Real Exchange Rates and Primary Commodity Prices: Mussa Meets Backus-Smith

Joao Ayres (Inter-American Development Bank)

Constantino Hevia (Universidad Torcuato Di Tella)

Juan Pablo Nicolini (Federal Reserve Bank of Minneapolis/UTDT)

Real Exchange Rates and Primary Commodity Prices: Mussa Meets Backus-Smith[*]

Joao Ayres

Inter-American Development Bank

joaoay@iadb.org

Constantino Hevia

Universidad Torcuato Di Tella

chevia@utdt.edu

Juan Pablo Nicolini

Federal Reserve Bank of Minneapolis and Universidad Torcuato Di Tella

juanpa@minneapolisfed.org

September 22, 2021

Abstract

We show that explicitly modeling primary commodities in an otherwise totally standard incomplete markets open economy model can go a long way in explaining the Mussa puzzle and the Backus-Smith puzzle, two of the main puzzles in the international economics literature.

Keywords: primary commodity prices, Mussa puzzle, Backus-Smith puzzle.
JEL classification: F31, F41.

1   Introduction

Primary commodity prices are known to be very volatile and very persistent. This has been especially true since the oil price shocks of the early 1970s. The small open economy literature has long recognized the effect that shocks to commodity prices have on the economy in general and on the real exchange rate in particular. Their importance has been identified so strongly in the literature that it is commonly acknowledged in the policy debate in small commodity-producing countries.

In contrast, in studies of the behavior of the real exchange rate between large developed economies, the role of these markets has been largely ignored, particularly so during the last three decades.[1] The most probable reason is that the value added of primary commodities in total economic activity is small in these countries.

We show that although primary commodities' share in total output may appear small, the volatility of primary commodity prices is so high that it can potentially have large effects on real exchange rates.

In a recent paper (Ayres, Hevia and Nicolini, 2020), we documented a strong and robust co-movement between the real exchange rates of Germany, Japan and the UK against the US dollar and a handful of primary commodity prices during the last half century. We also showed that a simple static model had the potential to deliver much higher volatility and persistence in real exchange rates than a model that ignores the commodity sector.[2] In this paper, we go a few steps beyond, solve a truly dynamic model, and quantitatively address two famous puzzles in international economics.

The first one, known as the Mussa puzzle, documents a substantial increase in the volatil-

---

[1] An earlier literature did discuss and evaluate the role of primary commodity markets in the behavior of real exchange rates among developed economies. 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 3. On the empirical side, Sachs (1985) and Dornbusch (1985*a*,*b*, 1987) are important contributions.

[2] The model is formally dynamic, but we impose a zero trade balance every period, so the solution of the model is static.

ity of real exchange rates since the breakdown of the Bretton Woods system of fixed exchange rates in 1973.

The second one, known as the Backus-Smith puzzle, documents a low correlation, and in many cases a negative one, between the bilateral real exchange rate of any two given countries and the ratio of consumption between those same two countries. In complete markets models, this correlation ought to be very close to 1.[3]

We also show another strong pattern in the data, which we call the "Mussa meets Backus-Smith puzzle". This is the generalized fact that the correlation between the real exchange rate and the ratio of consumption is even lower in the period following 1973 when compared with the period before. In other words, the Backus-Smith puzzle becomes quantitatively larger in the post-1973 period than in the period before. This feature of the data has received much less attention in the literature.[4]

The key ingredient of the model is the interaction between incomplete markets and shocks that move prices of primary commodities.[5] We document how the transmission mechanism implied by this interaction helps reconcile the main puzzles.

We study a labor-only open economy three-country model. Each country produces a final nontraded good, a traded intermediate good and (potentially) three primary commodities. Labor is used in all technologies, commodities are used to produce the intermediate good, and the intermediate goods are used to produce the final good. We show that as long as countries have different production structures, the real exchange rate is affected by shocks that change primary commodity prices.

We illustrate the mechanism in a simplified version of the model that can be solved analytically. We then calibrate the general version of the model for the USA and Japan, because as we show, this is the country pair for which the puzzles are quantitatively more

---

[3]Results similar to the ones in Backus and Smith (1993) were independently reported in Robert Kollman's unpublished PhD dissertation in 1990. The main results were published in Kollmann (1995).

[4]Two notable exceptions are Colacito and Croce (2013) and Itskhoki and Mukhin (2019).

[5]Incomplete markets are essential to study the Backus-Smith puzzle and are also important in generating volatile real exchange rates, as shown by Heathcote and Perri (2002).

striking, and because the Ministry of Economy, Trade, and Industry of Japan publishes a bilateral Japan-US input-output table that we use to discipline our calibration. The third country is taken to be the rest of the world. The rest of the world is subject to shocks in its excess demand for commodities. The key step in our calibration is to choose the volatility of these shocks in order to match the volatility and persistence of the primary commodity prices we see in the data for each subperiod. We want to emphasize at the outset that we do not explain why the prices of primary commodities are so volatile.

A key aspect of our calibration that we wish to highlight is that the sizes of value added created in the primary commodity sectors in the model are as small as they are in the data. Our analysis shows that, in spite its size, the high volatility of the shocks to primary commodity markets has a substantial effect on the real exchange rate and on its correlation with the ratio of consumption. And these effects are more pronounced after 1973, when primary commodity prices became more volatile. In a nutshell, the benchmark calibrated model can account for a large fraction of the puzzles mentioned above.

Specifically, we show that by calibrating the model pre- and post-1973 (the end of the Bretton Woods period), the model generates substantially more volatile real exchange rates in the post-Bretton Woods period (the Mussa puzzle), as observed in the data. The calibrated model also exhibits a correlation between the real exchange rate and the ratio of consumption much lower than one. Interestingly, the model also generates a lower correlation of the real exchange rate and the ratio of consumption for the post-1973 period (the Mussa meets Backus-Smith puzzle), as in the data.

In describing the stylized facts, we first show that the puzzles hold very generally. We illustrate this fact using data for 15 bilateral pairs between the USA and 15 other OECD countries.

We also show that these moments have substantial statistical uncertainty due to the high persistence of both the bilateral real exchange rates and the primary commodity prices. For this reason, we focus on the US-Japan bilateral relation and compute the entire small sample

distribution of the statistics of interest.

We find these distributions to be quite dispersed and skewed. For example, the point estimate of the standard deviation of the US-Japan real exchange rate after 1973 is about 19 percentage points, but the associated small sample distribution implies that one could quite likely observe values that are 50 percent smaller or 50 percent larger than that. For this reason, we focus on the entire small sample distributions of the statistics of interest instead of the point estimate. To evaluate the performance of our model, we compare these estimated distributions with analogous distributions obtained from model simulations.

The model has no frictions other than the lack of contingent asset markets. Thus, it is unable to match the deviations from the uncovered interest rate parity observed in the data— just one risk-free bond, as we assume, is enough for the UIP to hold in the model. Thus, we see the results of our paper as complementary to efforts in the literature that explore frictions in the setting of prices (as in Chari, Kehoe and McGrattan (2002) as well as many others) or segmentation in asset markets (such as the recent work by Itskhoki and Mukhin (2019, 2021)).

To highlight the transmission mechanism, we choose to keep the model as simple as possible. Thus, there is no physical capital in the model, and production functions are assumed to be Cobb-Douglas except in the production of primary commodities, for which we allow the elasticity of substitution to be smaller than one, as suggested by the evidence. We also ignore realistic microeconomic features like time-to-build restrictions in investment, the role of inventories and political economy considerations, which are very relevant for understanding primary commodity markets, as emphasized, for instance, in Baumeister and Kilian (2016).

The paper proceeds as follows. In Section 2, we describe the main features of the data. In Section 3, we briefly present the model, which, as mentioned earlier, is very standard. We discuss in detail a simplified version of the model that can be solved analytically and helps develop intuition for the transmission mechanism. It also highlights the type of demand

and supply shocks that can change primary commodity prices. In Section 4, we discuss the calibration. We calibrate the volatility and persistence of supply shocks to match the volatility and persistence of primary commodity prices in the data in each subperiod (pre- and post-1973). This is without loss of generality, since as the simple model that can be solved in closed form shows, the same volatility can be generated by demand shocks.[6] Section 5 presents the results.

2   The stylized facts

In this section, we document the main features of the data. We first describe point estimates for bilateral relationships between the USA and a group of 15 OECD countries. We then estimate the distributions for those estimates between the USA and Japan. We start our analysis in 1960 and focus on the differential behavior of the statistics of interest before and after 1973. Before entering into specifics, we discuss next how we choose to detrend the data.

After World War II, most industrial countries displayed a strong convergence in their income level to that of the United States. This convergence in income levels was simultane- ously accompanied by a Balassa-Samuelson effect whereby the bilateral real exchange rates against the USA appreciated over time. Importantly, most of the convergence in output levels and its corresponding Balassa-Samuelson effect occurred in the first decade and a half. Since the mid-1970s, relative incomes and consumption have tended to be roughly stable across developed countries, and the bilateral real exchange rates have tended to fluctuate around a relatively stable level.

To illustrate these points, the left panel of Figure 1 shows the log of the bilateral real exchange rate between the USA and Japan at a quarterly frequency from 1960:Q1 through 2019:Q4. The right panel shows the log of consumption per working age population in Japan

---

[6]This is the reason why we do not need to take a stand on the fascinating ongoing empirical debate over the nature of the shocks driving the volatility of oil prices. See Baumeister and Kilian (2016), Baumeister and Hamilton (2019), and the references therein.

Figure 1: USA and Japan: real exchange rates and relative consumptions

(a) Real exchange rate

(b) Relative consumption
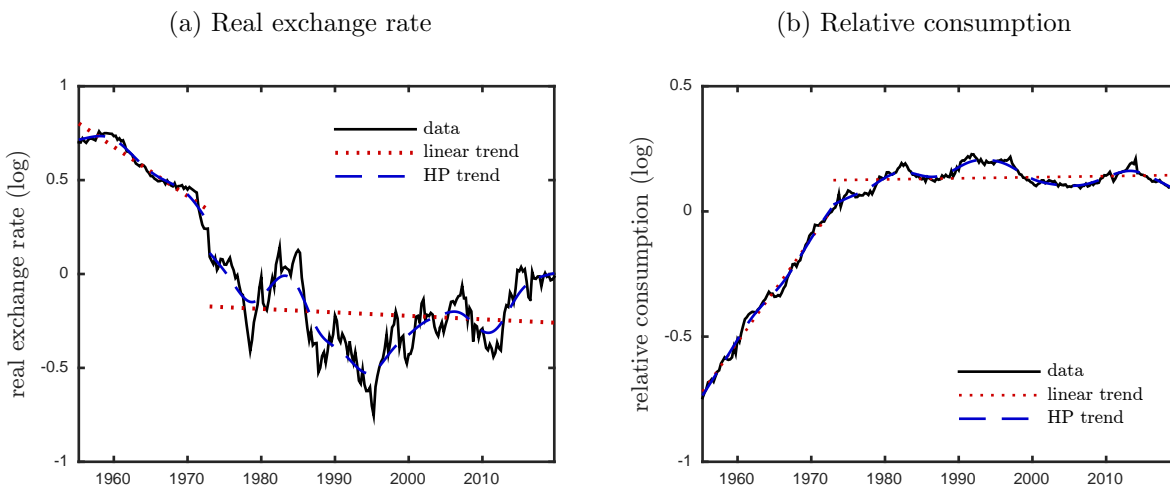
relative to that of the USA over the same period.[7] There is a clear break around 1973 in the trend of both the bilateral real exchange rate and the ratio of consumption. After that point, both values fluctuate around roughly stable levels, consistent with the presumption that the convergence of Japan to the USA had already occurred by that date. This observation is a fortunate coincidence, which we exploit to detrend each subperiod separately, as we discuss below.

Besides the raw data, Figure 1 also shows two alternative trends for both variables. The dotted red lines are log-linear trends estimated separately, one before 1973 and the other from 1973:Q1 until 2019:Q4. The dashed blue lines are Hodrick-Prescott trends constructed using a smoothing parameter of 1600. The HP-filter assigns to the trend economically relevant fluctuations that are stationary but of relatively low frequencies, which we do not want to discard, so we focus our empirical analysis on log-linear trends instead.

Therefore, we first construct a database of quarterly real exchange rates and consumption for 16 OECD countries over the period 1960:Q1 through 2019:Q4.[8] The country of reference

---

[7]Here, we show data for the USA and Japan because we use this country pair to calibrate our model below. But the pattern is similar for the other countries.

[8]The countries are Australia, Austria, Belgium, Canada, Finland, France, Germany, Italy, Japan, the Netherlands, New Zealand, Norway, Spain, Sweden, the United Kingdom, and the United States. The series for Canada and the Netherlands start in 1961:Q1 and 1960:Q2, respectively.

is the USA, so we construct 15 bilateral real exchange rates against the US dollar. We also use commodity price data at a quarterly frequency that we describe later on. We divide the sample into two subperiods: the Bretton Woods period, from 1960:Q1 (our first observation) until 1972:Q4, and the flexible exchange rate period, from 1973:Q1 through 2019:Q4. Then, we compute a log-linear trend of each variable in each subperiod and construct the detrended data by subtracting the log-linear trend from the log of the corresponding variable in each subperiod.

In Subsection 2.1, we report a number of statistics for the 15 bilateral relationships and for a series of primary commodity prices. Next, in Subsection 2.2, we provide a deeper analysis of the puzzles for the USA and Japan, the country pair that we use to calibrate our model.

2.1   The Mussa and Backus-Smith puzzles

Mussa (1986) observed that following the end of the Bretton Woods system in the early 1970s, the volatility of real exchange rates increased dramatically.

