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

March 5, 2021

Abstract

In this paper we show that explicitly modeling primary commodities in an otherwise totally standard incomplete markets open economy model can go a long way in explaining some of the main puzzles in the international economics literature.

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

*We thank Tim Kehoe, Andy Neumeyer, and Pedro Teles for comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the Inter-American Development Bank.

1 Introduction

In this paper we show that explicitly modeling primary commodities in an otherwise totally standard incomplete markets open economy model can go a long way in explaining some of the main puzzles in the international economics literature.

Primary commodity prices are known to be very volatile and very persistent. A large literature has developed that studies those markets and to identify and estimate the underlying shocks that make those markets particularly volatile.

The small open economy literature has long recognized the role that shocks to commodity prices have in the economy in general and on the real exchange rate in particular. This literature analyzed and identified so strongly this role in small commodity producing countries that it is acknowledged and debated in policy circles.

However, when studying 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 on total economic activity is quite small in these countries. The analysis of this paper shows that, although their share on total output may seem small, the volatility of those prices is so high that they 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 USA 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 than a model that ignores the commodity sector.[2] In this paper we go a few steps beyond, solve a truly dynamic model, and address quantitatively some of the famous puzzles.

The key ingredient of the model is the interaction of incomplete markets and shocks that move prices of primary commodities. We document how the transmission mechanism implied by this interaction generates moments that can help to partially reconcile two main puzzles.

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

[1] An earlier literature did discuss and evaluated the role of primary commodity markets in real exchange rate 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 imposed a zero trade balanced every period, so the solution of the model is really static. This property of the model allowed us to solve the model globally. Instead, in this paper we use a log-linearized version to compute the solution.

ity of real exchange rates since the breakdown of the Bretton-Woods period of fixed nominal exchange rates, where, coincidentally, nominal exchange rates also became very volatile.

The second one, known as the Backus-Smith puzzle, documents a low correlation, even negative in some cases, between the times series of the bilateral real exchange rate between any two given countries and the ratio of consumption between those same two countries—in complete markets models, this correlation ought to be equal to 1.

We also show that the same transmission mechanism generates another strong pattern in the data, that we call the Mussa meets Backus-Smith puzzle. This is the quite generalized fact that the deviation of the correlation between the real exchange rate and the ratio of consumption from 1 (the Backus-Smith puzzle) is higher in the period following the breakdown of the Bretton-Woods period when compared with the period before. In other words, the Backus-Smith puzzle becomes a quantitatively larger puzzle in the post 1973 period than in the period before. This feature of the data has received much less attention in the literature—to the best of our knowledge, we are the first ones to document this systematic behavior of the data and offer a potential explanation of it.

We study a labor-only open economy 3-country model. Each country produces a final non-traded 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 if countries have different production structures, the real exchange rate is affected by shocks that change primary commodity prices.

Since the model that we consider is real, it is unable to explain if the real exchange rate movements come about by changes in the nominal exchange rate or changes in the relative inflation rates. However, a straightforward adaptation of the model that includes a monetary authority that uses monetary policy to replicate the inflation rate observed in these countries implies that most of the movements ought to come about through changes in the nominal exchange rate, as observed originally by Mussa.

We calibrate the model for the USA and Japan because, as we show, this is the country pair for which the puzzles are quantitatively more striking and because the Ministry of Economy, Trade, and Industry of Japan publishes a bilateral Japan-USA 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 the supply of 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. Thus, we emphasize that we do not explain why the prices of primary commodities are so volatile.[3]

[3] In an ongoing paper, we are exploiting the supply and demand shocks estimated in the empirical literature

A key aspect of our calibration that we wish to highlight, is that the size of value added created in the primary commodity sectors in the model are as small as it is 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.

In a nutshell, the benchmark calibrated model can account for a large fraction of the Mussa puzzle and the Mussa meets Backus-Smith puzzle. We also show how small variations on the parametrization can deliver an even better performance.

Our contribution is to show that if primary commodity markets are exposed to shocks (like the weather, OPEC policies, armed conflict or trade disruptions for example) that are volatile enough so as to be able to account for the volatility of the prices observed in the data, the transmission mechanism in the model is enough to explain a sizable fraction of the puzzles mentioned above. We show that the model, so calibrated, generates volatile and persistent real exchange rates and correlations between the real exchange rates and the ratio of consumption substantially lower than one (the Backus-Smith puzzle), in the direction of matching the data.

We also 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. In addition, the calibrated model exhibits a correlation between the real exchange rate and the ratio of consumption much lower than one, as observed in the data. Interestingly, the model also generates a lower correlation of the real exchange rate and the ratio of consumption for the post Bretton-Woods period (the Backus-Smith meets Mussa puzzle). As we document, the data also features this property.

In describing the stylized facts, we perform two types of exercises. First, we show that the properties of the data and aforementioned puzzles are qualitatively quite general and common to most industrialized countries. For this first part of the analysis we use data for 17 OECD countries and compute standard statistics such as standard deviations and correlations. The estimated quantitative strength of the puzzles, however, does depend on the particular country pair.

We note, however, that the moments used to describe the stylized facts have a lot of statistical uncertainty due to the high persistence of both the bilateral real exchange rates and the primary commodity prices. For this reason, we then focus on the USA-Japan bilateral relation and compute the entire small sample distribution of the statistics of interest, such

that used VARs (Caldara, Cavallo and Iacoviello, 2019) or gigantic oil discoveries (Arezki, Ramey and Sheng, 2017) to identify the shocks.

as the volatility of the real exchange rate, and find that they are quite dispersed and skewed. For example, the point estimate of the standard deviation of the USA-Japan real exchange rate 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. This observation means that focusing on point estimates that could be the result of an awkward random draw from these distributions could be misleading. 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 our our model, we construct the analogous distributions obtained from artificial time series generated by the model and compare them to those estimated using the data.

The model has no frictions other than the lack of contingent asset markets. As such, it is unable to match another major puzzle, namely, the deviations from the uncovered interest rate parity—we show in the appendix that just one risk free bond, as we assume, is enough for the UIP to hold in the model. This is the reason why we see the results of our paper as complementary to efforts in the literature that explore frictions in the setting of prices (Chari, Kehoe and McGrattan (2002) as well as many other) or segmentation in asset markets (such as the recent work by Itskhoki and Mukhin (2020) and Itskhoki and Mukhin (2019)).

In view of this, we chose to keep the model as simple as possible, with the exception of the incomplete asset markets assumption. 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, were we allow for the elasticity of substitution to be smaller than one, as suggested by the literature.

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, is very standard. Section 4 discusses the calibration and Section 5 presents the results.

2 The stylized facts

In documenting the stylized facts of real exchange rates and the Mussa and Backus-Smith puzzles, we need to choose a particular way to filter the data. To extract long-run trends, most of the literature uses the Hodrick-Prescott (HP) filter or some sort of differencing of the data. As is well known, however, one problem with these filtering approaches is that they put more weight to higher frequencies relative to lower frequency fluctuations in the data that do not necessarily represent the long-run trend, as we document below. We believe that these lower frequencies are equally important in order to discipline the type of models we want to use to understand these type of phenomena. This observation is particularly important for variables with a lot of persistence such as real exchange rates and primary commodity prices.

