 + \frac{P_t^{e_1^1}}{P_t^{y_1}}e_1^1 + \frac{P_t^{e_3^1}}{P_t^{y_1}}e_3^1,$$

where n^1 is the endowment of labor in country 1. In equilibrium, $Y_t^1 = C_t^1$, so the previous equation also determines the level of final output in country 1. Next, we compute the inputs in the final good sector from the optimality conditions of the firms' problem:

$$q^1_{1,t} = \alpha^1_1 \frac{P^{y_1}_t Y^1_t}{P^{q_1}_t},$$

$$q^1_{2,t} = \alpha^1_2 \frac{P^{y_1}_t Y^1_t}{P^{q_2}_t},$$

$$n^1_{y,t} = \alpha^1_3 \frac{P^{y_1}_t Y^1_t}{W^1_t}.$$

Similarly, we can solve for the allocations in the primary commodity sectors using the optimality conditions. In this case, however, the inputs of natural resources are equal to the exogenously given endowments, so we can solve for the labor inputs and total production of primary commodities $X^1_{1,t}$ and $X^1_{3,t}$ as follows:

$$X^1_{1,t} = \frac{P^{e^1_1}_t e^1_1}{(1-\phi^1_1) P^{x_1}_t},$$

$$n^1_{x_1,t} = \phi^1_1 \frac{P^{x_1}_t X^1_{1,t}}{W^1_t},$$

$$X^1_{3,t} = \frac{P^{e^1_3}_t e^1_3}{(1-\phi^1_3) P^{x_3}_t},$$

$$n^1_{x_3,t} = \phi^1_3 \frac{P^{x_3}_t X^1_{3,t}}{W^1_t}.$$

Next, we use market clearing in the labor market to solve for the labor input allocated to the intermediate good sector:

$$n^1_{q_1,t} = n^1 - n^1_{y_1,t} - n^1_{x_1,t} - n^1_{x_3,t}.$$

With the labor input in the intermediate good sector, we can solve for the total production of intermediate good Q^1_t, as well as for the demand for primary commodities $q^1_{1,t}$, $q^1_{2,t}$, and

$q^1_{3,t}$ using the optimality conditions of the firms' problem:

$$
\begin{aligned}
Q^1_t &= \frac{W^1_t n^1_{q_1,t}}{\beta^1_4 P^{q_1}_t}, \\
x^1_{1,t} &= \beta^1_1 \frac{P^{q_1}_t Q^1_t}{P^{x_1}_t}, \\
x^1_{2,t} &= \beta^1_2 \frac{P^{q_1}_t Q^1_t}{P^{x_2}_t}, \\
x^1_{3,t} &= \beta^1_3 \frac{P^{q_1}_t Q^1_t}{P^{x_3}_t}.
\end{aligned}
$$

Using an analogous procedure, we compute prices and allocations in country 2. So we now turn to the computation of prices and allocations in country 3. The optimality condition of the firms' problem with respect to labor input, the household budget constraint, and the market clearing condition for labor implies that the following conditions must hold in equilibrium:

$$
\begin{aligned}
\pi_4 P^{y_3}_t Y^3_t &= W^3_t n^3_{y,t}, \\
P^{y_3}_t Y^3_t &= P^{x_1}_t X^3_{1,t} + P^{x_2}_t X^3_{2,t} + P^{x_3}_t X^3_{3,t} + W^3_t n^3, \\
n^3_{y,t} &= n^3.
\end{aligned}
$$

We can use these equations to solve for nominal wages W^3_t as a function of the commodity prices and endowments of primary commodities in country 3:

$$
W^3_t = \left(\frac{\pi_4}{1-\pi_4}\right) \frac{P^{x_1}_t X^3_{1,t} + P^{x_2}_t X^3_{2,t} + P^{x_3}_t X^3_{3,t}}{n^3}.
$$

From the cost minimization problem of the final good firms, it follows that the price of the final good in country 3 is given by

$$
P^{y_3}_t = \left(\frac{P^{x_1}_t}{\pi_1}\right)^{\pi_1} \left(\frac{P^{x_2}_t}{\pi_2}\right)^{\pi_2} \left(\frac{P^{x_3}_t}{\pi_3}\right)^{\pi_3} \left(\frac{W^3_t}{\pi_4}\right)^{\pi_4}.
$$