The left panel of Figure 2 displays the Mussa puzzle. The figure shows a scatter plot of the standard deviation of real exchange rates against the US dollar for 15 OECD countries. The horizontal axis displays the volatility before 1973, while the vertical axis measures the volatility afterwards. Evidently, the volatility of all bilateral real exchange rates against the US dollar increased dramatically after the demise of the Bretton Woods system, as all points are well above the 45-degree line. Moreover, although discussions of the Mussa puzzle usually center on changes in volatility, there is also a simultaneous increase in the persistence of real exchange rates, as shown in the right panel of the figure. Below, we focus the quantitative analysis on the US-Japan pair, so the corresponding moments are highlighted in both panels.

It is often claimed (e.g., by Itskhoki and Mukhin, 2019) that it is difficult to find similar breaks in the volatility of other macroeconomic variables, such as consumption, GDP, and net exports. Yet, there is a set of primary commodities such as oil, aluminum, and soybeans

Figure 2: Volatility and persistence of real exchange rates before and after 1973

(a) Volatility

(b) Persistence
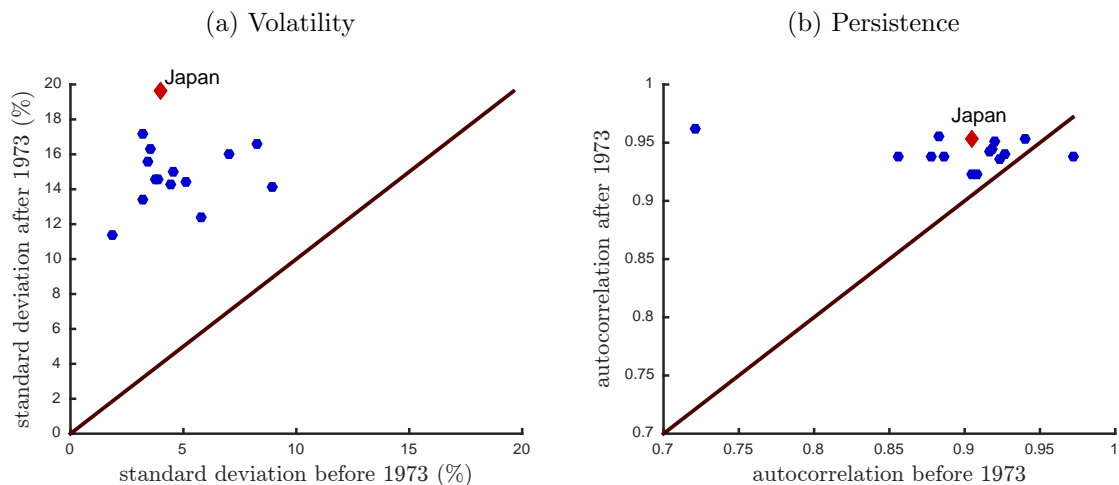
whose relative prices also show a substantial increase in volatilities after 1973.

We document this fact in Figure 3. The left panel reports the standard deviation of a set of primary commodity price indices measured in US dollars and deflated by the US CPI before and after 1973.[9] The three highlighted points in the figure represent the volatility of the commodity price indices of energy, agriculture, and metals and minerals, which are the commodity groups that we use to calibrate our model. Together, they account for the bulk of world production and trade of primary commodities.
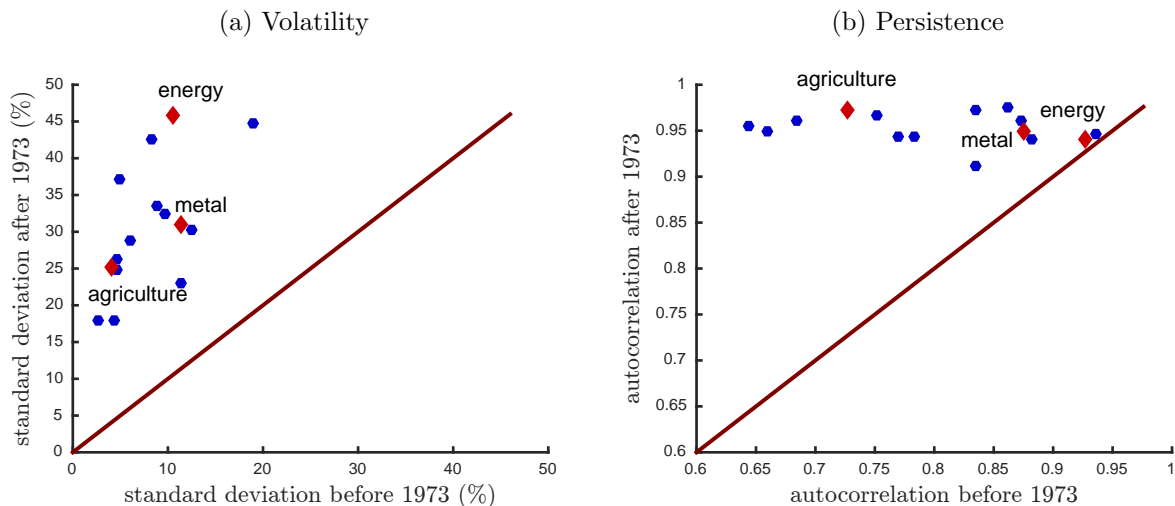
The increase in the volatility of primary commodity prices after 1973 is proportionally larger than the increase in the real exchange rate volatility documented above.[10] As was the case with the real exchange rates, relative commodity prices also display a large increase in their persistence after 1973, as the right panel of Figure 3 shows.

In sum, these figures show that the increase in the volatility and persistence of real

---

[9]The commodity price indices are from the World Bank Commodity Price Data (The Pink Sheet). We construct quarterly series from monthly data using end-of-period values. The 15 indices are energy, non-energy, agriculture, beverages, food, oils and meals, grains, other food, agricultural raw materials, timber, other raw materials, fertilizers, metals and minerals, base metals, and precious metals. The same pattern emerges if we display the standard deviations of the underlying individual commodity price series instead.

[10]The increase in the volatility of commodity prices is not an artifact of deflating nominal prices by the US CPI. If we express all commodities in deutsche marks (euros after 2000) and deflate nominal prices by the German CPI, we obtain similar results. In addition, deflating nominal commodity prices in US dollars by the price of another commodity, such as wheat, we also observe a substantial increase in volatility after 1973. We report the results in the Online Appendix.

Figure 3: Volatility and persistence of primary commodity prices (in constant USD) before and after 1973

(a) Volatility

(b) Persistence

exchange rates after 1973 is accompanied by an increase in the volatility and persistence of primary commodity relative prices.

2.1.1 Mussa meets Backus-Smith

In models with complete financial markets, international consumption risk sharing implies that the ratio of the marginal utilities of consumption between country pairs should be proportional to the bilateral real exchange rates. This property implies that relative consumption across countries should be strongly positively correlated with the real exchange rate (Backus and Smith, 1993).[11] As is well known, this observation is violated in the data.

We now show that the quantitative importance of the Backus-Smith puzzle interacts with the Mussa puzzle in that the deviation of the correlation between the real exchange rate and the ratio of consumption from one is higher in the period following 1973 than before, as Figure 4 shows. This fact has already been recognized by Colacito and Croce (2013) for the US-UK pair and by Itskhoki and Mukhin (2019) for the USA against seven

---

[11]The correlation should be one in models with time separable utility functions and constant relative risk aversion in an aggregate consumption bundle of traded and nontraded goods.

Figure 4: Correlation between real exchange rates and relative consumption

OECD countries.[12] We call this observation the Mussa meets Backus-Smith puzzle.[13]

2.2   The US-Japan stylized facts

Here we dig deeper into the properties of the bilateral real exchange rate between the USA and Japan. We choose this country pair to calibrate our model for two reasons. First, US-Japan is a country pair for which the puzzles are particularly striking, as shown in Figures 2 and 4. Second, the Ministry of Economy, Trade, and Industry of Japan publishes a bilateral Japan-US input-output table that allows us to easily calibrate key parameters of the model, as we describe below.[14]

As noted above, real exchange rates and commodity prices are highly persistent, so computing statistics using short time series leads to large statistical uncertainty. For example, the point estimate of the standard deviation of the bilateral real exchange rate between the USA and Japan after 1973 is 19.6 percentage points, but the associated small sample distribution implies that one could quite likely observe values that are 50 percent smaller or 50 percent larger than that. To deal with this issue, we compute the small sample distribution of the volatility (standard deviation) and persistence (autocorrelation) of the bilateral US-

---

[12]Canada, France, Germany, Italy, Japan, Spain, and the UK.

[13]USA-Germany is the only country pair for which this observation does not hold.

[14]The input-output table is available at https://www.meti.go.jp.

Japan real exchange rate, the distribution of the Backus-Smith correlation for the USA and Japan, and the distribution of the volatility and persistence of US-Japan relative consumptions. We also compute the distribution of the volatility and persistence of the three primary commodity price indices that we use to calibrate our model below (energy, agriculture, and metals and minerals).

We estimate these small sample distributions using a bootstrapping procedure with the following statistical model. We divide our sample before and after 1973:Q1 and detrend the data using the log-linear trend discussed above. The first subsample has 52 quarters of observations and the second has 188. For each subsample, we run a VAR of order 1 using the following five variables: i) a price index of energy normalized by the US CPI, ii) a price index of agriculture normalized by the US CPI, iii) a price index of metals and minerals normalized by the US CPI, iv) the bilateral US-Japan real exchange rate, and v) the relative per-capita consumption of the USA and Japan.[15] With the estimated parameters, we construct artificial time series of the five variables in the VAR, with a length of 55 for the first subperiod and 188 for the second, by drawing with replacement from the fitted residuals of the VAR. For each artificial time series, we compute the statistics of interest, such as the standard deviation of the bilateral real exchange rate. We repeat the procedure and construct 5,000 artificial time series for each of the five variables in the VAR, from which we estimate the small sample distribution of the statistics.

The top panel of Figure 5 displays the estimated small sample distribution of the standard deviation of the US-Japan real exchange rate before (in blue) and after (in orange) 1973. In this type of figure, the claim that the real exchange rate became more volatile is represented by a shift to the right of the estimated distributions. As noted above, the point estimate in the second subperiod is 19.6, but the distribution of the statistic has a wide range, from about 10 to 30 percentage points. Moreover, the distribution is right-skewed, with the mode (of about 15 percentage points) smaller than the mean (of about 17 percentage points), which, in

---

[15]We choose the lag length of the VAR using the Schwarz information criterion. Increasing the number of lags does not lead to significant differences.

Figure 5: Small sample distributions of the standard deviation and autocorrelation of the US-Japan real exchange rate

(a) Standard deviation of the US-Japan real exchange rate

(b) Autocorrelation of the US-Japan real exchange rate
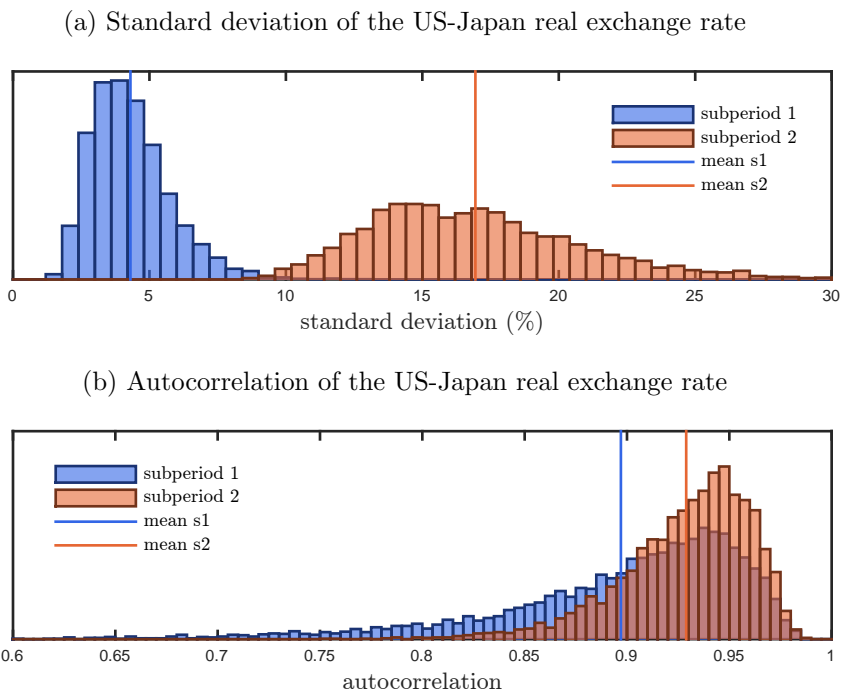
turn, is smaller than the point estimate of 19.6 percentage points. For the first subperiod, the small sample distribution of the volatility of the real exchange rate is more concentrated at lower values of the standard deviation and is less skewed than the distribution of the second subperiod. The point estimate of 4 is close to both the mean and mode of the distribution.

The bottom panel of the figure shows the estimated small sample distributions of the coefficient of autocorrelation of the bilateral real exchange rate before and after 1973. While there is a large area of overlap between the two distributions, that for the first subperiod has more mass in lower values of the autocorrelation coefficient, consistent with the observation of the previous subsection that the persistence of real exchange rates may have increased after 1973.

The moral of Figure 5 is that it could be misleading to evaluate the performance of the model using just the point estimates. Therefore, to evaluate the model, we will compare the small sample distribution of the statistic estimated from the data, as discussed above,

13

Figure 6: Small sample distribution of the correlation between the US-Japan real exchange rate and relative consumption

with the analogous distribution generated by simulating the model 5,000 times, drawing histories of artificial time series of 55 and 188 observations for the first and second subperiods, respectively, and computing the equivalent statistics that we computed using the actual data.