Following the post-World War II period, most industrial countries displayed a strong convergence in income levels to that of the United States. This convergence was also accompanied by a Balassa-Samuelson effect whereby the bilateral real exchange rates of these countries against the USA appreciated over time as the economies converged. Importantly, the convergence in income levels over the post-World War II period occurred mostly during the Bretton-Woods period of fixed exchange rates. Since the mid 1970s, relative incomes and consumption tended to be roughly stable across developed countries and the bilateral real exchange rates tended to fluctuate around a relatively stable level.[4]

Figure 1: Detrending real exchange rates and Balassa-Samuelson

(a) Japan

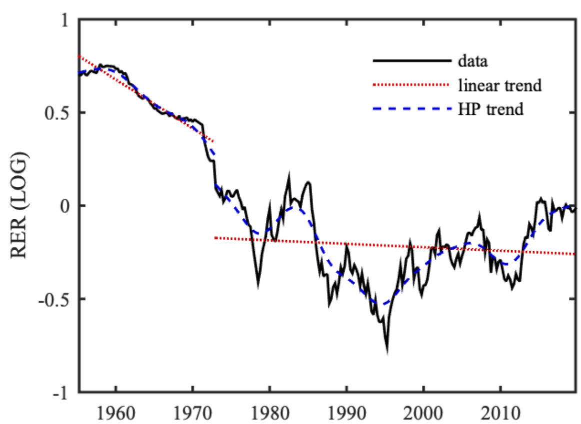

(b) Germany

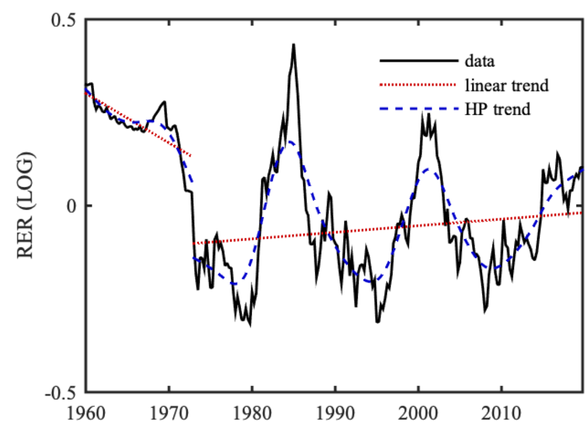

(c) United Kingdom

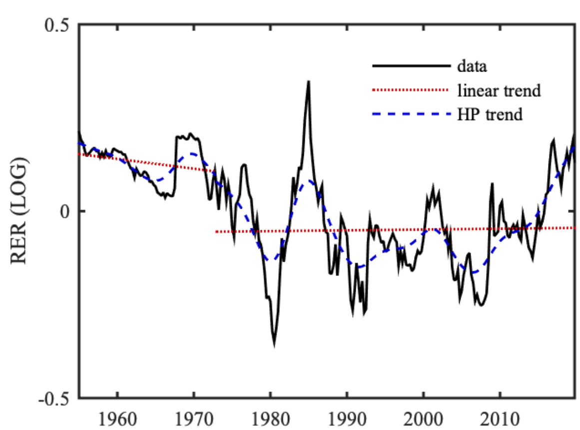

[4]In the model that we describe in Section 3, a Balassa-Samuelson effect emerges naturally from different growth rates over time in the total factor productivity of a country relative to that of the USA.

To illustrate this point, the solid lines in Figure 1 shows the log of the bilateral real exchange rates between the USA and Japan, Germany, and UK at a quarterly frequency from 1955 through 2019. Those three countries are ordered in terms of how strong was the Balassa-Samuelson effect during the Bretton-Woods era, with Japan featuring the largest appreciation of its currency, followed by Germany and then by the United Kingdom. Note that there is a clear break in the trend of the real exchange rate that roughly coincides with the end of Bretton-Woods. After the early 1970s, real exchange rates fluctuate around a roughly stable level, consistent with the presumption that convergence to the USA has already occurred by that date. This observation is a fortunate coincidence that we exploit to detrend each sub-period separately, as we discuss below.

Besides the raw data, Figure 1 also shows two broken trends for the real exchange rates estimated separately, one for the Bretton-Woods period and the other from 1973 until 2019. We split the sample in this way in order to capture the observation that the Balassa-Samuelson effect occurred mostly during the Bretton-Woods period and not afterwards. The dotted red lines are log-linear trends and the dashed lines are Hodrick-Prescott trends constructed using a smoothing parameter of 1600. Note how the HP filter assigns to the trend economically relevant fluctuations that are stationary but of relatively low frequencies (cicles of about 10 years of duration or more) that we do not want to discard in our empirical analysis.

For these reasons, we adopt the following strategy in our data analysis. First, we construct a data base for quarterly real exchange rates and consumption for 17 OECD countries over the period 1955:Q1 through 2019:Q4.[5] The country of reference is the USA, so we construct 16 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 in two sub-periods: the Bretton-Woods period, from 1955:Q1 (our first observation) until 1972:Q4, and the flexible exchange rate period, from 1973:Q1 until 2019:Q4. Then we compute a log-linear trend of each variable in each of the sub-periods and construct the detrended data by subtracting the log-linear trend from the log of each variable in each sub-period, as we did for the real exchange rates between Japan, Germany and the UK against the USA in Figure 1.

With these detrended data, we perform two types of exercises. First, in Subsection 2.1 we report point estimates for standard deviations and for correlations of variables for the 16 bilateral relationships and for a series of primary commodity prices. The purpose of this subsection is to document that the puzzles we analyze are qualitatively quite general, and do not depend on the specific country pair. On the other hand, as it will become clear, the quantitative strength of the puzzle does depend on the country pair.

[5]The countries are Australia, Austria, Belgium, Canada, Finland, France, Germany, Italy, Japan, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, United Kingdom, and United States.

Figure 2: The Mussa Puzzle

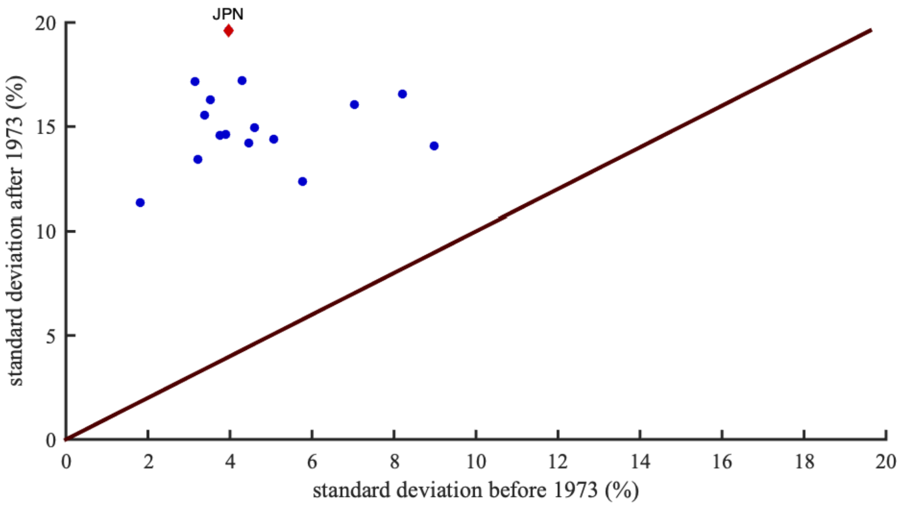

Then, in Subsection 2.2, we provide a deeper analysis of the puzzles for the Japan-USA pair. The reason to do so is that the core of the paper is the analysis of a model that is calibrated to precisely those two countries.