With the price of the final good P^{y_3} and nominal wage W^3, we compute the output of the final good in country 3:

$$
Y^3_t = \frac{W^3_t n^3}{\pi_4 P^{y_3}_t}.
$$

Finally, with the price of the final good and the exogenously given endowments of primary commodities in country 3, we solve for the demand for each primary commodity:

$$x_{1,t}^3 = \pi_1 \frac{P_t^{y_3} Y_t^3}{P_t^{x_1}},$$

$$x_{2,t}^3 = \pi_2 \frac{P_t^{y_3} Y_t^3}{P_t^{x_2}},$$

$$x_{3,t}^3 = \pi_3 \frac{P_t^{y_3} Y_t^3}{P_t^{x_3}}.$$

All the previous prices and quantities were computed given a guess for the vector of prices $[P_t^{x_1}, P_t^{x_2}, P_t^{q_1}, P_t^{q_2}]$. We solve for these prices imposing market clearing in the world market of commodities and intermediate goods,[25]

$$\begin{aligned} X_{1,t}^1 + X_{1,t}^3 &= x_{1,t}^1 + x_{1,t}^2 + x_{1,t}^3, \\ X_{3,t}^1 + X_{3,t}^2 + X_{3,t}^3 &= x_{3,t}^1 + x_{3,t}^2 + x_{3,t}^3, \\ Q_t^1 &= q_{1,t}^1 + q_{1,t}^2, \\ Q_t^2 &= q_{2,t}^1 + q_{2,t}^2. \end{aligned}$$

Real exchange rate. The bilateral real exchange rate between countries 1 and 2 is defined as the price of the final good in country 1, $P_t^{y_1}$, relative to the price of the final good in country 2, $P_t^{y_2}$. Let lowercase letters denote log values. Using equation (B.1) and ignoring the constant terms, the log of the real exchange rate, ξ, can be expressed as

$$\xi_t = z_t^2 - z_t^1 + (\alpha_1^1 - \alpha_1^2) p_t^{q_1} + (\alpha_2^1 - \alpha_2^2) p_t^{q_2} + \alpha_3^1 w_t^1 - \alpha_3^2 w_t^2. \tag{B.4}$$

Next, we use equation (B.2) for country 1 and the analogous equation for country 2 to substitute for the price of intermediate goods in equation (B.4). Ignoring the constant terms, the expression for the real exchange rate becomes

$$\begin{aligned} \xi_t = & \left(1 - (\alpha_2^1 - \alpha_2^2)\right) z_t^2 - \left(1 + (\alpha_1^1 - \alpha_1^2)\right) z_t^1 + \left((\alpha_1^1 - \alpha_1^2)\beta_1^1 + (\alpha_2^1 - \alpha_2^2)\beta_1^2\right) p_t^{x_1} \\ & + \left((\alpha_1^1 - \alpha_1^2)\beta_2^1 + (\alpha_2^1 - \alpha_2^2)\beta_2^2\right) p_t^{x_2} + \left((\alpha_1^1 - \alpha_1^2)\beta_3^1 + (\alpha_2^1 - \alpha_2^2)\beta_3^2\right) p_t^{x_3} \\ & + \left((\alpha_1^1 - \alpha_1^2)\beta_4^1 + \alpha_3^1\right) w_t^1 + \left((\alpha_2^1 - \alpha_2^2)\beta_4^2 - \alpha_3^2\right) w_t^2. \end{aligned}$$

Finally, we substitute for wages using the log of equation (B.3) for country 1 and the

[25] Walras's law implies that the market clearing condition for the primary commodity X_2 is also satisfied.

analogous equation for country 2. Ignoring the constant terms, that leads to equation (1) in the main paper:

$$\begin{aligned}
\xi_t = {} & \left(1 - (\alpha_2^1 - \alpha_2^2) + \frac{(\alpha_2^1 - \alpha_2^2)\beta_4^2 - \alpha_3^2}{\phi_2^2}\right) z_t^2 - \left(1 + (\alpha_1^1 - \alpha_1^2) - \frac{(\alpha_1^1 - \alpha_1^2)\beta_4^1 + \alpha_3^1}{\phi_1^1}\right) z_t^1 \\
& - \frac{1-\phi_1^1}{\phi_1^1}\left((\alpha_1^1 - \alpha_1^2)\beta_4^1 + \alpha_3^1\right) p_t^{e_1} - \frac{1-\phi_2^2}{\phi_2^2}\left((\alpha_2^1 - \alpha_2^2)\beta_4^2 - \alpha_3^2\right) p_t^{e_2} \\
& + \left((\alpha_1^1 - \alpha_1^2)\beta_1^1 + (\alpha_2^1 - \alpha_2^2)\beta_1^2 + \frac{(\alpha_1^1 - \alpha_1^2)\beta_4^1 + \alpha_3^1}{\phi_1^1}\right) p_t^{x_1} \\
& + \left((\alpha_1^1 - \alpha_1^2)\beta_2^1 + (\alpha_2^1 - \alpha_2^2)\beta_2^2 + \frac{(\alpha_2^1 - \alpha_2^2)\beta_4^2 - \alpha_3^2}{\phi_2^2}\right) p_t^{x_2} \\
& + \left((\alpha_1^1 - \alpha_1^2)\beta_3^1 + (\alpha_2^1 - \alpha_2^2)\beta_3^2\right) p_t^{x_3}.
\end{aligned} \tag{B.5}$$

ONLINE APPENDIX

(not for publication)

C Additional figures and tables

Table C.1 reports the results of unit root tests for the data in levels and in three-, four-, and five-year differences. There is evidence of unit roots for the raw data, but it vanishes for the data in four-year differences.