As for the Backus-Smith puzzle, Figure 6 shows that for both subperiods, the distributions of the correlation between the US-Japan real exchange rate and the ratio of consumption per capita have wide supports. Yet, the Mussa meets Backus-Smith puzzle is evident in that most of the mass of the distribution for the second subperiod is concentrated in negative values of the correlation, while that for the first subperiod is concentrated in positive values. But the dispersion is huge: while the point estimate of the correlation for the second subperiod is -0.47, it could be as low as -0.8 or as high as 0.3. In any case, these values are far away from one, which constitutes the Backus-Smith puzzle. On the other hand, the point estimate of the correlation in the first subperiod is 0.38, but there is a non-trivial mass of the distribution with values above 0.8, in which case the correlation could hardly be called a puzzle. In sum, the figure shows that there is a lot of statistical uncertainty in the estimate of the Backus-Smith correlation, even larger than that of the volatility of the real exchange rate, but also that the observation that the correlation was less of a puzzle during the Bretton Woods period still holds.

Finally, Figure 7 shows the estimated small sample distributions of the standard deviation and autocorrelation of the three relative commodity prices before and after 1973. In the case of the volatility, the distributions are more concentrated and take smaller values in the

14

Figure 7: Small sample distribution of the standard deviation and autocorrelation of commodity prices

(a) Energy: standard deviation

(b) Energy: autocorrelation

(c) Agriculture: standard deviation

(d) Agriculture: autocorrelation

(e) Metals and minerals: standard deviation

(f) Metals and minerals: autocorrelation

first subperiod than in the second one, similar to what we observe with the small sample distributions of the volatility of the real exchange rate. Regarding the autocorrelation, the higher persistence of commodity prices after 1973 is reflected by the fact that the small sample distributions have more mass at higher values of the autocorrelation coefficient after 1973 than before.

3   The model

We study a simple dynamic general equilibrium model featuring three large countries. There are three sectors of production in each country: nontradable final goods, tradable intermediate goods, and tradable primary commodities. Countries 1 and 2 represent the USA and Japan, respectively, while country 3 is the rest of the world.

There are two types of shocks in the model. First, there are standard productivity shocks in countries 1 and 2. For simplicity, we assume an aggregate productivity shock that affects all sectors in each country.[16] In addition, there are shocks to the endowments of primary commodities in the rest of the world, which are designed to generate volatile and persistent primary commodity prices. We could also introduce demand shocks for commodities in the rest of the world. We do so in a simplified version of the model below and show that they are equivalent ways to generate volatile commodity prices, since all that matters are shocks to the excess demand from the rest of the world. Thus, for parsimony, the calibrated model has only shocks to the endowment of commodities.

Time is discrete and denoted by $t = 0, 1, 2, ...$ Households in each country consume a single nontradable final good and can internationally trade a single non-contingent bond that pays in units of primary commodity 3.[17] Households in country $i = 1, 2, 3$ make their consumption-savings decision in order to maximize the expected lifetime utility

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \delta^t \frac{(C_t^i)^{1-\gamma}}{1-\gamma} \right],$$

where $\mathbb{E}_0 [\cdot]$ is the expectation operator conditional on information at time $t = 0$, $\delta \in (0, 1)$ is the discount factor, and $\gamma > 0$ is the risk aversion parameter.

All technologies in the model feature constant elasticities of substitution, and all goods

---

[16] These shocks are known to generate too little volatility in real exchange rates. Thus, allowing for different shocks across sectors has little hope of being of any quantitative relevance.

[17] Since we log-linearize the model, the unit in which the non-contingent bond pays is irrelevant.

and inputs are traded in competitive markets. Country $i = 1, 2, 3$ produces a final non-tradable good, C_t^i, using labor and three intermediate tradable goods, each produced in a different country.[18] The technology to produce the intermediate good, denoted by Q_t^i, uses labor and three tradable primary commodities. In addition, countries 1 and 2 are able to produce the three primary commodities using labor and a commodity-specific fixed endowment of natural resources, while country 3 receives stochastic endowments of the three primary commodities. 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, $x_{3,t}^1$ is the demand for commodity 3 by country 1 at time t, and so on.

We assume Cobb-Douglas technologies—unit elasticity of substitution—for the final goods and intermediate goods to reduce the number of parameters in the model. With Cobb-Douglas technologies, we can calibrate all parameters directly from observed input-output tables. However, we do allow for arbitrary elasticities of substitution in the production of the commodities, since there is ample micro evidence that it is substantially lower than one.

Thus, the production functions are given by

$$C_t^i = Z_t^i \left(q_{1,t}^i\right)^{\alpha_1^i} \left(q_{2,t}^i\right)^{\alpha_2^i} \left(q_{3,t}^i\right)^{\alpha_3^i} \left(n_{c,t}^i\right)^{\alpha_4^i},$$

$$Q_t^i = Z_t^i \left(x_{1,t}^i\right)^{\beta_1^i} \left(x_{2,t}^i\right)^{\beta_2^i} \left(x_{3,t}^i\right)^{\beta_3^i} \left(n_{q,t}^i\right)^{\beta_4^i},$$

$$X_{j,t}^i = Z_t^i \left[\left(1 - \phi_j^i\right) \left(e_{j,t}^i\right)^{\frac{\sigma_{x_j}^i - 1}{\sigma_{x_j}^i}} + \phi_j^i \left(n_{x_j,t}^i\right)^{\frac{\sigma_{x_j}^i - 1}{\sigma_{x_j}^i}} \right]^{\frac{\sigma_{x_j}^i}{\sigma_{x_j}^i - 1}},$$

where Z_t^i denotes aggregate productivity (TFP) in country i, which is common across sectors; $q_{1,t}^i$, $q_{2,t}^i$, and $q_{3,t}^i$ are the inputs of intermediate goods used to produce final goods; $x_{1,t}^i$, $x_{2,t}^i$, and $x_{3,t}^i$ are the inputs of primary commodities used to produce the intermediate good; $n_{c,t}^i$, $n_{q,t}^i$, and $n_{x_j,t}^i$ are the labor inputs allocated to each sector; $X_{j,t}^i$ denotes the production of commodity j; and $e_{j,t}^i$ is a commodity-specific fixed endowment of natural resources. The

---

[18]The final good is nontradable and can be used only for consumption, so we simply denote it by C.

parameters α_j^i and β_j^i are the factor shares, which sum to one in each sector, and $\phi_j^i \in (0, 1)$. The parameter $\sigma_{x_j}^i$ denotes the elasticity of substitution between inputs.

Our assumption that the rest of the world does not produce the commodities but rather receives endowments is equivalent to assuming that $\phi_j^3 = 0$ for $j = 1, 2, 3$.[19]

Uncertainty is then represented by the aggregate productivity shocks in countries 1 and 2 and the stochastic endowments of primary commodities in country 3. We assume the following (stationary) autoregressive processes:

$$\ln\left(Z_t^1\right) = (1 - \rho^{z_1})\ln\left(Z^1\right) + \rho^{z_1}\ln\left(Z_{t-1}^1\right) + \varepsilon_t^{z_1},$$

$$\ln\left(Z_t^2\right) = (1 - \rho^{z_2})\ln\left(Z^2\right) + \rho^{z_2}\ln\left(Z_{t-1}^2\right) + \varepsilon_t^{z_2},$$

$$\ln\left(X_{1,t}^3\right) = (1 - \rho^{x_1^3})\ln\left(X_1^3\right) + \rho^{x_1^3}\ln\left(X_{1,t-1}^3\right) + \varepsilon_t^{x_1},$$

$$\ln\left(X_{2,t}^3\right) = (1 - \rho^{x_2^3})\ln\left(X_2^3\right) + \rho^{x_2^3}\ln\left(X_{2,t-1}^3\right) + \varepsilon_t^{x_2},$$

$$\ln\left(X_{3,t}^3\right) = (1 - \rho^{x_3^3})\ln\left(X_3^3\right) + \rho^{x_3^3}\ln\left(X_{3,t-1}^3\right) + \varepsilon_t^{x_3},$$

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 zero mean and arbitrary covariance matrix. Variables without time subscripts represent long-run means.

Markets are competitive, and given prices, firms solve a static profit maximization problem in every period. For the calibration, we assume that in steady state, the three countries have a zero net asset position. To simulate the model, we impose a small quadratic adjustment cost to changes in the stock of the single non-contingent bond, as is customary in the literature, to avoid the inherent unit root when using linearization techniques. As the model and the solution technique are totally standard, we leave the complete description of the equilibrium and the details of the computation to the Online Appendix.

Before showing the quantitative results, we discuss the closed-form solution of a simplified version of the model. This version is useful for understanding the transmission mechanisms of the different shocks that we study in the dynamic quantitative model later on.

---

[19]This assumption simplifies the calibration, since it allows direct control of the supply of the commodities in the rest of the world.

3.1 A simplified economy

We make three simplifying assumptions to the general model. First, we assume financial autarky, so all countries must satisfy a zero trade balance condition. In this case, the solution of the model becomes static and preferences are irrelevant. So, in what follows, we do not make explicit the time dependence of the variables. Second, we simplify the production structure and eliminate the intermediate goods. And third, we consider only two commodities. Country 1 produces commodity 1, and country 2 produces commodity 2. The rest of the world receives endowments of the two commodities. To simplify notation, we use X to denote commodity 1 and M to denote commodity 2. In addition, all prices are measured in an international numeraire that is common across countries.[20]

We derive explicit expressions for the real exchange rate and for the ratio of consumption between the two countries to theoretically address the Mussa and Backus-Smith puzzles. Moreover, for this example, we allow for different productivity shocks across sectors to clarify why they have little chance of accounting for the puzzles quantitatively.

3.1.1 Country 1

The technology to produce the final good in country 1 is given by

$$ C^1 = Z^1 \left(x^1 \right)^{\alpha_x} \left(m^1 \right)^{\alpha_m} \left(l^1 \right)^{\alpha_l} , $$

where Z^1 is productivity, C^1 is the final good, and x^1, m^1 and l^1 are the inputs of commodity X, commodity M, and labor allocated to the production of the final good. The technology is constant returns to scale, so $\alpha_x + \alpha_m + \alpha_l = 1$.

---

[20]Defining national currencies and nominal exchange rates is redundant.

The cost minimization conditions imply

$$P^x x^1 = \frac{\alpha_x}{\alpha_l} W^1 l^1, \tag{1}$$

$$P^m m^1 = \frac{\alpha_m}{\alpha_l} W^1 l^1, \tag{2}$$

where P^x and P^m are the prices of the two commodities, and W^1 is the wage.

Moreover, since markets are competitive, the price level in country 1 satisfies

$$P^1 = \frac{1}{Z^1} \left(\frac{W^1}{\alpha_l}\right)^{\alpha_l} \left(\frac{P^x}{\alpha_x}\right)^{\alpha_x} \left(\frac{P^m}{\alpha_m}\right)^{\alpha_m}.$$

Zero trade balance means that

$$P^x(X - x) = P^m m^1,$$

where X is the total amount of the commodity produced in country 1. Combining this expression with equations (1) and (2) gives

$$P^x X = \frac{\alpha_x + \alpha_m}{\alpha_l} W^1 l^1. \tag{3}$$

We use this equation to replace W^1 in the expression for the price level and obtain

$$P^1 = k^1 \frac{1}{Z^1} \left(\frac{X}{l^1}\right)^{\alpha_l} (P^x)^{\alpha_x + \alpha_l} (P^m)^{\alpha_m}, \tag{4}$$

where k^1 is a constant. Equation (4) makes clear that the price of final consumption in country 1 is homogeneous of degree one in commodity prices. The reason is that since country 1 produces commodity X, there is a direct link between the domestic wage and P^x, described in equation (3). This is the reason why the share of P^x in the price level is given by its direct effect, represented by the parameter α_x, plus an indirect effect, given by α_l.

In addition, trade balance combined with equations (1) and (2) implies

$$x^1 = X \frac{\alpha_x}{\alpha_x + \alpha_m} \quad \text{and} \quad m^1 = X \frac{\alpha_m}{\alpha_x + \alpha_m} \frac{P^x}{P^m}. \tag{5}$$

The technology to produce commodity X is given by

$$X = Z^x \left[(1 - \phi_1) \left(e^1 \right)^{\frac{\sigma^1 - 1}{\sigma^1}} + \phi_1 \left(n^1 \right)^{\frac{\sigma^1 - 1}{\sigma^1}} \right]^{\frac{\sigma^1}{\sigma^1 - 1}},$$

where X is total production; Z^x is a stochastic productivity parameter; e^1 is the endowment of a natural resource, which we assume constant over time; and n^1 is labor allocated to the production of commodities.

The optimality condition with respect to labor is

$$W^1 \left(n^1 \right)^{\frac{1}{\sigma^1}} = \phi_1 P^x \left(Z^x \right)^{\frac{\sigma^1 - 1}{\sigma^1}} \left(X \right)^{\frac{1}{\sigma^1}}. \tag{6}$$

If we use equation (3) to replace the wage W^1 above, we obtain

$$\frac{\alpha_l}{\alpha_x + \alpha_m} \left(\frac{X}{Z^x n^1} \right)^{\frac{\sigma^1 - 1}{\sigma^1}} = \phi_1 \frac{l^1}{n^1},$$

which, using the production function of the commodity, can be written as

$$\frac{\alpha_l}{\alpha_x + \alpha_m} \left[(1 - \phi_1) \left(\frac{e^1}{n^1} \right)^{\frac{\sigma^1 - 1}{\sigma^1}} + \phi_1 \right] = \phi_1 \frac{l^1}{n^1}. \tag{7}$$

Equation (7) and the equilibrium condition in the labor market,

$$n^1 + l^1 = N^1, \tag{8}$$

determine the allocation of labor, where N^1 is the endowment of labor in country 1.