2.1 The Mussa and Backus-Smith puzzles

Mussa (1986) observed that, by the end of the Bretton-Woods system in the early 1970s, the volatility of both *nominal* and *real* exchange rates increased dramatically as industrial countries moved from fixed to flexible exchange rates. This high correlation between nominal and real exchange rates has been taken by many as evidence in favor of monetary non-neutrality and price stickiness.

Figure 2 displays the Mussa puzzle. The figure shows a scatter plot of the volatility (standard deviation) of the bilateral real exchange rate against the USA dollar for 16 OECD countries before and after 1973. The horizontal axis represents the volatility during the Bretton-Woods period while the vertical axis measures the volatility afterwards. Evidently, the volatility of all bilateral real exchange rates against the USA dollar increased dramatically after the demise of Bretton-Woods system, as all points are well above the 45 degree line.

In a recent paper, Itskhoki and Mukhin (2019) claim that this break in the behavior of the real (and nominal) exchange rates does not have a counterpart in other macro variables, such as consumption, GDP, and net export; or other nominal variables, such as inflation. There is, however, a set of relative prices which do show a large break in their volatility before and after 1973: the relative prices of primary commodities such as oil, gold, soybeans,

Table 1: Standard deviation of primary commodity prices normalized by USA CPI (%)

Commodity	1960-1972	1973-2019
Index - Energy	10.6	46.0
Index - Agriculture	4.0	25.2
Index - Metals and minerals	11.3	30.8
Average (36 series)	8.2	30.8

and so on, as shown in Table 1 and Figure 3.[6]

Table 1 reports the volatility of primary commodity prices measured in USA dollars and deflated by the USA CPI. The table shows the volatility of the commodity price indices of energy, agriculture, and metals and minerals, which is the set of commodities that we use to calibrate our model, as well as the average volatility of 36 individual commodity prices in our database. There is a large increase in the volatility of primary commodity prices after 1973, even proportionally larger than the increase in the real exchange rate volatility documented above. This increase the volatility of commodity prices is not an artifact of deflating commodity prices by the USA CPI. If we deflate nominal commodity prices by the price of another commodity, such as wheat, the volatility of relative commodity prices also increases substantially after 1973. Figure 3 shows the increase in the volatility of the 15 primary commodity price indices in our database after 1973.[7],[8]

2.1.1 Mussa meets Backus-Smith

In models with complete financial markets, international consumption risk sharing implies that the ratio of marginal utilities of consumption between country pairs should be proportional to their 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).[9] This observation is strongly violated in the data: the corre-

[6]The commodity price series are from the World Bank Commodity Price Data (The Pink Sheet). The series are monthly and we construct quarterly series using end-of-period values. The 15 series in Table 1 and Figure 3, including Energy, Agriculture, and Metals and Minerals, correspond to price indices of commodity groups. The pattern is similar if we use the individual price series instead.

[7]As above, the same pattern emerges if we deflate nominal commodity prices by the price of other commodity, such as wheat.

[8]A less well known fact, related to the Mussa puzzle, is that not only the volatility but also the persistence of real exchange rates increased after 1973. Moreover, as it was the case with the volatility, the persistence of the primary commodity prices also increased substantially after 1973.

[9]The correlation should be one in models in which the utility functions features constant relative risk aversion in an aggregate consumption bundle of traded and non-traded goods.

Figure 3: Standard deviation of primary commodity prices normalized by US CPI

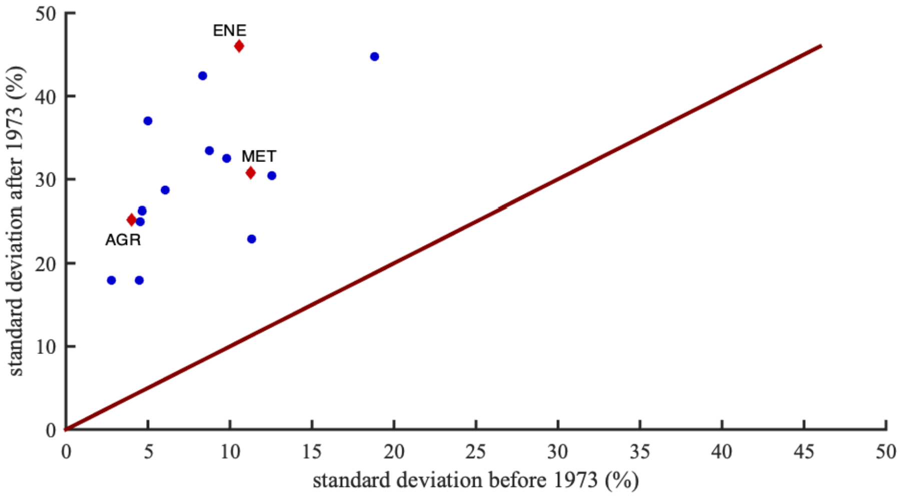

lation between relative consumptions and the bilateral real exchange rate tends to be mildly negative, although this correlation is often not statistically significant.

We note, however, 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 the breakdown of Bretton-Woods than before, as shown in Figure 4. In other words, the Backus-Smith puzzle becomes a quantitatively larger puzzle after 1973 than in the period before. We call this observation the Mussa meets the Backus-Smith puzzle.[10]

2.2 The USA-Japan stylized facts

In this subsection we dig deeper into the properties of the bilateral real exchange rate between the USA and Japan. As mentioned above, this is the pair chosen to calibrate the model in the next Section. In there, we choose USA and Japan for two reasons. First, USA-Japan is the country pair for which the puzzles are quantitatively more striking, as shown in Figures 2 and 4. And second, the Ministry of Economy, Trade, and Industry (METI) of Japan publishes a bilateral Japan-USA Input-Output table that allows us to easily calibrate key parameters of the model that we describe below.[11]

As noted above, real exchange rates and commodity prices are highly persistent. As is well known, computing statistics using highly persistent data using relatively short time series, as we have, lead to very large statistical uncertainty. For example, the point estimate

[10] USA-Germany is the only country pair for which this observation does not hold.
[11] The Input-Output table is available at https://www.meti.go.jp.