Table C.1: Unit root tests (p-values)

	Level	three-year differences	four-year differences	five-year differences
Real Exchange Rates				
US-UK	0.018	0.001	0.003	0.003
US-DEU	0.117	0.005	0.028	0.027
US-JPN	0.809	0.001	0.018	0.027
Commodities				
Oil	0.485	0.079	0.128	0.356
Fish	0.352	0.001	0.027	0.009
Meat	0.523	0.019	0.047	0.304
Aluminum	0.145	0.001	0.001	0.003
Copper	0.319	0.009	0.025	0.103
Gold	0.508	0.001	0.016	0.025
Wheat	0.226	0.001	0.005	0.009
Maize	0.269	0.001	0.010	0.013
Timber	0.047	0.003	0.018	0.047
Cotton	0.592	0.005	0.016	0.015

Notes: Variables are in logs, and commodity prices are normalized by US CPI. We use the Dickey-Fuller test, in which the p-values are under the null hypothesis that the series follows a unit root process. The lag length is selected according to the Ng-Perron test. We assume a trend in the case of Japan. Cointegration tests such as Johansen (1991) or Stock and Watson (1993) do not provide evidence of cointegration between real exchange rates and primary commodity prices.

In Table C.2, we report the first-order autocorrelation for all the series in four-year differences, which is our benchmark case. As can be clearly seen, the high persistence of real exchange rates is also present in the commodity prices.

Table C.3 shows the volatility (standard deviation) of the bilateral real exchange rates and the commodity prices in log-levels. This table provides the same messages as Table 1 in the main paper.

Table C.2: First-order autocorrelation of four-year differences

US-UK	US-DEU	US-JPN	Oil	Fish	Meat	Aluminum
0.98	0.98	0.98	0.97	0.98	0.97	0.98
Copper	Gold	Wheat	Maize	Timber	Cotton	
0.98	0.99	0.98	0.97	0.97	0.98	

Note: Variables are in logs, and commodity prices are normalized by US CPI.

Table C.3: Volatilities of real exchange rates and primary commodity prices

	1960–2014	1960–1972	1973–1985	1986–1998	1999–2014
			(a) Levels		
Real exchange rates:					
US-UK	0.12	0.06	0.15	0.08	0.08
US-DEU	0.18	0.07	0.21	0.09	0.13
US-JPN	0.37	0.13	0.13	0.12	0.11
Average across commodities:					
Simple	0.44	0.13	0.31	0.22	0.36
Trade weighted	0.57	0.13	0.37	0.23	0.46

Notes: Variables are in logs, and commodity prices are normalized by US CPI. Weights are based on the share of total trade in 1990. The set of primary commodities is oil, fish, meat, aluminum, copper, gold, wheat, maize, timber, and cotton.

Table C.4 shows the simple correlations of each of the bilateral RER and all the commodity prices we use. As can be seen, all simple correlations between the prices and the RER are sizable. In addition, the correlations across the PCP are also sizable in many cases.

One concern about regression (2) is that the variables are expressed in constant US dollars, so the US CPI appears on both sides of the equation. If its volatility is sufficiently large relative to the volatility of the nominal exchange rate and foreign CPI, that would imply large R^2s. In panel (e) of Table C.5, we show that this is not the case. The table shows the same results as in panel (b) of Table 2, but for the case in which we use variables expressed in current US dollars; that is, we do not subtract the log of US CPI from either side of equation (2). [26] The results are invariant to whether we use variables in current or constant US dollars.

Another concern regarding regression (2) is that commodity prices are expressed in US dollars, so they might contain the nominal exchange rate, which in turn would imply that the nominal exchange rate appears on both sides of equation (2). Again, Table C.5 shows

[26] This procedure is correct to the extent that the sum of the coefficients that multiply the price of the US CPI on the right-hand side is 1 in all cases. The model in Section 2 rationalizes that restriction.