Since the endowment of natural resources is constant, it follows that the labor allocation

is constant over time. We have thus proved the following proposition.

Proposition 1: *The allocation of labor in country 1 depends only on production function parameters and the local endowments. Therefore, it is invariant to productivity shocks and shocks to the endowment of commodities in the rest of the world.*

Corollary 1: *The equilibrium production of the commodity X is proportional to the productivity shock Z^x.*

The proof of the corollary follows from inspection of the production function for X.

3.1.2 Country 2

Country 2 is similar to country 1, except that it produces commodity M rather than commodity X. The notation for both prices and quantities is defined as in the case of country 1. The technology to produce the final nontraded good is given by

$$C^2 = Z^2 \left(x^2\right)^{\beta_x} \left(m^2\right)^{\beta_m} \left(l^2\right)^{\beta_n}.$$

The optimality conditions are similar to the ones obtained in country 1, and the price level is given by

$$P^2 = \frac{1}{Z^2} \left(\frac{W^2}{\beta_l}\right)^{\beta_l} \left(\frac{P^x}{\beta_x}\right)^{\beta_x} \left(\frac{P^m}{\beta_m}\right)^{\beta_m}.$$

Using the same logic that we applied to country 1, we can write the price level as

$$P^2 = k^2 \frac{1}{Z^2} \left(\frac{M}{l^2}\right)^{\beta_l} \left(P^x\right)^{\beta_x} \left(P^m\right)^{\beta_m + \beta_l}, \tag{9}$$

where k^2 is a constant.

Notice that for country 2, the price that inherits the share of labor is P^m, while for country 1, it was P^x. The reason for this difference is that country 2 specializes in the production of commodity M, so its domestic wage will be connected to P^m by an equation similar to (3). As we show below, this is the main feature that relates the primary commodity

prices to the real exchange rate.

The technology to produce commodity M is given by

$$M = Z^m \left[(1 - \phi_2) \left(e^2 \right)^{\frac{\sigma^2 - 1}{\sigma^2}} + \phi_2 \left(n^2 \right)^{\frac{\sigma^2 - 1}{\sigma^2}} \right]^{\frac{\sigma^2}{\sigma^2 - 1}}.$$

Following the same steps we used for country 1, we obtain

$$m^2 = M \frac{\beta_m}{\beta_x + \beta_m} \quad \text{and} \quad x^2 = M \frac{\beta_x}{\beta_x + \beta_m} \frac{P^m}{P^x}. \tag{10}$$

As the countries are symmetric, we have the following results.

Proposition 2: *The allocation of labor in country 2 depends only on production function parameters and the local endowments. Therefore, it is invariant to local productivity shocks and shocks to the endowment of commodities in the rest of the world.*

Corollary 2: *The equilibrium production of the commodity M is proportional to the productivity shock Z^m.*

Before describing the rest of the world (Country 3), we use the optimality conditions for these two countries to relate both the real exchange rate and the ratio of consumption of the two countries to the relative prices of the commodities. Those formulas help develop intuition for the transmission mechanism that we exploit in the quantitative dynamic model of the next section.

3.1.3 Real exchange rate and commodity prices

The real exchange rate is the ratio of the price levels in both countries, so

$$\frac{P^1}{P^2} = \frac{k^1}{k^2} \frac{Z^2}{Z^1} \frac{\left(\frac{X}{l^1} \right)^{\alpha_l} (P^x)^{\alpha_x + \alpha_l} (P^m)^{\alpha_m}}{\left(\frac{M}{l^2} \right)^{\beta_l} (P^x)^{\beta_x} (P^m)^{\beta_m + \beta_l}}.$$

Using the propositions and two corollaries, we can write the real exchange rate as

$$\frac{P^1}{P^2} = K^{RER}\frac{Z^2}{Z^1}\frac{(Z^x)^{\alpha_l}}{(Z^m)^{\beta_l}}\left(\frac{P^x}{P^m}\right)^{1-\alpha_m-\beta_x}. \tag{11}$$

where K^{RER} is a constant that depends on parameters, the endowments e^1, e^2, and the (constant) allocation of labor in the two countries.

Equation (11) shows the connection between the real exchange rate and the primary commodity prices.[21] In actual economies, the exponents α_m and β_x are small, since the value added of the commodity sectors is a small fraction of GDP. In this case, the exponent $(1 - \alpha_m - \beta_x)$ is close to 1. Therefore, equation (11) implies that relative commodity prices move almost one-for-one with the real exchange rate.

The key mechanism driving this result is the connection between primary commodity prices and domestic wages. Changes in world commodity prices affect the cost of final non-traded goods not only directly, through their own coefficient in the production function, but also indirectly, through the impact that commodity prices have on wages in the commodity sector, since labor is used to produce the commodities. To the extent that countries produce different commodities, their domestic wages respond differently to a shock in a particular commodity. The differential direct effect of the commodity prices on the costs of final goods and the differential indirect effect that they have on domestic wages become the source of the real exchange rate fluctuations.

In the model, other than the endowments used in commodity production, labor is the only input. In a more general model with more diverse sets of inputs (different types of labor or of human and physical capital), the key for the mechanism to operate is that commodity production and the final good production compete in the markets for these common inputs. In addition, in this crude example, the specialization of each country in the production of each

---

[21]In more general models, in which the labor allocation is not invariant to changes in commodity prices, the terms k^1 and k^2 would not be constant, since they depend on the labor intensity used in the production of commodities. And these ratios would depend on the commodity prices, making their relationship with real exchange rates subtler.

commodity generates the close to one-for-one relationship between the real exchange rate and commodity prices. It is reasonable to expect that in more complex models, this relationship is not necessarily close to one-for-one. In fact, this simple model would spectacularly fail to match the data, since real exchange rates are substantially less volatile than primary commodity prices, as we show in the next section.

As an example, in the special case of $\alpha_m = \beta_x = 0.5$, the exponent on the relative price of commodities in expression (11) is zero. In this case, changes in relative commodity prices would be uncorrelated with the real exchange rate. The reason is that fluctuations in commodity prices have exactly the same proportional impact on the final good prices of both countries. As a result, the real exchange rate is unaffected by shocks that move only primary commodity prices.

3.1.4 Relative consumption and commodity prices

In equilibrium, the ratio of consumption is given by

$$\frac{C^2}{C^1} = \frac{Z^2 (x^2)^{\beta_x} (m^2)^{\beta_m} (l^2)^{\beta_l}}{Z^1 (x^1)^{\alpha_x} (m^1)^{\alpha_m} (l^1)^{\alpha_l}}.$$

Noting that the allocation of labor is constant and using the optimal solutions for the inputs summarized by equations (5) and (10) yields

$$\frac{C^2}{C^1} = \frac{Z^2 \left(M \frac{\beta_x}{\beta_x + \beta_m} \frac{P^m}{P^x} \right)^{\beta_x} \left(M \frac{\beta_m}{\beta_x + \beta_m} \right)^{\beta_m} (l^2)^{\beta_l}}{Z^1 \left(X \frac{\alpha_x}{\alpha_x + \alpha_m} \right)^{\alpha_x} \left(X \frac{\alpha_m}{\alpha_x + \alpha_m} \frac{P^x}{P^m} \right)^{\alpha_m} (l^1)^{\alpha_l}}.$$

Corollaries 1 and 2 above imply that both X and M are proportional to the productivities Z^x and Z^m, respectively. Therefore, we can write the ratio of consumption as

$$\frac{C^2}{C^1} = K^{CC} \frac{Z^2 (Z^m)^{\beta_x + \beta_m}}{Z^1 (Z^x)^{\alpha_x + \alpha_m}} \left(\frac{P^x}{P^m} \right)^{-(\alpha_m + \beta_x)}, \tag{12}$$

where K^{CC} is a constant. Thus, changes in the relative prices of commodities are related to the ratio of consumption. While standard measures of output do not depend on terms of trade shocks, for reasons spelled out in Kehoe and Ruhl (2008), consumption does. The reason is that the changes in relative prices of commodities change the quantities of primary commodities used in the production of the final good differentially across the two countries.

3.1.5 Country 3 and equilibrium prices

We now continue with the computation of the equilibrium. The production function of the final good in country 3 is similar to those in countries 1 and 2, and it is given by

$$C^3 = Z^3 \left(x^3\right)^{\gamma_x} \left(m^3\right)^{\gamma_m} \left(l^3\right)^{\gamma_l}.$$

As before, $\gamma_x + \gamma_m + \gamma_l = 1$. But contrary to the cases of countries 1 and 2, we allow for shocks to γ_x and γ_m, with γ_l adjusting so that the sum of the three shares equals 1. Assuming stochastic shares is a simple way to introduce shocks to the demand of primary commodities in the rest of the world.

The optimality conditions are given by

$$P^x x^3 = \gamma_x P^3 C^3,$$
$$P^m m^3 = \gamma_m P^3 C^3,$$
$$W^3 l^3 = \gamma_l P^3 C^3.$$

In contrast with the way we modeled countries 1 and 2, we assume that the rest of the world receives a stochastic endowment of the two commodities, which we denote by X^3 and M^3. This assumption allows us to directly interpret the endowments as supply shocks to commodities in the rest of the world.

Trade balance implies

$$P^x \left(X^3 - x^3 \right) + P^m \left(M^3 - m^3 \right) = 0.$$

Combining this expression with the first order conditions, we obtain

$$m^3 = \frac{\gamma_m}{\gamma_x + \gamma_m} \left(\frac{P^x}{P^m} X^3 + M^3 \right),$$

$$x^3 = \frac{\gamma_x}{\gamma_x + \gamma_m} \left(X^3 + \frac{P^m}{P^x} M^3 \right).$$

In equilibrium, it must be true that

$$x^1 + x^2 + x^3 = X + X^3$$

holds in every period.[22] Replacing the input demands x^j for $j = 1, 2, 3$ into the previous condition yields

$$X \frac{\alpha_x}{\alpha_x + \alpha_m} + M \frac{\beta_x}{\beta_x + \beta_m} \frac{P^m}{P^x} + \frac{\gamma_x}{\gamma_m + \gamma_x} \left(X^3 + \frac{P^m}{P^x} M^3 \right) = X + X^3.$$

From this equation, we obtain the equilibrium relative prices

$$\frac{P^x}{P^m} = \frac{M \left(\frac{\beta_x}{\beta_x + \beta_m} \right) + M^3 \left(\frac{\gamma_x}{\gamma_x + \gamma_m} \right)}{X \left(\frac{\alpha_m}{\alpha_x + \alpha_m} \right) + X^3 \left(\frac{\gamma_m}{\gamma_x + \gamma_m} \right)}. \tag{13}$$

Recall that Corollaries 1 and 2 imply that X and M are proportional to the corresponding productivity shocks in countries 1 and 2. Those supply shocks clearly affect the relative price in the standard fashion.[23] In addition, shocks to the world supply of commodity X as well as to the world demand of commodity M, driven by an increase in γ_m, decrease the relative

---

[22]Walras's law implies that the market for M will also be in equilibrium.

[23]For simplicity, we did not consider demand shocks in countries 1 and 2 but they can be incorporated by allowing for stochastic shares in the production of the final good.

price of commodity X. On the other hand, shocks to the world supply of commodity M as well as to the world demand of commodity X, driven by an increase in γ_x, make the relative price of commodity X to increase.

3.1.6 The Mussa puzzle

If we put the equilibrium relative prices (13) into the real exchange rate equation (11), use Corollaries 1 and 2, and take logs, we obtain

$$\ln\left(\frac{P^1}{P^2}\right) \propto \ln\left(\frac{Z^2}{Z^1}\right) + \ln\left(\frac{(Z^x)^{\alpha_l}}{(Z^m)^{\beta_l}}\right) + (1 - (\alpha_m + \beta_x))\ln\left(\frac{\mathrm{a_m}Z^m + \left(\frac{\gamma_x}{\gamma_m + \gamma_x}\right)M^3}{\mathrm{a_x}Z^x + \left(\frac{\gamma_m}{\gamma_m + \gamma_x}\right)X^3}\right), \quad (14)$$

where $\mathrm{a_m}$ and $\mathrm{a_x}$ are positive constants and the symbol $\propto$ means that we are ignoring an additive constant.[24]

The first term on the right hand side, $\ln(Z^2/Z^1)$, captures the direct effect of productivity shocks on the relative marginal costs to produce final goods in each country. The second term captures the impact that productivity shocks in the domestic commodity sectors translate, through a common labor market, into changes in wages and, hence, in the marginal cost to produce final goods. These effects of productivity shocks are standard in open economy models and known to generate too little volatility in real exchange rates because of their relatively low volatility in the data.

The third term is the new effect brought about by shocks to primary commodity markets. This term is multiplied by $(1 - (\alpha_m + \beta_x))$, which is close to 1 to the extent that commodities have a low share in the production of final goods in the two countries. This term captures how shocks to the world supply (M^3 and X^3) or demand (γ_m and γ_x) of commodities can affect the bilateral real exchange rate between countries 1 and 2 through changes in equilibrium relative prices. To the extent that these shocks are volatile and persistent, the real exchange rate will inherit these same properties, somewhat dampened by the term $(1 - (\alpha_m + \beta_x))$,

---

[24]Given that the labor inputs and natural resources are fixed, we can normalize variables so that the output of X and M are equal to Z^x and Z^m, respectively.

which is lower than but close to one for most economies. Notice, also, that the productivity shocks Z^x and Z^m in the commodity sectors also affect equilibrium prices and therefore appear in this last term. Yet, since these shocks are of relatively low volatility, they have a much smaller impact on equilibrium relative prices than shocks to the world supply or demand of commodities.