Figure 4: Correlation between relative consumption and RER for all country pairs against the USA

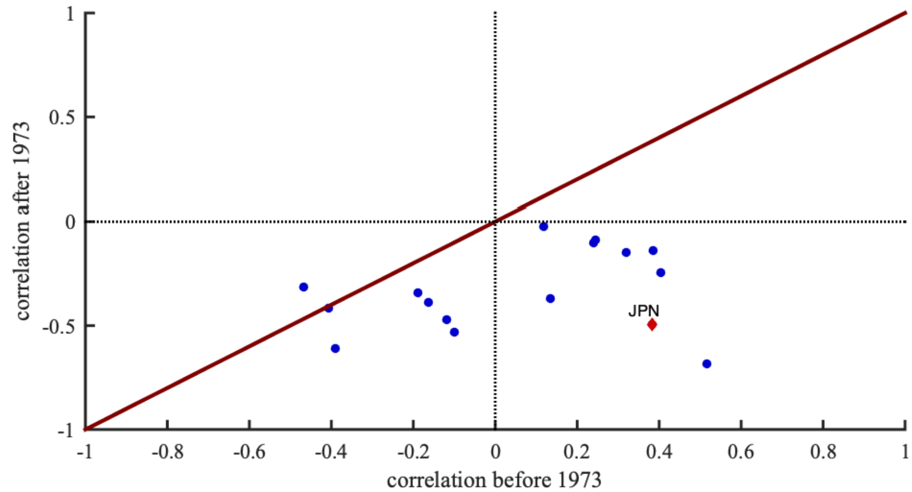

of the standard deviation of the bilateral real exchange rate between the US and Japan after 1973 is 19.6 percentage points. Is that a good estimate of the population moment or just an awkward random draw from a highly persistent data generating process? To deal with this issue, in this subsection we estimate the small sample distribution of the estimates that we are interested in, focusing our attention on the USA-Japan pair. In particular, we compute the small sample distribution of the volatility (standard deviation) of the bilateral US-Japan real exchange rate, the distribution of the Backus-Smith correlation for USA and Japan, and the distribution of the volatility of the main three primary commodity prices (energy, agriculture, metals and minerals) that we use to calibrating our model below.

We estimate the aforementioned small sample distributions using a bootstrapping procedure with the following statistical model. We divide our sample before and since 1973:Q1. The first subsample has 52 quarters of observations and the second, 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 USA CPI, ii) a price index of agriculture normalized by the USA CPI, iii) a price index of metal and minerals normalized by the USA CPI, iv) the bilateral US-Japan real exchange rate, and v) the relative per-capita consumption of US and Japan.[12] With the estimated parameters, we construct artificial time series of the five variables in the VAR, of length 55 for the first sub-period and 188 for the second sub-period, by drawing with replacement the fitted residuals of the estimated VARs. For each artificial time series, we compute the statistics of interest, such as the standard deviation of the bilateral real

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

Figure 5: Estimated distributions of the standard deviation of the USA-Japan real exchange rate

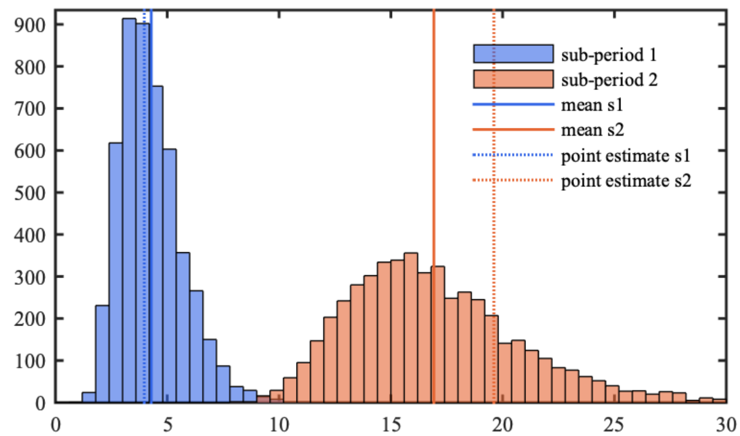

exchange rate. We repeat the procedure and construct 5000 artificial time series for each of the five variables in the VAR from which we estimate the small sample distribution of the statistics of interest.

To illustrate the importance of this procedure, Figure 5 displays the estimated small sample distribution of the standard deviation of the USA-Japan real exchange rate before (in blue) and after (in orange) 1973. As noted above, the point estimate of the second sub-period is 19.6 (shown as a vertical dotted line in the plot) 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 turn, is smaller than the point estimate of 19.6 percentage points. For the first sub-period, the small sample distribution of the volatility of the real exchange rate is more concentrated and less skewed than that of the second sub-period. The point estimate of 4.0 is close to both the mean and the mode of the distribution.

The moral of Figure 5 is that it could be misleading to evaluate the performance of the model using just the point estimate of 19.6. Therefore, to evaluate the model we will compare the small sample distribution of the statistic estimated from the data, as discussed above, with the analogous distribution generated by simulating the model 5000 times, drawing histories of artificial time series of 55 and 188 observations for the first and second sub-periods, 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 sub-periods, the distributions of the correlation between the USA-Japan real exchange rate and the ratio of consumptions per capita have wide supports. Yet, the Mussa meets Backus-Smith puzzle

Figure 6: Estimated distributions of the Backus-Smith correlation for USA and Japan

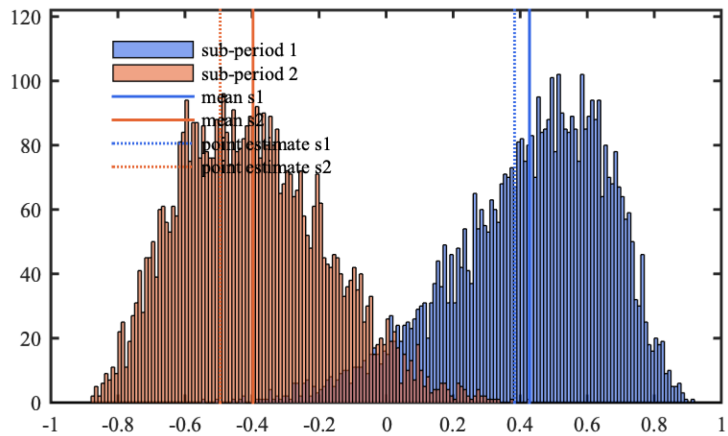

is evident in that most of the mass of the distribution for the second sub-period is mostly concentrated in negative values of the correlation while that for the first sub-period is concentrated in positive values. But the dispersion is huge: while the point estimate of the correlation for the second sub-period is -0.47 (dotted orange line) it could be as low as -0.8 or as high as 0.2. 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 sub-period is 0.38 (dotted blue line) 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 that the observation that the correlation was less of a puzzle during the Bretton-Woods period still holds.

Finally, Figure 7 shows the estimated distributions of the volatility of the three relative commodity prices before and after 1973. In all cases, the distributions are more concentrated and take smaller values in the first sub-period than in the second one, similar to we observe with the distributions of the volatility of the real exchange rate.