Table C.4: Contemporaneous correlations (1960–2014)

	Oil	Fish	Meat	Alum.	Copper	Gold	Wheat	Maize	Timber	Cotton
RER										
US-UK	-0.47	0.00	0.30	0.11	0.09	-0.53	0.26	0.36	-0.40	0.30
US-DEU	-0.51	-0.24	0.16	0.08	-0.08	-0.62	0.06	0.14	-0.58	0.11
US-JPN	-0.49	0.25	0.59	0.52	0.41	-0.63	0.59	0.63	-0.44	0.55
Commodities										
Oil	1.00									
Fish	0.28	1.00								
Meat	-0.17	0.45	1.00							
Alum.	-0.20	0.36	0.73	1.00						
Copper	0.07	0.72	0.57	0.52	1.00					
Gold	0.88	0.25	-0.16	-0.22	0.03	1.00				
Wheat	-0.05	0.57	0.78	0.70	0.60	-0.07	1.00			
Maize	-0.11	0.58	0.81	0.70	0.62	-0.13	0.94	1.00		
Timber	0.39	0.10	0.18	0.17	0.10	0.56	0.21	0.15	1.00	
Cotton	-0.23	0.39	0.83	0.78	0.43	-0.21	0.84	0.86	0.24	1.00

Note: Variables are in logs, and primary commodity prices are normalized by US CPI.

that this is not the case. Panel (f) of the table shows the results for the case in which we run the regressions for the bilateral real exchange rates without including the United States. That is, we run the regression in (2) for the bilateral real exchange rates of the United Kingdom versus Germany and Japan, and for Germany versus Japan. The results show that four primary commodities still account for a large fraction of these bilateral real exchange rate fluctuations. Moreover, this result holds true for the whole period as well as for the subperiods.

C.1 Selecting commodities based on US trade data

In Section 5, we showed the results with the four PCP that have the best fit with the real exchange rates. This set (possibly) varies by country pair and subperiod, and whether we use data in levels or in four-year differences. In this section, we explore an alternative approach based on the theory presented above: we choose the set of commodities based on US trade data and keep it fixed for all subperiods and country pairs.

Equation (1) shows that a necessary condition for the primary commodity price to explain a large fraction of real exchange rate fluctuations, is that the commodity price must be an important input (i.e. must have a large share) in the production structure of one of the economies in the country pair. As we mentioned before, the difference in shares is what is crucial, but it can be observed only if the primary commodity is an important input in at

Table C.5: Coefficients of determination R^2

	1960–2014	1960–1972	1973–1985	1986–1998	1999–2014
(a) 10 commodities, level					
United Kingdom	0.50	0.76	0.73	0.67	0.54
Germany	0.59	0.87	0.86	0.57	0.67
Japan	0.81	0.87	0.60	0.75	0.75
(b) 10 commodities, 4-year differences					
United Kingdom	0.48	0.90	0.90	0.81	0.60
Germany	0.63	0.95	0.87	0.83	0.75
Japan	0.57	0.92	0.84	0.92	0.82
(c) 4 commodities (best fit), level					
United Kingdom	0.39	0.66	0.70	0.51	0.51
Germany	0.56	0.77	0.83	0.38	0.66
Japan	0.79	0.85	0.57	0.67	0.66
(d) 4 commodities (best fit), 4-year differences					
United Kingdom	0.33	0.72	0.82	0.63	0.58
Germany	0.56	0.84	0.87	0.81	0.74
Japan	0.48	0.88	0.76	0.86	0.80
(e) Nominal values					
United Kingdom	0.41	0.60	0.82	0.67	0.63
Germany	0.59	0.86	0.87	0.81	0.75
Japan	0.57	0.89	0.76	0.77	0.79
(f) Non-US pairs					
UK-DEU	0.37	0.73	0.53	0.68	0.72
UK-JPN	0.42	0.64	0.56	0.80	0.68
DEU-JPN	0.35	0.79	0.46	0.69	0.68

least one of the economies. Based on that information, in this section we show the same set of results as in Section 5, but this time we choose as regressors the four commodities with the largest trade share for the United States. This is, admittedly, a very crude approximation to the data using the model of the previous section, but it has the advantage that the four commodities have not been chosen to fit the data. We see this exercise as a first approximation to using the model to discipline our choices in the empirical analysis.

Table A.2 in Appendix A.1 shows the trade data for each country and commodity that we analyze in this paper. We choose the four commodities that are the most traded in the United States according to Table A.2: petroleum, fish, timber, and gold. We report the results in Tables C.6 and C.7 and in Figure C.1.

As can be seen, the results when we choose the set of four PCP based on US trade data are still very good. However, it is important to discuss some of the differences. First, the R^2s in Table C.7 are lower than in the case for the best-fit commodities (Table 3). This result should be expected, since the four commodities were chosen to maximize R^2. But, with the exception of the United Kingdom, the differences are not very large. Second, the bootstrap exercise based on the out-of-sample-fit exercise shows very similar results for Germany and Japan and somewhat worse results for the United Kingdom. Indeed, the main difference with the previous analysis is the out-of-sample-fit for the United Kingdom: the curve is within the 90% confidence interval, very marginally so for the first months but getting worse as the period length grows longer.