3.1.7 The Backus-Smith puzzle

Likewise, using the equilibrium relative prices (13) in equation (12) we obtain

$$\ln\left(\frac{C^2}{C^1}\right) \propto \ln\left(\frac{Z^2}{Z^1}\right) + \ln\left(\frac{(Z^m)^{1-\beta_l}}{(Z^x)^{1-\alpha_l}}\right) - (\alpha_m + \beta_x)\ln\left(\frac{a_m Z^m + \left(\frac{\gamma_x}{\gamma_m+\gamma_x}\right) M^3}{a_x Z^x + \left(\frac{\gamma_m}{\gamma_m+\gamma_x}\right) X^3}\right). \quad (15)$$

Note that world supply and demand shocks for commodities also affect the ratio of consumption, but these shocks are muted, since the last term is multiplied by $(\alpha_m + \beta_x)$, which is a relatively small number for the economies we are studying.

The sign and magnitude of the correlation between the (log) real exchange rate and the (log) ratio of consumption depend on the type of shock hitting the economy, as can be seen by comparing equations (14) and (15). First, productivity shocks in the final goods sectors, captured by the term $\log(Z^2/Z^1)$, impart a correlation of 1 between the real exchange rate and the ratio of consumption. This result, which parallels that in an economy with complete financial markets, is obtained in an economy with an extreme form of market incompleteness. Intuitively, a TFP shock to the final good's technology has the simultaneous effect of decreasing the domestic price level and increasing domestic consumption because of a positive wealth effect. This behavior results in a positive correlation between the real exchange rate and the ratio of consumption that equals 1 in this example. This same intuition applies to an economy with a single non-contingent bond, as the productivity shock still has a wealth effect that increases consumption and simultaneously reduces the price level.

Second, in the economically relevant case that $\alpha_m + \beta_x < 1$—so that commodity sectors

29

are relatively small—all shocks (supply or demand) to primary commodity markets in the rest of the world impart a correlation of -1 between the log real exchange rate and log ratio of consumption. Mechanically, this result follows because the term $1 - (\alpha_m + \beta_x)$ is positive in equation (14) and $-(\alpha_m + \beta_x)$ is negative in equation (15). Intuitively, a world shock that increases the relative price P^x/P^m generates a positive wealth effect in country 1, which produces commodity X, and a negative wealth effect in country 2, which imports commodity X. As a result, consumption increases in country 1 and decreases in country 2. Simultaneously, an increase in P^x/P^m increases the price of final goods in country 1 relative to those in country 2, leading to a correlation of -1 between the real exchange rate and the ratio of consumption.

The impact of productivity shocks to the commodity sectors is subtler, as there are direct and indirect effects through changes in relative prices. For simplicity, consider a two-country version of this model, with $X^3 = M^3 = 0$. In this case, and ignoring irrelevant constants, equations (14) and (15) collapse to

$$\ln \left(\frac{P^1}{P^2} \right) \propto (\beta_m - \alpha_m) \ln Z^x - (\alpha_x - \beta_x) \ln Z^m,$$

$$\ln \left(\frac{C^2}{C^1} \right) \propto 2 \left(\alpha_m + \frac{\alpha_x + \beta_x}{2} \right) \ln Z^x - 2 \left(\beta_x + \frac{\alpha_m + \beta_m}{2} \right) \ln Z^m.$$

The sign of the correlation between $\ln(P^1/P^2)$ and $\ln(C^2/C^1)$ depends on the symmetry of technologies between countries. If the share of commodities in the production function of both countries is the same, so that $\alpha_m = \beta_m$ and $\alpha_x = \beta_x$, the real exchange rate does not move with Z^x or Z^m and the correlation is zero. But if each country uses in the production of its final goods more of the commodity that it produces than that of the other country, so that $\alpha_x > \beta_x$ and $\beta_m > \alpha_m$, then shocks to Z^x and Z^m impart a correlation of -1 between the real exchange rate and the ratio of consumption.

The previous discussion shows that there are forces that affect the correlation in both directions. So the correlation will depend on the covariance matrix of the vector of shocks to

30

the economy. To the extent that shocks to primary commodity markets dominate, one would expect a negative correlation between the real exchange rate and the ratio of consumption. This stark conclusion depends critically on the assumption of financial autarky, since it is well known that with complete markets, the correlation ought to be equal to 1. In the incomplete markets, one-bond economy we consider next, the equilibrium exhibits insurance, a feature that makes the correlation move closer to 1. Yet, if productivity shocks and shocks to primary commodity markets are highly persistent, as we argue below, the ability of a single-bond economy to efficiently smooth consumption over states of nature deteriorates, and the intuition derived from this simplified model extends to the one-bond economy. On the other hand, if shocks have little persistence, a single non-contingent bond does a good job of smoothing consumption, and the economy looks more like a model with complete financial markets, as we show in the quantitative exercises below.

3.2 The role of nominal variables

The model discussed above has no frictions in the setting of prices, so only relative prices matter. As money is neutral, monetary policy need not be specified, so we do not specify it. It is trivial to show that the model is consistent with inflation targeting policies, whereby inflation rates are made roughly constant, as they have been in most countries during the last decades. In this case, almost all of the volatility in real exchange rates would come about through changes in the nominal exchange rate, as has also been documented by Mussa (1986).

It is certainly not the purpose of this paper to argue that price frictions are irrelevant in explaining the puzzles we address. In fact, Mussa (1986) also presents evidence from the endings of fixed exchange rate regimes, in which the real devaluation follows the nominal devaluation. It is this evidence that has been used more forcefully to argue for the frictions in the setting of prices. This evidence is much harder to reconcile with our model. Our purpose is to explore the role of real shocks arising from primary commodity markets in explaining the behavior of real exchange rates. To highlight that role, we adopted a flexible

prices model, in which the exchange rate regime is irrelevant.

The previous discussion highlights that our theory of real exchange rate fluctuations departs from most of the literature in that the breakdown of the Bretton Woods system plays no role. The interpretation that we endorse of the events is that the collapse of the Bretton Woods system happened to occur at roughly the same period in which primary commodity markets started operating in a very different way, for independent reasons.

There is ample evidence supporting the notion that the oil market went through a transformation that started slowly in the 1960s, took speed by 1970, and had fully taken effect by the time of the first oil price shock in 1973. A very compelling case is developed in detail in a fascinating book by Garavini (2019).

The series of events described by Garavini (2019) transformed a market that had traditionally been controlled by a cartel of a few international firms known as the majors. The oil producing countries played a very passive role. The relationship between the companies and the oil producing countries was based on long-term contracts with fixed posted prices. These posted prices were closely related to the price in the United States, which was fixed by the government.

The foundation of the OPEC by the end of the 1950s was an attempt by the oil producing countries to change the relationship with the majors. But for the OPEC countries, progress was slow, and only by the early 1970s did the manage to change the rules of the game. This occurred at a time in which global demand for oil was increasing dramatically because of the boom in total world output, as it has been thoroughly documented by Baumeister and Kilian (2016), for instance.

The oil market after 1973 became one without long-term contracts and no posted prices, with oil producing countries playing a prominent role. At the same time, swings in global demand became larger as a result of the growth miracles in Southern Europe, Southeast Asia and eventually China and India.

Other primary commodity markets went through similar transformations. For instance,

the tin market was regulated through the International Tin Agreement, signed by over 20 countries in the 1950s. The agreement was meant to regulate the tin market and avoid excessive fluctuations in prices. As described in detail in Mallory (1990), the agreement formally collapsed in 1985, after a few years of disagreements among its members.

4   Calibration

We divide the calibration of the model in two blocks. In the first block, we calibrate all parameters of preferences and technologies using steady state conditions or standard values in the literature. The second block of the calibration is concerned with the parameters that govern the evolution of the five stochastic processes in the model. We set those parameters by matching moments in the data with the equivalent moments generated by the model. Importantly, we do not use data on real exchange rates or consumption to calibrate the parameters of the model. They are used to evaluate the performance of the model.

4.1   Matching steady state conditions

To calibrate the first block of parameters, we proceed in three steps. First, we set the discount factor and coefficient of risk aversion to standard values: $\delta = 0.99$ and $\gamma = 2$.

The second step consists of calibrating the parameters of the production functions, which correspond to the factor shares and elasticities of substitution. As our benchmark case, we set all elasticities of substitution to be one (Cobb-Douglas technologies). In Section 5.1.3 we show that using lower elasticities of substitutions in the production function of commodities actually improves the quantitative performance of the model. To calibrate the share parameters, we use the 2005 Japan-US input-output table published by the Ministry of Economy, Trade, and Industry 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. We then further divide the primary commodity sectors into three

groups: energy, metals and minerals (referred to as metal), and the rest (referred to as agriculture). The exact mapping is discussed in the Online Appendix. The group of all final goods in the United States is assumed to be C^1, and the group of all intermediate goods is assumed to be Q^1. We do the same for C^2 and Q^2 in the case of Japan.

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 share parameters of the production functions described in Section 3. 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. With this information, we calibrate the parameters of the production functions in the USA, country 1, and Japan, country 2. They are reported in the first and second columns of Table 1.

Table 1: Calibration: factor shares (%)

	Country 1 (USA)	Country 2 (JPN)	Country 3 (ROW)
Final good C^i			
intermediate good q_1^i	$\alpha_1^1 = 23.5$	$\alpha_1^2 = 0.4$	$\alpha_1^3 = 2.8$
intermediate good q_2^i	$\alpha_2^1 = 0.2$	$\alpha_2^2 = 26.3$	$\alpha_2^3 = 1.3$
intermediate good q_3^i	$\alpha_3^1 = 3.3$	$\alpha_3^2 = 1.6$	$\alpha_3^3 = 34.9$
labor n_c^i	$\alpha_4^1 = 73.0$	$\alpha_4^2 = 71.7$	$\alpha_4^3 = 61.0$
Intermediate good Q^i			
primary commodity x_1^i	$\beta_1^1 = 8.6$	$\beta_1^2 = 5.4$	$\beta_1^3 = 9.9$
primary commodity x_2^i	$\beta_2^1 = 3.9$	$\beta_2^2 = 4.2$	$\beta_2^3 = 7.0$
primary commodity x_3^i	$\beta_3^1 = 3.3$	$\beta_3^2 = 5.7$	$\beta_3^3 = 5.1$
labor n_q^i	$\beta_4^1 = 84.2$	$\beta_4^2 = 84.7$	$\beta_4^3 = 78.0$
Primary commodity X_1^i			
labor $n_{x_1}^i$	$\phi_1^1 = 29.6$	$\phi_2^2 = 48.4$	
natural resource e_1^i	$1 - \phi_1^1 = 70.4$	$1 - \phi_2^2 = 51.6$	
Primary commodity X_2^i			
labor $n_{x_2}^i$	$\phi_1^1 = 34.3$	$\phi_2^2 = 20.3$	
natural resource e_2^i	$1 - \phi_1^1 = 65.7$	$1 - \phi_2^2 = 79.7$	
Primary commodity X_3^i			
labor $n_{x_3}^i$	$\phi_3^1 = 68.5$	$\phi_3^2 = 38.7$	
natural resource e_3^i	$1 - \phi_3^1 = 31.5$	$1 - \phi_3^2 = 61.3$	

We do not have a similar input-output table for the rest of the world (ROW). Besides the information regarding the transactions with the rest of the world in the Japan-US input-output table, we use data from the 10-sector database available from the Groningen Growth and Development Center, trade data from Comtrade, and nominal GDP and population data from the World Bank Development Indicators to pin down the remaining parameters.

We impose zero trade balance in the steady state, so the share of the final good sector in GDP is equal to the share of labor in the final good sector, α_4^i. Using the 10-sector database, we set α_4^3 equal to the GDP-weighted average of the share of the final good sector in the rest of the world.[25] Next, the input-output table reports the rest of the world's consumption of intermediate goods produced in the USA and Japan. We pin down α_1^3 and α_2^3 using data from the World Development Indicators to compute the GDP for the rest of the world. Finally, given that factor shares sum to one, we set the parameter α_3^3 as a residual.

Next, we move to the parameters of the production of intermediate goods in country 3. We calibrate β_4^3 together with the endowments of natural resources in commodities to match the share of primary commodities in the rest of the world's GDP together with other moments. That process is described in the third step below. The remaining shares, β_1^3, β_2^3, and β_3^3, are distributed according to their shares in primary commodity trade in 2005 from Comtrade. The resulting factor shares are presented in the third column of Table 1.

The third and final step consists of calibrating the relative size of each economy in steady state together with the composition of the primary commodity sectors in each country and the share of the primary commodity sector in the rest of the world's GDP, as mentioned above. Again, we use data from the World Development Indicators to compute the shares of the rest of the world's GDP, the 10-sector database to compute the shares of the primary commodity sector in rest of the world's GDP, and the US-Japan input-output table to compute the composition of GDP in the primary commodity sector in the USA and Japan.

We normalize aggregate productivity in each country to one and use population data in

---

[25]We compute the weights using nominal GDP in USD from the 2005 World Development Indicators.