3 The Model

We study a simple general equilibrium model featuring two large countries and a third one that is interpreted as the rest of the world. In each country there are three sectors of production: nontradable final goods, tradable intermediate goods, and tradable primary commodities. The third economy (the rest of the world) is used as a device to generate

Figure 7: Estimated distributions of the standard deviation of commodity prices

(a) Energy

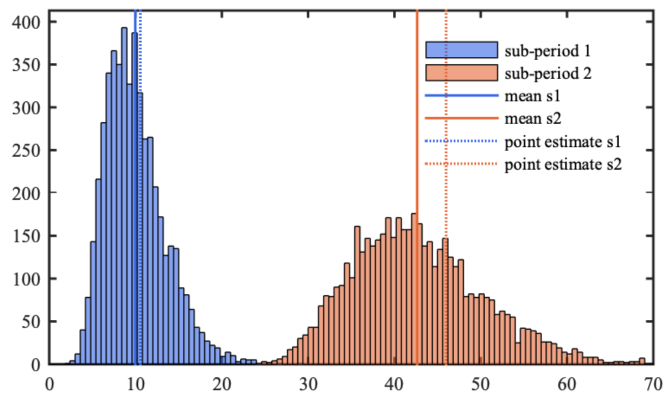

(b) Agriculture

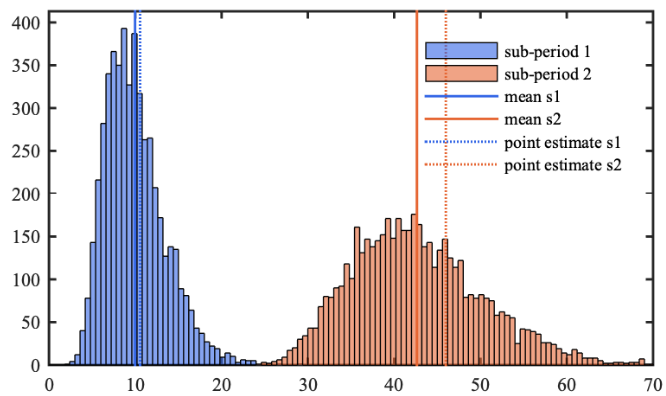

(c) Metals and minerals

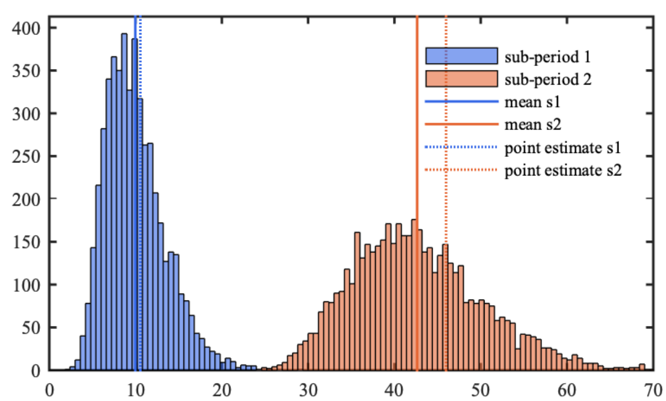

excess demands for primary commodities that can be shocked to generate commodity price fluctuations.

There are two types of shocks to the economy. First, there are standard productivity shocks in the two large economies. These shocks are know to generate too little volatility in real exchange rates, so, for simplicity, we assume a single productivity shock in each country, that affects all sectors in that country in the same way. On top of these two shocks, there are shocks to the endowments of primary commodities in the rest of the world, which are designed to generate volatile primary commodity prices.

Time is discrete and denoted by $t = 0, 1, 2, ...$ Households in each country consume a single nontradable final good. Households in countries 1 and 2 inelastically supply their endowments of labor and commodity-specific fixed endowments of natural resources, whereas households in country 3 (the rest of the world) inelastically supply their fixed endowment of labor and stochastic endowments of primary commodities. In addition, all households trade a single non-contingent bond that pays in units of primary commodity 3. Households in country $i = 1, 2, 3$ make their consumption-savings decision in order to maximize the expected lifetime utility

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

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

All technologies in the model are CES and all goods and inputs are traded in competitive markets. Countries produce a final nontradable good, Y_t^i, using labor and three intermediate tradable goods, each produced in one country, Q_t^i. The technology to produce the intermediate good 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, whereas country 3 receives stochastic endowments of all 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. The production functions are given by

$$Y_t^i = Z_t^i \left[\alpha_1^i \left(q_{1,t}^i \right)^{\frac{\sigma_y^i - 1}{\sigma_y^i}} + \alpha_2^i \left(q_{2,t}^i \right)^{\frac{\sigma_y^i - 1}{\sigma_y^i}} + \alpha_3^i \left(q_{3,t}^i \right)^{\frac{\sigma_y^i - 1}{\sigma_y^i}} + \alpha_4^i \left(n_{y,t}^i \right)^{\frac{\sigma_y^i - 1}{\sigma_y^i}} \right]^{\frac{\sigma_y^i}{\sigma_y^i - 1}},$$

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

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

where Z_t^i denotes total factor productivity (TFP) in country i, 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; $x_{j,t}^i$ for $j = 1, 2, 3$ are the inputs of primary commodities used to produce the intermediate good; $n_{y,t}^i$, $n_{q,t}^i$, 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 the commodity-specific fixed endowment of natural resources. The parameters α^i,β^i, $\phi_j^i > 0$ are the factor shares and sum to one in each sector. The parameters σ_y^i, σ_q^i, and $\sigma_{x_j}^i$ denote the elasticities of substitution between the inputs and vary from zero (perfect complements) to positive infinity (perfect substitutes).

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

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

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

Markets are competitive and, given prices, firms solve a static profit maximization problem in every period. As we now show, these static conditions are enough to obtain an equilibrium relationship between the real exchange rate and the primary commodity prices.

Note that the bilateral real exchange rate is defined as $\left(P_t^{y_1} S_t^{1,2}\right) / P_t^{y_2}$, where $P_t^{y_i}$ is the final good price index in country $i = 1, 2$, and $S_t^{1,2}$ transforms units of account in country 1 into units of account in country 2. We let ξ_t be the logarithm of the real exchange rate between countries 1 and 2, which is then given by

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

where $x = \ln(X)$.