Table C.6: R^2 with four commodities (largest US trade share)

	1960–2014	1960–1972	1973–1985	1986–1998	1999–2014
			(a) Level		
United Kingdom	0.32	0.52	0.63	0.38	0.34
Germany	0.48	0.50	0.55	0.19	0.56
Japan	0.59	0.21	0.18	0.55	0.58
			(b) Four-year differences		
United Kingdom	0.25	0.63	0.73	0.25	0.46
Germany	0.53	0.89	0.71	0.50	0.69
Japan	0.44	0.71	0.52	0.82	0.68

Table C.7: Bootstrap distributions of R^2 under the null hypothesis of orthogonality, with four commodities (largest US trade share) in four-year differences

	$\hat{R}^2$	Percentiles distribution of R^2				$\Pr(R^2 \geq \hat{R}^2)$
		Median	75	90	95	
United Kingdom						
1960-2014	0.25	0.13	0.20	0.27	0.31	0.134
1960-1972	0.63	0.51	0.64	0.74	0.79	0.274
1973-1985	0.73	0.44	0.59	0.70	0.76	0.070
1986-1998	0.25	0.37	0.50	0.62	0.68	0.734
1999-2014	0.46	0.30	0.43	0.54	0.59	0.188
Germany						
1960-2014	0.53	0.13	0.20	0.27	0.33	0.000
1960-1972	0.89	0.53	0.68	0.77	0.81	0.005
1973-1985	0.71	0.45	0.60	0.72	0.77	0.106
1986-1998	0.50	0.40	0.54	0.66	0.71	0.313
1999-2014	0.69	0.33	0.47	0.58	0.64	0.024
Japan						
1960-2014	0.44	0.13	0.20	0.27	0.32	0.007
1960-1972	0.71	0.54	0.67	0.77	0.82	0.190
1973-1985	0.52	0.47	0.62	0.74	0.79	0.421
1986-1998	0.82	0.42	0.56	0.67	0.73	0.007
1999-2014	0.68	0.35	0.48	0.59	0.65	0.033

Figure C.1: Fitted correlations and bootstrap bands under the null hypothesis of orthogonality (with four commodities, largest US trade share)