Table 2: Composition of GDP within and across countries

	USA (Country 1)		JPN (Country 2)		ROW (Country 3)	
	Model	Data	Model	Data	Model	Data
Shares (%)						
Sectoral composition of GDP						
Final good	74	61	72	58	61	61
Intermediate good	23	36	26	39	29	29
Primary commodity	3	3	2	3	10	10
Composition of commodity sector GDP						
X1: Energy	59	39	0	0	46	45
X2: Agriculture	16	43	58	61	31	32
X3: Metals and minerals	25	18	42	39	23	23
Share of world GDP	37	27	16	10	47	63

2005 to compute the relative endowment of labor in each country. So we are left with the endowments of natural resources in countries 1 and 2, the endowments of primary commodities in country 3, and the share of labor in the production of the intermediate goods in country 3. Because of strong non-linearities, an exact match between the moments in the model and in the data is not feasible, so we set their values to achieve a good approximation by means of minimizing the distance between the model generated moments and the value observed in the data. We report the endowment values in the Online Appendix.

The corresponding GDP composition within and across countries is reported in Table 2, which also includes some non-targeted moments, such as the share of the final good sectors in countries 1 and 2. Table 2 shows that our model does a decent job in matching the data on GDP composition in the world. And, more importantly, it shows that our calibrated model does not overestimate the size of the primary commodity sector in the economies. It is as small as in the data.

4.2 Calibration of the stochastic processes

There are two types of shocks in the model. First, there are country-specific productivity shocks, $\{\ln Z_t^1, \ln Z_t^2\}$. We choose the parameters of these processes so as to match

the standard deviation and autocorrelation of output in each country and their correlation over the period 1973–2019. Since these parameters can be chosen to perfectly match those moments, we do not report them.[26] Importantly, we assume that the stochastic process for the productivity shocks is orthogonal to the one for commodity shocks and remains the same over the entire sample period, from 1960 to 2019.

Next, we calibrate the shocks to the rest of the world's supply of primary commodities. While we formally assume shocks to the endowment of commodities in the rest of the world, what really matters in the mechanism that we propose is to generate shifts in the *excess demand* of commodities in the rest of the world. Fluctuations in these shocks then drive fluctuations in equilibrium relative primary commodity prices.[27] Given the calibrated values for the productivity shocks, we choose the parameters of the stochastic process for the endowment of primary commodities in the rest of the world to match the standard deviation, first-order autocorrelation, and cross correlations of the three primary commodity price indices that we observe in the data. The target moments that we consider are the average values of the associated small sample distributions of the corresponding estimates that we constructed in Section 2.2.

More specifically, we divide our sample in two. The first subperiod goes from 1960 until 1973. We remove a linear trend from the three series of primary commodity prices, just as we did for the real exchange rates, and compute their persistence, volatility, and correlations. We then calibrate the parameters of the stochastic processes for the endowments of commodities in the rest of the world so as to minimize the distance between the moments in the model and the moments in the data. We repeat the same procedure for the second subperiod, the one that goes from 1973 until 2019.

---

[26]The calibration of the two types of shocks can be done independently because relative price shocks do not affect the proper measure of output, as Kehoe and Ruhl (2008) explain in detail. The Online Appendix provides a formal proof. In contrast, relative price shocks do affect relative consumption.

[27]Of course, productivity shocks also drive fluctuations in equilibrium commodity prices, but to a lesser extent, because of their lower volatility.

5 Results

In this section, we evaluate the quantitative performance of the model. We compare distributions of standard deviations and autocorrelations of commodity prices, the real exchange rate, and relative consumption generated by the model with those estimated using the actual data, as discussed in Subsection 2.2. To compute the distributions from the model, we run 5,000 simulations of length 52 and 188 for the first and second subperiods, respectively.

We first discuss the results regarding the targeted moments. Figures 8 and 9 show the distributions of the targeted moments of commodity prices used to calibrate the model. For example, the top panels in Figure 8 show the distribution of the standard deviation of the price of energy, P^{x_1}, in the first (blue bars) and second (orange bars) subperiods, both in the model (Figure 8a) and the data (Figure 8b). The vertical lines represent the means of each distribution, and these are the moments that we use to calibrate the stochastic processes. As the figure shows, the mean standard deviation of energy prices generated by the model almost exactly matches the mean standard deviation in the data in both subperiods. Even though we did not try to match any other moments, the distributions are quite similar. The same happens for the standard deviation of the prices of agriculture P^{x_2} and metals P^{x_3} in Figures 8c–8f.

In Figures 9a and 9b, we show that the mean autocorrelation of the price of energy generated by the model is somewhat higher than that in the data and the model-based distribution has more mass in higher values of the autocorrelation coefficient relative to that in the data. For the other two commodities, however, the entire distributions of the coefficient of autocorrelation estimated in the data and generated from the model are quite similar in both subperiods.[28]

We now present the performance of the model in terms of the non-targeted moments.

---

[28]We also target the correlation between commodity prices. We show these results in the Online Appendix.

Figure 8: Targeted moments: Standard deviation of commodity prices

(a) Price of energy P^{x_1}: model

(b) Price of energy P^{x_1}: data

(c) Price of agriculture P^{x_2}: model

(d) Price of agriculture P^{x_2}: data

(e) Price of metals and minerals P^{x_3}: model

(f) Price of metals and minerals P^{x_3}: data

Figure 10a shows that the distribution of the real exchange rate volatility generated by the model in the first subperiod is almost identical to that in the data. In the second subperiod, the mass of the distribution generated by the model moves to the right, though somewhat less than the distribution using the data. Both distributions are right-skewed, and they have a large area of overlap. Likewise, the distribution of the persistence of the real exchange rates, shown in Figure 10b, also moves to the right, as it does in the data. Although there is a large area of overlap between the two distributions, the distribution obtained from the model has more mass concentrated in higher values of the autocorrelation than in the data.

Figure 9: Targeted moments: Autocorrelation of commodity prices

(a) Price of energy P^{x_1}: model

(b) Price of energy P^{x_1}: data

(c) Price of agriculture P^{x_2}: model

(d) Price of agriculture P^{x_2}: data

(e) Price of metals and minerals P^{x_3}: model

(f) Price of metals and minerals P^{x_3}: data

As for the Backus-Smith puzzle, Figure 11 shows the distributions of the correlation between the real exchange rates and relative consumption. The model accounts for both the low correlation observed in the data (the Backus-Smith puzzle) and the drop in the correlation across subperiods (the Mussa meets Backus-Smith puzzle), although the model-implied distributions are somewhat to the right, compared with the ones generated by the data. The mass of the distribution generated by the model in the first subperiod is concentrated mostly in positive values, as in the data, and then much of the mass moves to negative values, as it does in the data. Although the distribution generated by the model has a wide support,

Figure 10: Non-targeted moment: The US-JPN real exchange rate

(a) Standard deviation of RER

(b) Autocorrelation of RER
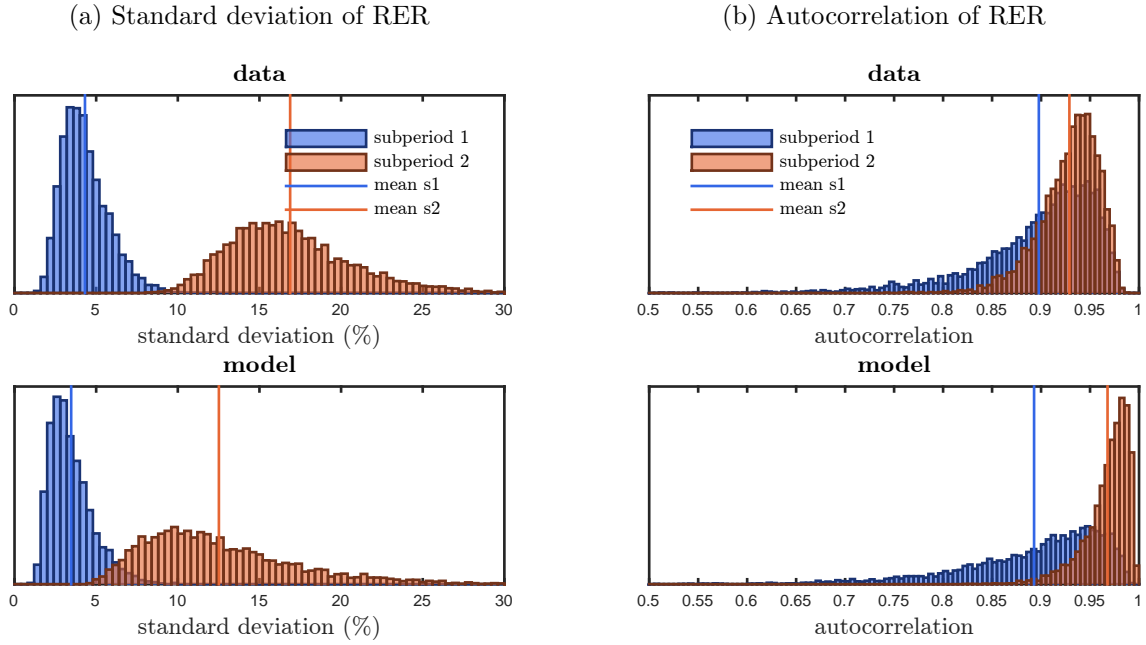
including negative and positive values, there is a large area of overlap with the distribution constructed using the actual data. Through equilibrium fluctuations in commodity prices, the mechanism that we explore can generate a wide range of values for the correlation between the real exchange rate and consumption ratios, consistent with the wide range of values observed across countries in Figure 4.

Figure 11: Correlation between the real exchange rate and relative consumption

(a) Model

(b) Data

In interpreting these results, it is important to emphasize that these are all non-targeted moments. The only thing that changes in the model between the first and second subperiods

is the parameters of the stochastic process of the shocks to the endowments of primary commodities in the rest of the world. They are chosen to match the moments of the commodity prices that we observe in the data, as Figures 8 and 9 show. No information regarding real exchange rates or consumption ratios was used in the calibration.

5.1 Robustness

The first conclusion that we draw from our experiments so far is that while the shares of primary commodities as a fraction of overall production in the USA and Japan are small, the variability and persistence of the shocks to world commodity sectors are large enough to explain some of the fluctuations that seem puzzling in the data.

In this subsection, we discuss a number of exercises to check the robustness of our results and clarify the mechanisms at work in the model. To simplify the exposition, we show results only for the volatility of the real exchange rate and for the correlation between the real exchange rate and relative consumption.[29]

In all cases, the volatility of the shocks that affect world commodity markets are recalibrated so as to always match the volatility of commodity prices observed in the data.

5.1.1 The size of the primary commodity sectors

Our first robustness exercise shows the effect of reducing the value added in the primary commodity sectors. We compare the benchmark calibration to the case in which we divide the share of primary commodities in the production of intermediate goods by 100 in each country, thereby bringing the share of commodities in the economy close to zero. This version of the model is closer to the type of models that are typically used in the literature.

Figure 12 shows the model simulation results, with the top figures representing the benchmark model discussed in the previous section, and the bottom ones the case with a lower share of commodities in the economy. The black vertical dotted lines represent the means

---

[29]The complete set of results is contained in the Online Appendix.

of the distributions in the data for the two subperiods.

Without commodities, the model is unable to account for the puzzles. In particular, the model without commodities generates low volatility of the real exchange rate and a distribution of the correlation between the real exchange rate and consumption ratios with a point mass concentrated at one. That is, the model without commodities generates a correlation between the real exchange rate and consumption ratios identical to that in complete market models. As noted in Chari, Kehoe and McGrattan (2002), having only productivity shocks and incomplete markets is not enough to explain the puzzles.[30]

Figure 12: Reducing the share of primary commodities

(a) Standard deviation of RER

(b) Correlation of RER and relative consumption
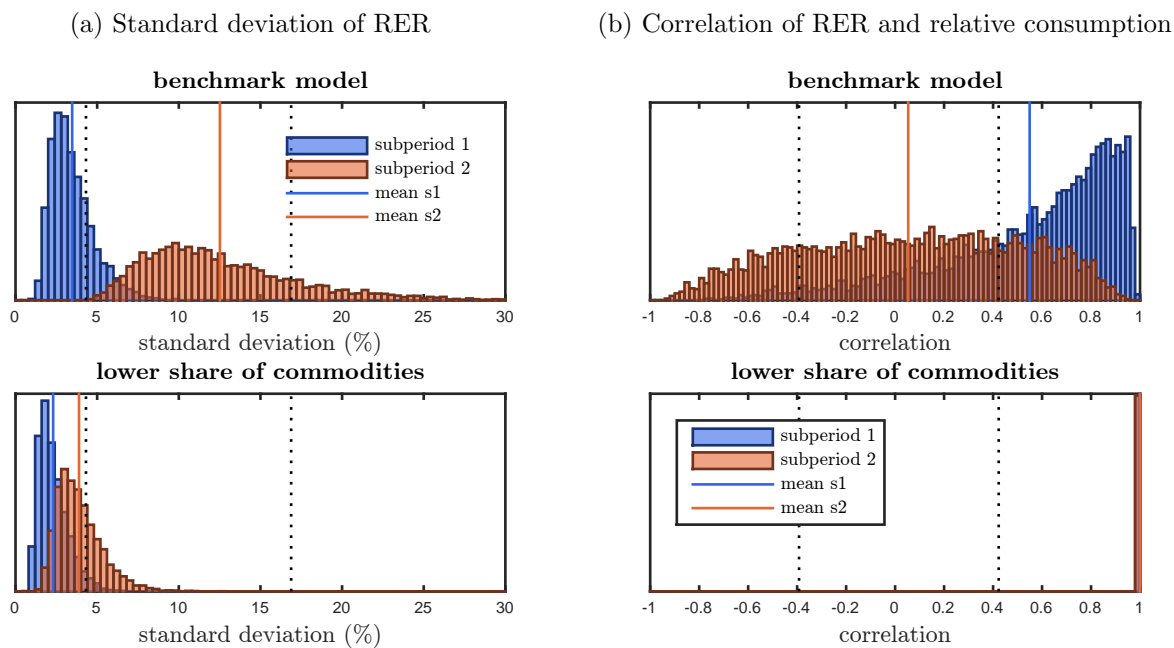
5.1.2 Heterogeneity in the commodity sectors