To obtain a closed form relationship that is linear in logs, we consider the particular case in which all elasticities of substitution are equal to one, so the production functions are Cobb-Douglas in the inputs. Perfect competition implies that prices equal marginal costs in all markets. Since all technologies are Cobb-Douglas, all marginal costs are Cobb-Douglas functions of their input prices as well. Thus, for example, for country 1, we obtain

$$P_t^{y_1} = \frac{1}{Z_t^1}\left(\frac{P^{q_t^1}}{\alpha_1^1}\right)^{\alpha_1^1}\left(\frac{P^{q_t^2}}{\alpha_2^1}\right)^{\alpha_2^1}\left(\frac{P^{q_t^3}}{\alpha_3^1}\right)^{\alpha_3^1}\left(\frac{W_t^1}{\alpha_4^1}\right)^{\alpha_4^1}$$

In a similar fashion, intermediate good prices $P^{q_t^1}$ and $P^{q_t^2}$ are also Cobb-Douglas functions of commodity prices and wages. But if we use the cost minimization conditions of the commodity-producing sectors, wages can be written as exponential functions of the commodity prices and the prices of the endowments. Using the law of one price for the commodities, we can write the term $(p_t^{y_1} + s_t^{12})$ as a linear function of the logarithms of the commodity prices and the prices of the endowments, all measured in units of account in country 2.[13]

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

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

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

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

$$
\begin{aligned}
\delta_{x_1} &= \left(\alpha_1^1-\alpha_1^2\right)\beta_1^1+\left(\alpha_2^1-\alpha_2^2\right)\beta_1^2+\left(\alpha_3^1-\alpha_3^2\right)\beta_1^3+\frac{\left(\alpha_1^1-\alpha_1^2\right)\beta_4^1+\alpha_4^1}{\phi_1^1},\\
\delta_{x_2} &= \left(\alpha_1^1-\alpha_1^2\right)\beta_2^1+\left(\alpha_2^1-\alpha_2^2\right)\beta_2^2+\left(\alpha_3^1-\alpha_3^2\right)\beta_2^3+\frac{\left(\alpha_2^1-\alpha_2^2\right)\beta_4^2-\alpha_4^2}{\phi_2^2},\\
\delta_{x_3} &= \left(\alpha_1^1-\alpha_1^2\right)\beta_3^1+\left(\alpha_2^1-\alpha_2^2\right)\beta_3^2+\left(\alpha_3^1-\alpha_3^2\right)\beta_3^3.
\end{aligned}
$$

As we made clear before, in equilibrium, all the prices on the right-hand side and the real exchange rate move together, so the coefficients just described do not measure the correlation between any given commodity price and the real exchange rate. But this equation makes explicit the connection between the real exchange rate and the primary commodity prices $p_t^{x_1}, p_t^{x_2}$, and $p_t^{x_3}$.

This relationship between the traded components and the real exchange rate depends on the asymmetries across countries. Notice that if the countries are the same, in the sense that $\alpha_1^1=\alpha_1^2$, $\alpha_2^1=\alpha_2^2$, and $\alpha_3^1=\alpha_3^2$ several terms in those coefficients do become zero. As we show below when we calibrate the model to the United States and Japan, we indeed find large asymmetries in factor shares between the two countries.

As we assumed all shocks to be stationary, all prices will also be stationary, so there is no reason to detrend the variables in the model. Primary commodity prices in the data are very persistent, but they do fluctuate around a relatively stable value, without obvious evidence of a trend. However, in exploring the ratio of output or consumption of Japan and the USA, the countries we will focus on, an obvious trend is visible during the fast convergence of output per capita of Japan to values close to the ones in the USA. This obvious trend in the data in output per capita imparts an obvious trend in the real exchange rate, as Figure 1 shows. This effect can be seen in equation (1) in the first two terms in the right hand side, that represent the total factor productivities in the two countries. Clearly, in the period in which productivity in Japan was converging fats to the value in the USA, a trend in the real exchange rate will become evident, exactly as described in Figure 1. As the period of convergence of Japan's productivity to the value of the USA is roughly consistent with the duration of the Bretton-Woods period, the piece-wise linear detrending we use in the paper seems consistent with the assumptions of the model.

Equation (1) is one of several possible representations of the real exchange rates that must hold in equilibrium. Other representations, which must also hold in equilibrium, can be derived by substituting different equilibrium conditions into the definition of the real exchange rate. Note also that, except for the aggregate productivity shocks z_t^1 and z_t^2, all the other variables in equation (1) are endogenous and, therefore, correlated with the productivity shocks and the other variables. The purpose of discussing this particular condition that must hold in an equilibrium is to provide intuition of how the input-output table does affect the transmission mechanism of shocks to the world supply of primary commodities into real exchange rates between large economies.

The relationship between the real exchange rate and the primary commodity prices just discussed, can be obtained using the static profit maximization problems of the firms in both countries. To compute an equilibrium, as it is standard in open economy models, we use the fact the the trade balance is determined by the savings decision of households in each country and the equilibrium condition that the sum of the trade balances of the three countries ought to be zero. For the calibration, we assume a zero net asset position for the three countries in steady-state, and we assume a small quadratic adjustment cost to the stock of the uncontingent international asset (the only asset available) to compute the transitions, as is customary in the literature.[14]

4 Calibration

To calibrate our model economy, we proceed in three steps. First, we specify the preference parameters β, the discount factor, and γ, the coefficient of risk aversion. These are standard parameters in the literature and we set $\beta = 0.99$ 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 unity (Cobb-Douglas). For the share parameters, we use the 2005 Japan-US Input-Output Table published by the Ministry of Economy, Trade, and Industry (METI) of Japan. We map each of the 174 sectors in the input-output table into the three sectors considered in our model: final goods, intermediate goods, and primary commodities. 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 presented in the Online Appendix. The group of all final goods in the United States is assumed to be Y^1, and the group of all its intermediate goods is assumed to be Q^1. We do the same for Y^2 and Q^2 in the case of Japan.

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

[14]We use the method of undetermined coefficients described in Uhlig (2003) to simulate the model.

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 function in the USA, country 1, and Japan, country 2. They are reported in the first and second columns of Table 2.

Table 2: Calibration: factor shares (%)

	Country 1 (USA)	**Country 2 (JPN)**	**Country 3 (ROW)**
Final good			
intermediate good Q_1	$\alpha_1^1 = 23.5$	$\alpha_1^2 = 0.4$	$\alpha_1^3 = 2.8$
intermediate good Q_2	$\alpha_2^1 = 0.2$	$\alpha_2^2 = 26.3$	$\alpha_2^3 = 1.3$
intermediate good Q_3	$\alpha_3^1 = 3.3$	$\alpha_3^2 = 1.6$	$\alpha_3^3 = 34.9$
labor n^y	$\alpha_4^1 = 73.0$	$\alpha_4^2 = 71.7$	$\alpha_4^3 = 61.0$
Intermediate good			
primary commodity X_1	$\beta_1^1 = 8.6$	$\beta_1^2 = 5.4$	$\beta_1^3 = 9.9$
primary commodity X_2	$\beta_2^1 = 3.9$	$\beta_2^2 = 4.2$	$\beta_2^3 = 7.0$
primary commodity X_3	$\beta_3^1 = 3.3$	$\beta_3^2 = 5.7$	$\beta_3^3 = 5.1$
labor n^q	$\beta_4^1 = 84.2$	$\beta_4^2 = 84.7$	$\beta_4^3 = 78.0$
Primary commodity X_1			
labor $n_i^{x_i}$	$\phi_1^1 = 29.6$	$\phi_2^2 = 48.4$	
natural resource e_i^i	$1 - \phi_1^1 = 70.4$	$1 - \phi_2^2 = 51.6$	
Primary commodity X_2			
labor $n_i^{x_i}$	$\phi_1^1 = 34.3$	$\phi_2^2 = 20.3$	
natural resource e_i^i	$1 - \phi_1^1 = 65.7$	$1 - \phi_2^2 = 79.7$	
Primary commodity X_3			
labor $n_i^{x_3}$	$\phi_3^1 = 68.5$	$\phi_3^2 = 38.7$	
natural resource e_3^i	$1 - \phi_3^1 = 31.5$	$1 - \phi_3^2 = 61.3$	

We do not have a similar dataset for the input-output table of country 3, the rest of the world (ROW). So 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 equilibrium, 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.[15] Next, the Input-Output Table reports how much of the intermediate goods from US and Japan the rest of the world consumes. Computing the rest of the world GDP using data from the World Development Indicators, we can pin down the factor shares α_1^3 and α_2^3. Finally, given that the factor shares sum to one, the parameter α_3^3 can be computed 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 and commodities to match the share of primary commodities in the ROW GDP together with other moments. That 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 2.