(a) United Kingdom

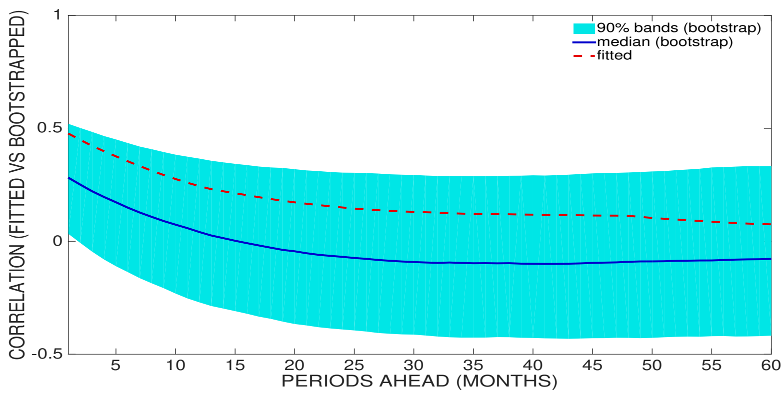

(b) Germany

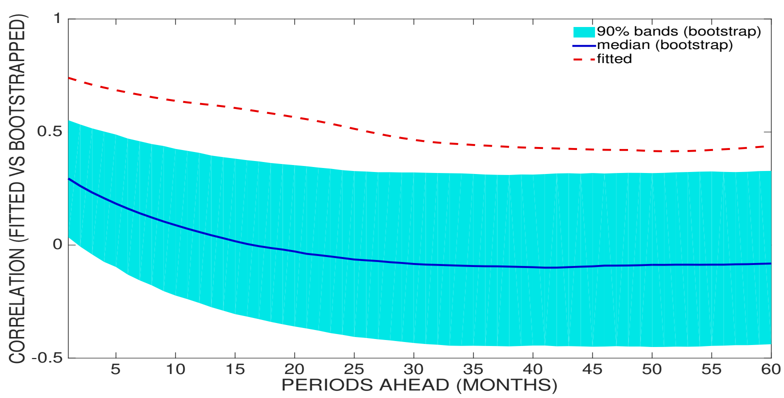

(c) Japan

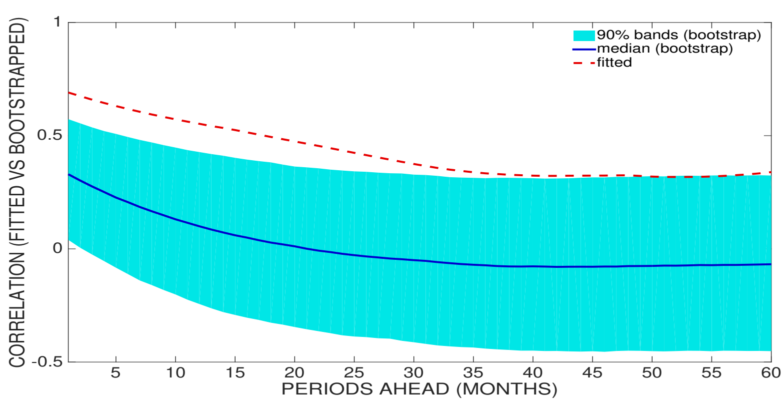

D Real exchange rates and fitted values (level)

Figure D.1: Real exchange rates and fitted values, level.

(a) United Kingdom

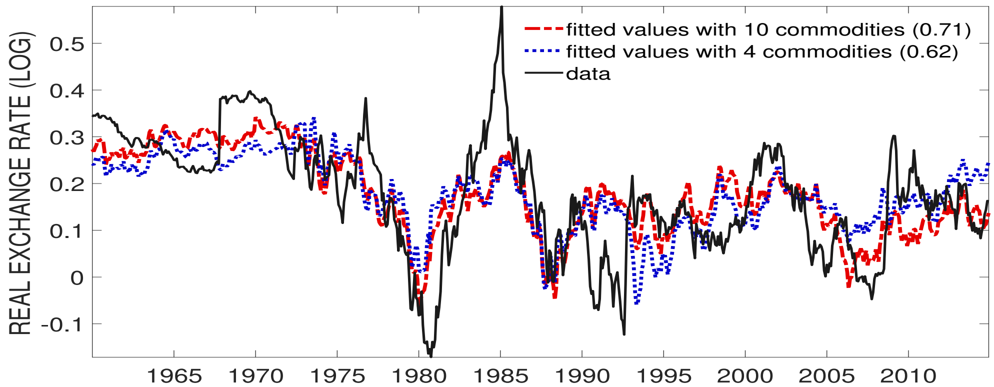

(b) Germany

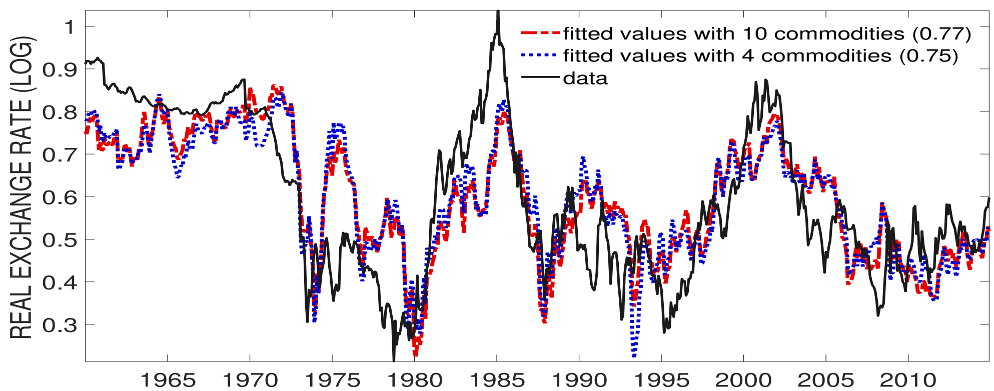

(c) Japan

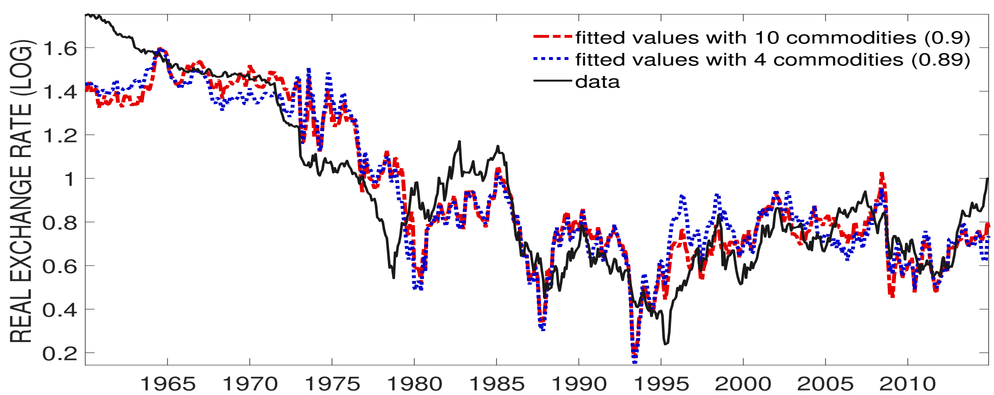

E R^2s as a function of the number of months for which we take the differences

Figure E.1: R^2s as a function of the number of months for which we take the differences, four commodities (best fit), 1960–2014