As emphasized in the discussion of the simplified model above, a key feature for generating fluctuations in real exchange rates and consumption is the heterogeneity in the production

---

[30]Note that Chari, Kehoe and McGrattan (2002) use the HP filter to detrend the data, whereas we remove log-linear trends.

structure between countries.[31]

The level of aggregation that we used in our model implies abstracting from many dimensions of heterogeneity contained in the data. For instance, when calibrating the size of the commodity sectors, we started with 26 different commodities contained in the input-output table and aggregated them into the three commodities of the model. This aggregation muted substantial degrees of heterogeneity that may affect the transmission mechanism of the model. To quantify this effect, let α_j^{USA} and α_j^{JPN} denote the share of commodity sector value added generated by commodity $j = 1, \ldots, 26$, for each country. We can define an index of heterogeneity by

$$H(USA, JPN) = \frac{1}{26} \sum_{j=1}^{26} \left| \alpha_j^{USA} - \alpha_j^{JPN} \right|,$$

where $|x|$ is the absolute value of x. This index has a minimum value of 0 whenever $\alpha_j^{USA} = \alpha_j^{JPN}$ for all j, so that countries are identical, and a maximum value of 1 when countries produce different sets of commodities. Using the 26 different commodity sectors in the input-output table, the value of the index is $H(USA, JPN) = 0.5$. But when we aggregate the 26 sectors into three, as we did to calibrate the model, the index falls to 0.3. This is one specific dimension in which aggregation reduces the true heterogeneity in the data.[32]

To provide a quantitative measure of the relevance of this heterogeneity, we simulated the model with the following variations. First, we choose the endowment of natural resources so that the relative sizes of the three commodity sectors are equal for Japan and the USA. In this homogeneous economy, the index $H(USA, JPN) = 0$. We also simulate a heterogeneous economy, in which we set the value added of commodities 2 and 3 in the USA to be 0, and the value added of commodities 1 and 2 in Japan to be 0. That is, the USA fully specializes in the production of energy, while Japan fully specializes in the production of agriculture. This is a case in which the countries are as heterogeneous as possible in the production of commodities, and the index $H(USA, JPN) = 1$.

---

[31]In fact, it is easy to prove that if the two countries are identical except for the productivity shocks, the real exchange rate is given by the ratio of productivities.

[32]This effect is also present when aggregating the intermediate goods.

Figure 13: Changes in the heterogeneity of the commodity sectors

(a) Standard deviation of RER   (b) Correlation of RER and relative consumption


Figure 13 displays the results of these experiments. The top panel shows the case in which the two countries are made homogenous, the middle panel the benchmark calibration, and the bottom panel the case of the extreme heterogeneity.

The figure shows that the homogeneous economy performs very poorly. The volatility of real exchange rates barely increases, and most of the mass of the correlation between the real exchange rate and the relative consumption is positive in both subperiods. At

45

the other extreme, by making the countries more heterogeneous in commodity production, the distribution of the volatility of the real exchange rate after 1973 moves further to the right, with a mean value around 22 percent, substantially higher than in the data.[33] The heterogenous case also generates an average correlation between the real exchange rate and the consumption ratios that is closer to the data both in the first and second subperiods.

The most remarkable feature of Figure 13 is that small changes to the model can have substantial effects in the standard deviation of real exchange rates and in the correlation between real exchange rates and relative consumption. We say these changes are small because they happen within a sector that represents only 3 percent of the economies we analyze.

### 5.1.3 Lower elasticity of substitution in the commodity sectors

In the baseline calibration, we assumed that the elasticity of substitution between labor and the natural resources in the commodity sectors is unity. The simple model in Section 3.1 implies that the results are invariant to changes in the elasticity of substitution. Here, we analyze to what extent this invariance result extends to the general model without complete specialization in commodity production and in which countries have access to financial markets through a risk-free bond.

Figure 14 shows the results of reducing the elasticity of substitution from 1 to 2/3 in the commodity sectors in both countries.[34] As becomes clear in the figure, with a lower elasticity of substitution, the model does a better job at matching the Mussa puzzle but barely changes the results regarding the Mussa meets Backus-Smith puzzle.

---

[33]The mean standard deviation also increases in the first subperiod, but since the volatility and persistence of primary commodity prices are much lower in that case, the difference is smaller.

[34]The algorithm had problems of convergence for lower values of the elasticity.

46

Figure 14: Lower elasticity of substitution in the commodity sectors

(a) Standard deviation of RER

(b) Correlation of RER and relative consumption
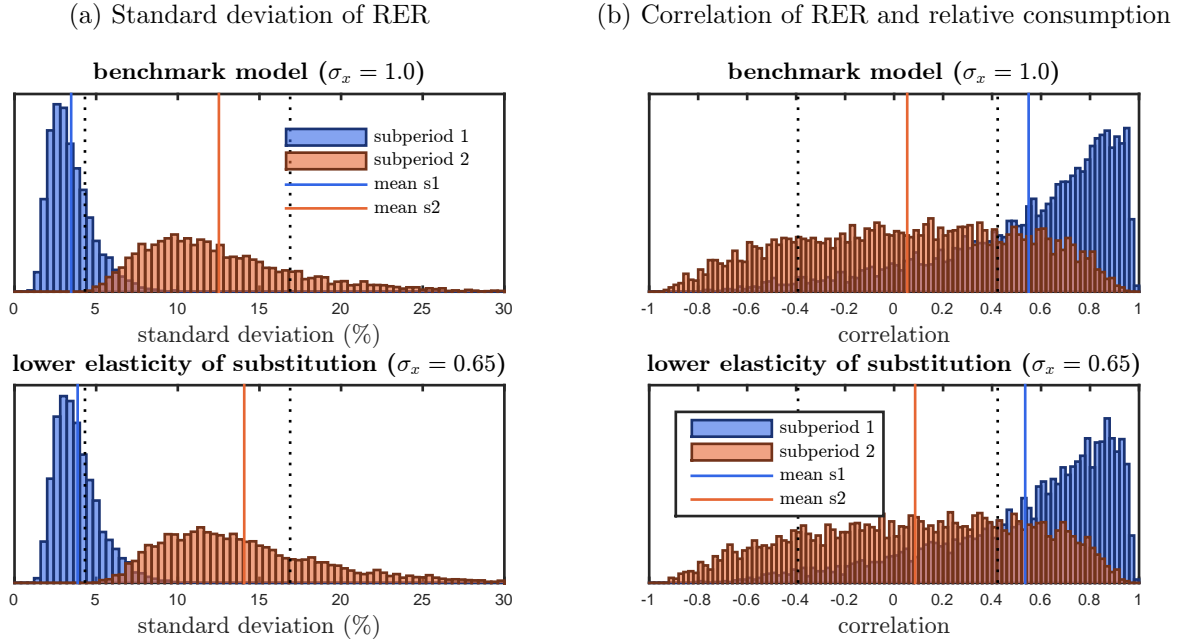

5.1.4 The structure of financial markets

In this section, we focus on the role of financial markets by considering two extremes: an economy with complete financial markets, and an economy in financial autarky. The model with financial autarky differs from the simple model discussed in Section 3.1 in that it has the more general production structure of the baseline economy.

Not surprisingly, the model with complete financial markets implies a counterfactual correlation of 1 between the real exchange rate and relative consumption, as confirmed in the upper right panel of Figure 15. But this model also fails to solve the Mussa puzzle, as the distribution of the volatility of the real exchange rate barely increases after 1973. The reason behind this result is that, relative to the baseline model, in the economy with complete financial markets, there is substantially more reallocation of labor across sectors that moderate the impact of commodity shocks on the relative prices of final goods between countries. At the other extreme, the model with financial autarky implies a larger shift to

Figure 15: The structure of financial markets

(a) Standard deviation of RER

(b) Correlation of RER and relative consumption
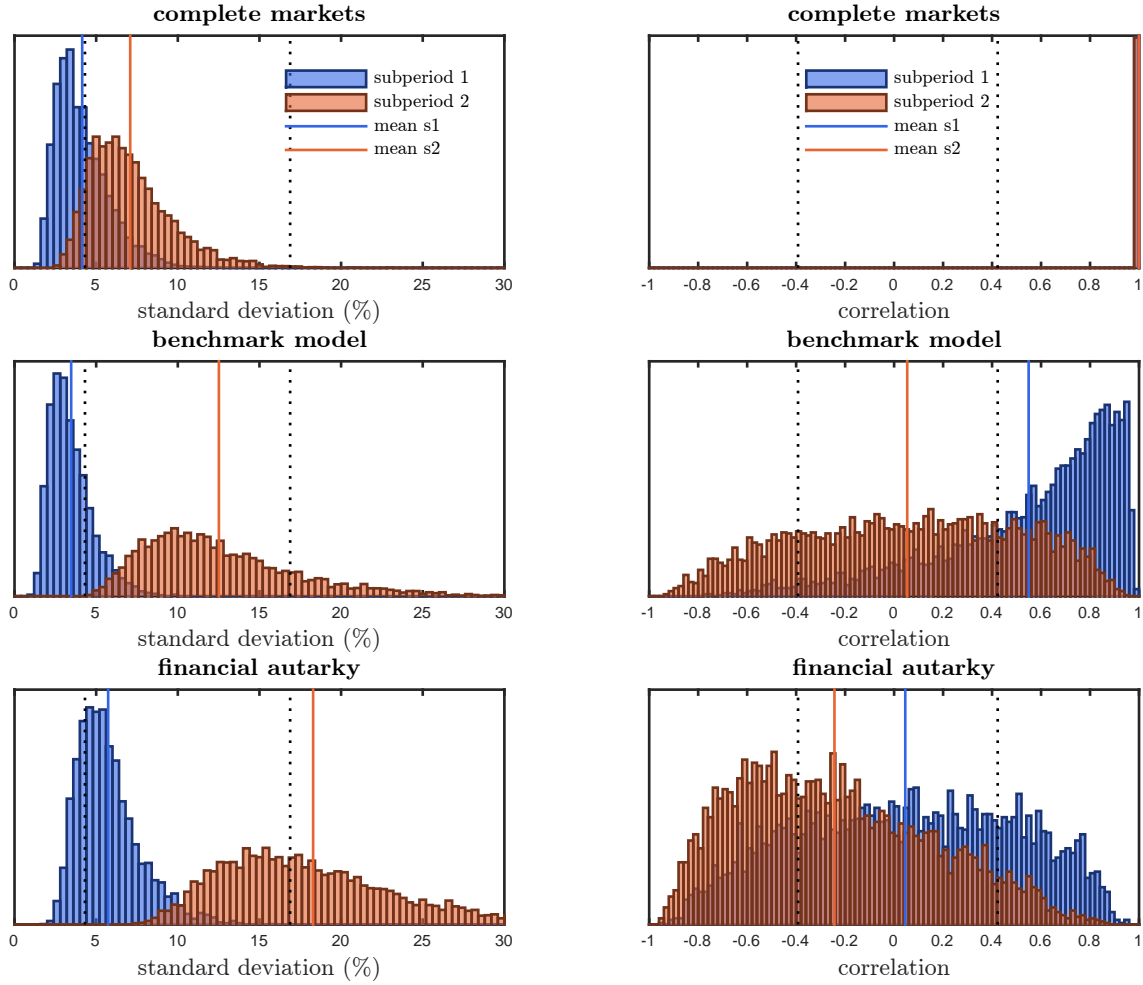
the right of the distribution of volatilities after 1973 and a more clear distinction between the distribution of the correlation between the real exchange rate and relative consumption, as shown in the bottom right panel of Figure 15. Intuitively, wealth effects in this economy are stronger and therefore also magnify the movements in relative consumption associated with equilibrium changes in commodity prices.
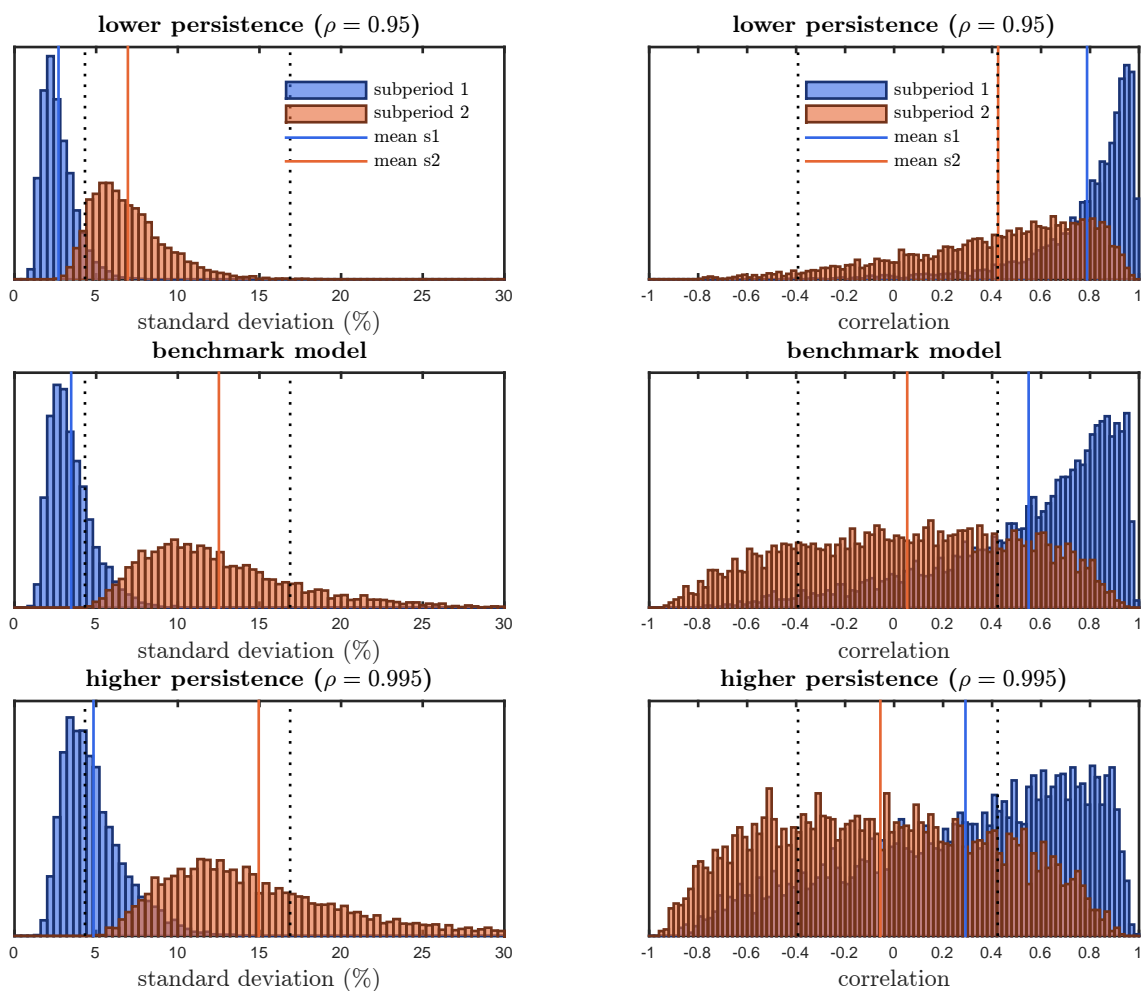
5.1.5 Persistence of shocks to world commodity markets