The third and final step consists of calibrating the relative size of each economy in steady state together with the composition of their primary commodity sectors and the share of the primary commodity sector in country 3 GDP as mentioned above. Again, we use data from the World Development Indicator to compute the shares of world GDP, the 10-sector Database to compute the shares of the primary commodity sector in ROW GDP, and the US-Japan Input-Output Table to compute the composition of GDP in the primary commodity sector in the US and Japan.

Table 3: Calibration: endowment distribution

USA (Country 1)	**JPN (Country 2)**	**ROW (Country 3)**
$N_1 = 4.5$	$N_2 = 2.0$	$N_3 = 93.5$
$e_{11} = 1.00$	$e_{21} = 0.01$	$X_{31} = 1.08$
$e_{12} = 0.24$	$e_{22} = 0.14$	$X_{32} = 0.60$
$e_{13} = 1.80$	$e_{23} = 0.13$	$X_{33} = 0.36$

We normalize the aggregate TFP in each country to be equal to one and we use population data in 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. An exact match between the moments in the model and in the data might not be feasible, so we simply set their values to achieve a good approximation. We report the endowment values in Table 3.

The corresponding GDP composition within and across countries are reported in Table

[15] We use nominal GDP in US dollars from the World Development Indicators in 2005 to compute the weights.

Table 4: Moments: GDP composition 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	73.9	61.1	72.1	57.9	61.0	61.0
Intermediate good	22.6	35.6	25.9	38.6	29.4	29.2
Primary commodity	3.5	3.3	2.0	3.4	9.6	9.8
Composition of commodity-sector GDP						
X1 - Energy	57.6	38.9	0.6	0.4	46.1	45.0
X2 - Agriculture	15.9	43.0	58.4	60.9	31.5	32.0
X3 - Metal	26.5	18.1	41.0	38.7	22.4	23.0
Share of world GDP	35.2	27.4	15.3	10.0	49.6	62.6

4, which also includes some non-targeted moments such as the share of the final good sectors in countries 1 and 2. Table 4 shows that our model does a good job in matching the data on GDP composition in the world. And most importantly, it shows that our calibrated model does not overestimate the size of the primary commodity sector in the economies. They are as small as in the data.

5 Results

In this section, we present the numerical performance of the model. We compare the distributions of the moments of commodity prices, real exchange rates, and relative consumption generated by the model to the ones in the data, discussed in Subsection 2.2 and illustrated in Figures 5, 6, and 7. As in the data, we compute the distributions in the model based on 5,000 simulations of 52 and 188 periods in the first and second sub-periods, respectively.

In the previous section, we showed how all structural parameters of preferences and technologies were pinned down using steady state conditions of the model. What remains to be calibrated are the parameters that describe the stochastic processes of the structural shocks. That is, the persistences, standard deviations, and correlations.

There are two types of shocks in our model. On one hand, there are productivity shocks that we allow to be country specific but not sector specific. These shocks are known to generate very little volatility to the real exchange rates, so we opted to avoid the complication of allowing for sector specific shocks. The moments that control these shocks have been chosen so as to match the volatility and persistence of output in each country in the period 1974–2019. They can be chosen so as to match exactly those moments, so they are not reported.[16]

On the other hand, we have shocks that affect the prices of primary commodities. For simplicity, we model those shocks as supply shocks in the rest of the world (country 3 in our model). To calibrate these shocks, we chose as targets the persistence and the volatility of primary commodity prices in the data as well as their correlations. As explained above, we aggregated the primary commodities in three groups, energy (essentially oil), agriculture and metals.[17]

More specifically, we divide our sample in two. The first sub-period goes from 1960 till 1973 and corresponds to the Bretton-Woods period. We remove a linear trend from the three series of primary commodity prices, just as we did for the real exchange rates and consumption ratios, and compute their persistence, volatility, and correlations. Then, we calibrate the supply shocks to the three commodities in the rest of the word so as to minimize the distance between the moments in the model and the moments in the data. We repeat the same exercise for the second sub-period, the one that goes from 1973 until 2019.[18]

The main set of results is depicted in Figures 8 and 9. Figure 8 shows the distributions of the targeted moments. For example, Figure 8a shows the distribution of the standard deviation of the price of energy, P_1^x, in the first (blue bars) and second (red bars) sub-periods, both in the data (top) and in the model (bottom). The vertical lines represent the means of each distribution. As the figure shows, the mean standard deviation of energy prices generated by the model matches the mean standard deviation in the data in both sub-periods almost exactly. The same happens for the standard deviation of the prices of agriculture P^{x_2} and metal P^{x_3} in Figures 8b and 8c, respectively. As mentioned before, we calibrated the commodity shocks such that the mean standard deviations of commodity prices in the model match the means in the data in each sub-period, and Figure 8 shows that we succeeded in doing so.[19] We do not target other moments of the distributions.

After matching the moments of commodity prices, we can now evaluate how the model

[16]Relative price shocks do not affect the proper measure of output (see Kehoe and Ruhl (2008) for details) so the calibration of the two types of shocks can be done independently.

[17]In the Online appendix we list all the primary commodities in our data set and we explain in detail how we aggregated them.

[18]We keep the calibration of the productivity shocks constant across sub-periods.

[19]We also matched the persistence of commodity prices and their correlations. We show these results in the Online Appendix.