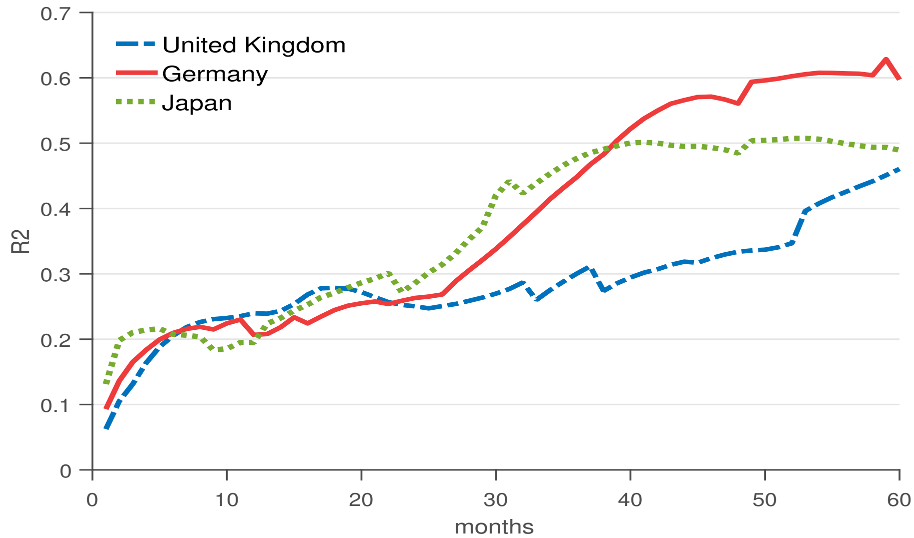

F Regression coefficients

Table F.1: Regression coefficients: United Kingdom

	1960-2014	1960-1972	1973-1985	1986-1998	1999-2014
(a) Level					
Oil		-0.317***		-0.033	
Fish				-0.189***	
Meat	0.078		-0.204***		-0.035
Aluminum	-0.113**	-0.045	-0.198**		
Copper				-0.178***	-0.186***
Gold			-0.250***		0.228
Wheat				0.254***	-0.122***
Maize	0.120***	-0.280***			
Timber	-0.230***	-0.010			
Cotton			0.055		
(b) Four-year differences					
Oil				-0.012	
Fish	-0.213***	-0.252***		-0.218***	-0.162**
Meat		-0.176**	-0.322***	-0.457***	-0.139***
Aluminum	-0.197***	-0.456***			
Copper			0.402***	-0.127***	-0.196***
Gold			-0.329***		0.290***
Wheat					
Maize	0.121**				
Timber			-0.211***		
Cotton	0.027	-0.439***			

Note: Superscripts *, **, and *** denote statistical significance at 10%, 5%, and 1%, respectively.

Table F.2: Regression coefficients: Germany

	1960-2014	1960-1972	1973-1985	1986-1998	1999-2014
(a) Level					
Oil		-0.338***			
Fish	-0.234***				
Meat	0.095			0.218***	0.193**
Aluminum		0.575***	-0.612***		
Copper		-0.085***		-0.150**	-0.151***
Gold			0.120***	-0.063	-0.033
Wheat			-0.449***		-0.106**
Maize	0.139		0.302*		
Timber	-0.450***	0.088*		-0.138***	
Cotton					
(b) Four-year differences					
Oil					
Fish	-0.319***	-0.095***	-0.194**		-0.375***
Meat			-0.005		
Aluminum			-0.541***		
Copper				-0.231***	
Gold				-0.138**	-0.186***
Wheat	-0.137*		-0.233***		-0.279***
Maize	0.223*	-0.032		0.252***	0.316**
Timber	-0.314***	0.282***		-0.149***	
Cotton		-0.537***			

Note: Superscripts *, **, and *** denote statistical significance at 10%, 5%, and 1%, respectively.

Table F.3: Regression coefficients: Japan

	1960-2014	1960-1972	1973-1985	1986-1998	1999-2014
			(a) Level		
Oil		-0.331***	0.018		
Fish					0.393***
Meat		-0.313**	-0.090*	0.299***	
Aluminum		0.932***	-0.375***		
Copper		-0.124**		-0.137**	
Gold	-0.238***			-0.122**	-0.327***
Wheat	0.289**				0.200***
Maize	0.282***		0.239***		
Timber	-0.529***			-0.371***	
Cotton					-0.171***
			(b) Four-year differences		
Oil	0.173**			0.270***	
Fish			-0.318***	0.241***	0.346***
Meat			-0.027		
Aluminum	-0.170**		-0.434***	-0.260***	
Copper		-0.029***	0.109		
Gold		-0.290***			-0.532***
Wheat					0.205***
Maize		-0.244***			
Timber	-0.396***			-0.352***	
Cotton	-0.124*	-0.318***			-0.130***

Note: Superscripts *, **, and *** denote statistical significance at 10%, 5%, and 1%, respectively.