We now consider the role of persistence. In particular, we consider an economy in which shocks to world commodity markets are less persistent than in the baseline model and another economy in which they are more persistent. As mentioned above, we recalibrate the volatility of the innovations to the commodity shocks so as to match the volatility of commodity prices observed in the data. The results of this experiment are shown in Figure 16.

Figure 16: Changing the persistence of shocks to commodity markets

(a) Standard deviation of RER

(b) Correlation of RER and relative consumption

The model with lower persistence in world commodity shocks performs worse, both in terms of solving the Mussa puzzle and in terms of solving the Backus-Smith puzzle. On the other hand, in the model with more persistent shocks, the volatility of the real exchange rate increases more after 1973, and the distribution of the correlation between the real exchange rate and relative consumption moves slightly to the left, with more mass in negative values of the correlation.

The reason for these results is that as we keep increasing the persistence of world commodity shocks, the model more and more resembles an economy in financial autarky because the risk-free bond is a less useful hedge against persistent shocks. Conversely, as we reduce the persistence of commodity shocks, the risk-free bond does a better job of insuring against shocks, and the economy tends to look more like an economy with complete financial markets.
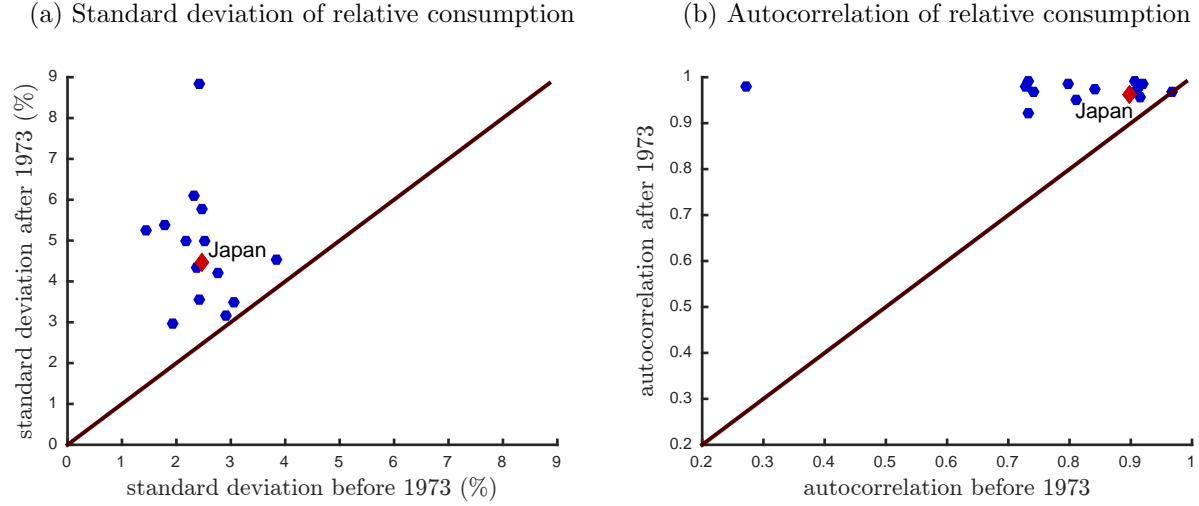
5.2   Relative consumption

Following the literature, so far we have focused the analysis on the fluctuations of the real exchange rate and its correlation with relative consumption. However, the model also makes predictions about the volatility and persistence of relative consumption before and after 1973. Therefore, in this section, we discuss those predictions and how they match features of the data. We first present evidence regarding the behavior of the ratio of consumption pre- and post-1973, as we did in Section 2. Then we show the predictions of the model.

The left panel of Figure 17 displays the volatility before and after 1973 of the log-difference of detrended consumption for the group of OECD countries analyzed in Section 2, relative to that of the USA. The right panel shows the persistence of relative consumption. As with the real exchange rates, we observe a clear increase in both the volatility and persistence of relative consumption after 1973.[35]

The same observation holds when we focus on the Japan-US pair, as shown in the upper panels of Figure 18. These figures show the bootstrapped small sample distributions of

---

[35]Portugal is the only country in which the volatility of relative consumption decreased after 1973.

Figure 17: Volatility and persistence of (log) relative consumption before and after 1973

(a) Standard deviation of relative consumption

(b) Autocorrelation of relative consumption

the standard deviation and first order autocorrelation of the (log) relative consumption of Japan against the USA. The increase in volatility and persistence is reflected as a shift of the distributions to higher values after 1973. These movements in the distributions of the volatility and persistence of relative consumption mimic those of the real exchange rate.

In the bottom panels of Figure 18, we see that the benchmark model is able to match the distributions of the volatility and persistence of relative consumption before and after 1973. We emphasize again that we did not use any data on consumption to calibrate the model, so the results are due solely to the increase in the volatility and persistence of primary commodity prices and the transmission mechanism generated by the model. In the simplified economy of Section 3.1, the increase in the volatility of shocks to world commodity markets (summarized by shocks to X^3 and M^3) translate into higher volatility of relative consumption, although it is of a smaller magnitude than that of the real exchange rate (compare equations (14) and (15), noting that $\alpha_m + \beta_x < 1 - (\alpha_m + \beta_x)$). Moreover, the persistence of relative consumption inherits the persistence of shocks to world commodity markets.

Figure 18: Volatility and persistence of relative consumption of Japan versus the USA

(a) Standard deviation of relative consumption

(b) Autocorrelation of relative consumption
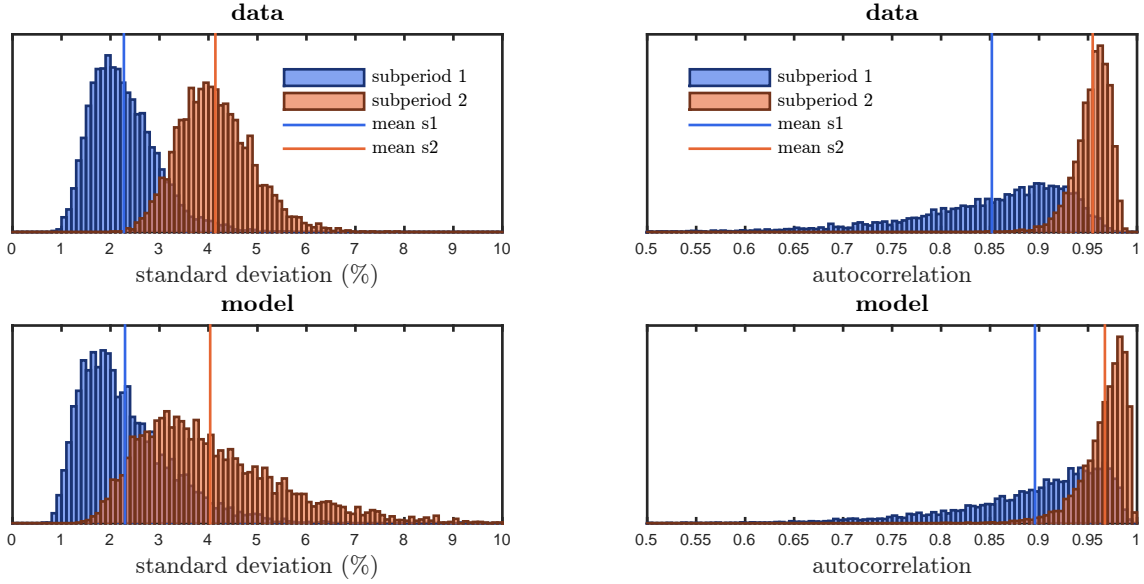
6 Conclusion

We showed that explicitly modeling primary commodities in an otherwise standard quantitative multi-country model can go a long way in explaining two puzzles in the international economics literature. This is the case even though both production and use of primary commodities account for a small fraction of the economy.

The key features of the model are incomplete markets and volatile shocks to the primary commodity markets. A calibrated model in which shocks to primary commodity markets are chosen so as to match the moments of primary commodity prices in the data, can explain most of the increase after 1973 in the volatility of real exchange rates between the USA and Japan. In addition, it also explains a sizable fraction of the low correlation between real exchange rates and the ratio of consumption. Importantly, in the calibrated model, the sizes of the primary commodity sectors are as low as in the data.

We also show, through a series of robustness exercises, that distributions of the key mo-

ments are quite sensitive to relatively small changes in parameters. Together with the high statistical uncertainty that we document, this fact strengthens the notion that documenting the puzzles with point estimates ought to be done with extreme caution and, more importantly, may be subject to substantial changes even following small changes in the production structure of countries.

For instance, the model implies that the fracking technology that changed the oil sector in the USA in the last couple of decades could have a substantial effect on the impulse response of the real exchange rate after a supply shock to the oil market.

In our theory of real exchange rate behavior, there is no role for exchange rate systems. Thus, the interpretation we adopted in the paper is that the breakdown of the Bretton Woods system in the early 1970s happened to coincide, by chance, with developments that made the primary commodity markets much more volatile.

There is an alternative interpretation, though, raised as a question by a discussant of this paper. According to that interpretation, the turbulence in primary commodity markets that became evident by the early 1970s may have contributed to the final collapse of the Bretton Woods system, which had already showed signs of weaknesses during the previous decade.

We find this a very interesting hypothesis, since it is typically the case that fixed exchange rates endogenously collapse owing to economic forces, and the end of the Bretton Woods system was not an exception. This view challenges the standard assumption in the literature, which takes the collapse of the Bretton Woods system as the exogenous change that cause the increased volatility of real exchange rates. The proper causality may be to go from changes in primary commodity markets to the collapse of the Bretton Woods system and more volatile real exchange rates. We leave this fascinating discussion for further research.

References

Ayres, Joao, Constantino Hevia, and Juan Pablo Nicolini. 2020. "Real exchange rates and primary commodity prices." *Journal of International Economics*, 122(C).

Backus, David K., and Gregor W. Smith. 1993. "Consumption and real exchange rates in dynamic economies with non-traded goods." *Journal of International Economics*, 35(3-4): 297–316.

Baumeister, Christiane, and James D. Hamilton. 2019. "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks." *American Economic Review*, 109(5): 1873–1910.

Baumeister, Christiane, and Lutz Kilian. 2016. "Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us." *Journal of Economic Perspectives*, 30(1): 139–60.

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.

Colacito, Riccardo, and Mariano M. Croce. 2013. "International Asset Pricing with Recursive Preferences." *The Journal of Finance*, 68(6): 2651–2686.

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.

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.

Garavini, Giuliano. 2019. *The Rise and Fall of OPEC in the Twentieth Century.* Oxford University Press.

Heathcote, Jonathan, and Fabrizio Perri. 2002. "Financial autarky and international business cycles." *Journal of Monetary Economics*, 49(3): 601–627.

Itskhoki, Oleg, and Dmitry Mukhin. 2019. "Mussa Puzzle Redux." 2019 Meeting Papers 1434.

Itskhoki, Oleg, and Dmitry Mukhin. 2021. "Exchange Rate Disconnect in General Equilibrium." *Journal of Political Economy*, 129(8): 2183–2232.

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.

Kollmann, Robert. 1995. "Consumption, real exchange rates and the structure of international asset markets." *Journal of International Money and Finance*, 14(2): 191–211.

Mallory, Ian A. 1990. "Conduct Unbecoming: The Collapse of the International Tin Agreement." *American University International Law Review*, 5(3): 835–892.

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.

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