G Selecting three commodities (best fit)

Table G.1: R^2 with three commodities (best fit)

	1960–2014	1960–1972	1973–1985	1986–1998	1999–2014
			(a) Level		
United Kingdom	0.37	0.66	0.70	0.51	0.51
Germany	0.54	0.66	0.81	0.37	0.65
Japan	0.78	0.78	0.42	0.66	0.38
			(b) Four-year differences		
United Kingdom	0.24	0.68	0.78	0.63	0.39
Germany	0.53	0.84	0.80	0.72	0.72
Japan	0.46	0.81	0.58	0.73	0.69

Table G.2: Bootstrap distributions of R^2 under the null hypothesis of orthogonality, with three commodities (best fit) in four-year differences

		Percentiles distribution of R^2				
	$\hat{R}^2$	Median	75	90	95	$\Pr(R^2 \geq \hat{R}^2)$
United Kingdom						
1960-2014	0.24	0.09	0.16	0.23	0.28	0.088
1960-1972	0.68	0.48	0.64	0.75	0.80	0.187
1973-1985	0.78	0.32	0.47	0.60	0.67	0.006
1986-1998	0.63	0.30	0.45	0.57	0.63	0.050
1999-2014	0.39	0.24	0.36	0.47	0.54	0.203
Germany						
1960-2014	0.53	0.10	0.17	0.24	0.29	0.000
1960-1972	0.84	0.37	0.52	0.64	0.70	0.002
1973-1985	0.80	0.31	0.47	0.60	0.66	0.005
1986-1998	0.72	0.33	0.48	0.61	0.67	0.024
1999-2014	0.72	0.23	0.36	0.49	0.57	0.005
Japan						
1960-2014	0.46	0.10	0.17	0.24	0.29	0.003
1960-1972	0.81	0.43	0.59	0.71	0.77	0.024
1973-1985	0.58	0.32	0.48	0.61	0.67	0.130
1986-1998	0.73	0.30	0.44	0.56	0.63	0.011
1999-2014	0.69	0.26	0.40	0.52	0.59	0.012

Figure G.1: Fitted correlations and bootstrap bands under the null hypothesis of orthogonality (with three commodities, best fit)

(a) United Kingdom

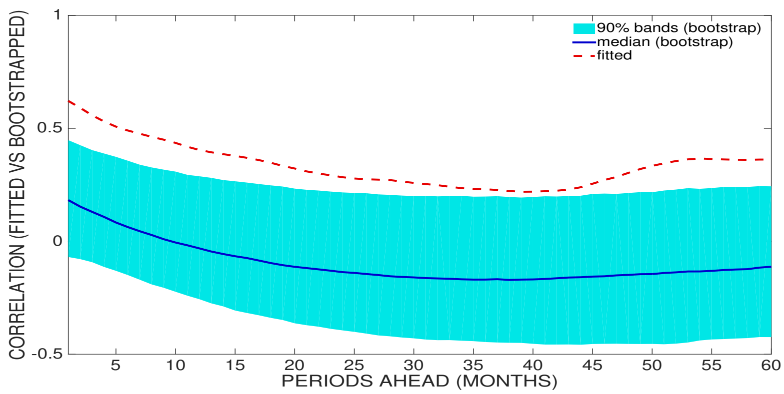

(b) Germany

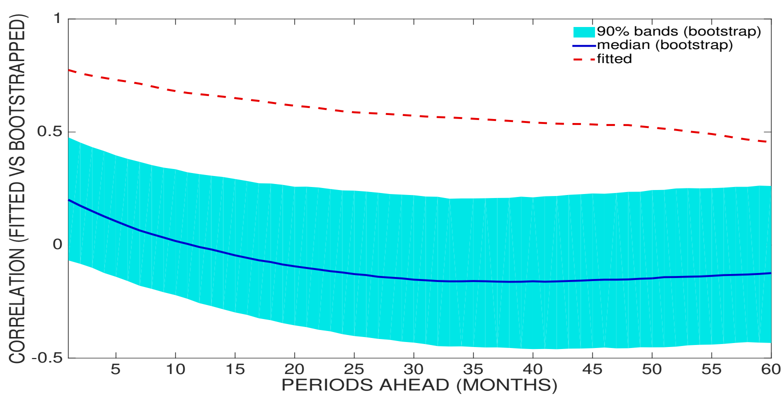

(c) Japan