Figure 8: Benchmark results: targeted moments

(a) Standard deviation: price of energy, P^{x_1}

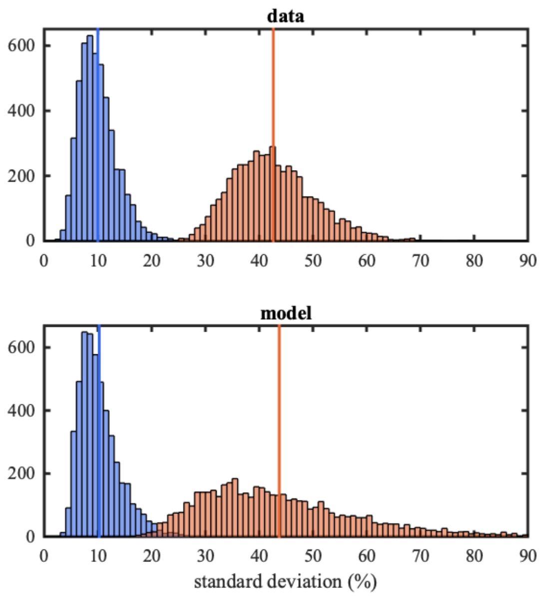

(b) Standard deviation: price of agriculture, P^{x_2}

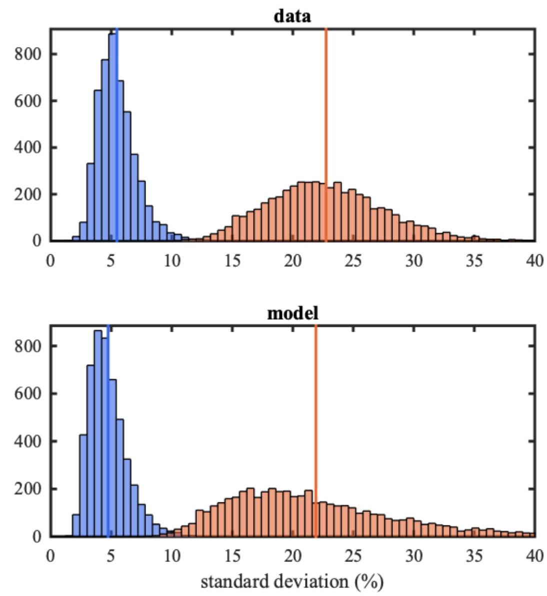

(c) Standard deviation: price of metal, P^{x_3}

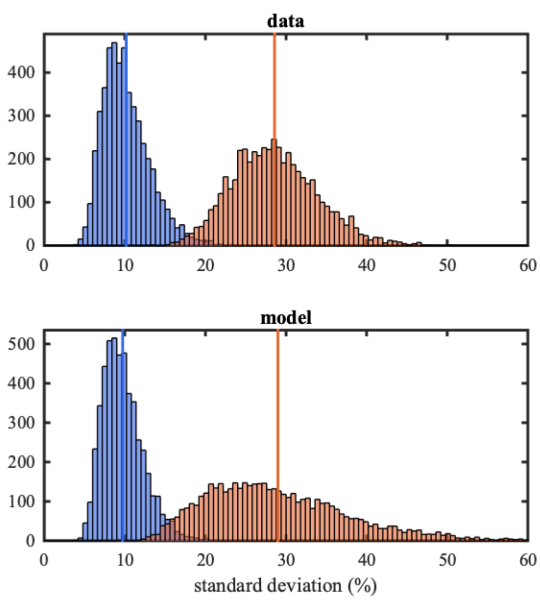

performs in terms of matching the non-targeted moments we are interested in. In particular, Figure 9 shows the distributions of the standard deviation of real exchange rates and the correlation between real exchange rates and consumption ratios.

In Figure 9a, we see that the model performs remarkably well in accounting for the volatility of the real exchange rate in both sub-periods, therefore accounting for the Mussa puzzle. The mean standard deviation of the RER is around 3 percent in the first sub-period, and it increases to 15 percent in the second sub-period, when the primary commodities become more volatile and persistent.

Figure 9: Benchmark results: non-targeted moments

(a) Standard deviation of RER

(b) Correlation of relative consumption and RER

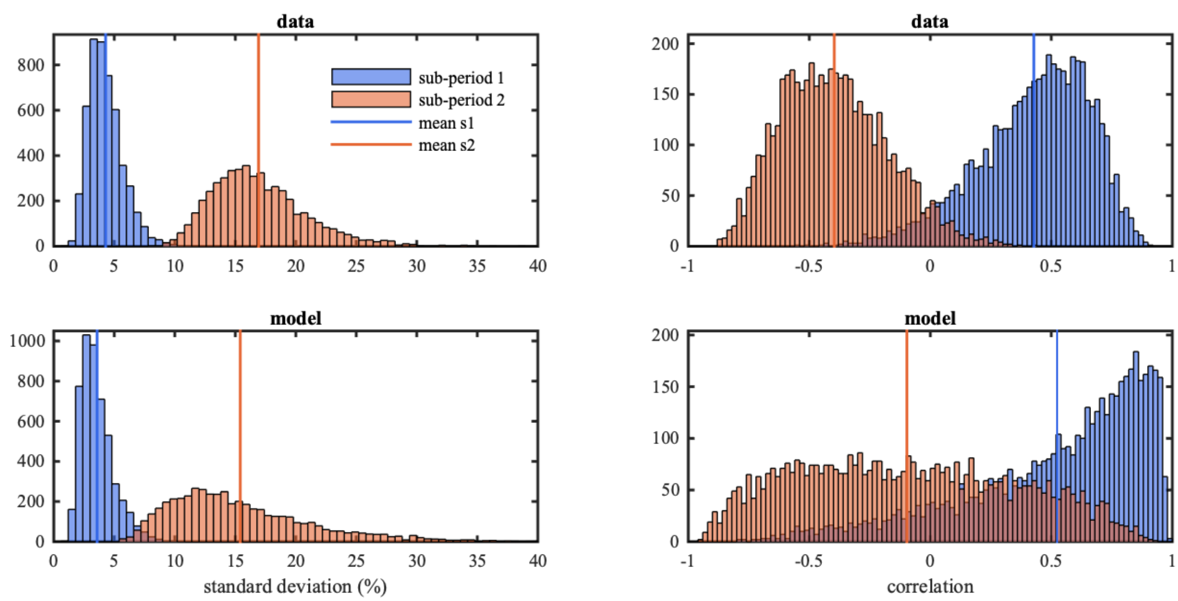

Figure 9b illustrates the distributions of the correlation between relative consumption and the RER. The model accounts for both the low correlation observed in the data, the Backus-Smith puzzle, and the changes in the correlation across sub-periods, the Mussa meets Backus-Smith puzzle. The mean of the distribution in the model in the first sub-period is around 0.5, just as in the data, and drops to around -0.1 in the second sub-period. While the mean of the distribution in the second sub-period is lower in the data, around -0.4, the distributions are wide enough and show a large area of overlap. This also shows that the mechanism we explore can generate a wide range of values for the correlation between real exchange rates and consumption ratios, consistent with the wide range of values observed across countries in Figure 4.

In interpreting these results, it is important to emphasize that they are all non-targeted moments. The only thing that varies in the model between the first and second sub-periods are the parameters of the stochastic processes of the shocks to the endowments of primary commodities in the rest of the world (Country 3). And they are not chosen freely, they are chosen to match the moments of the commodity prices in the data, as Figures 8a–8c show. No information regarding real exchange rates or consumption ratios was used.

The first conclusion we draw from this experiment is that while the shares of primary commodities are indeed quite small for both the USA and Japan in the data, the magnitude and persistence of these shocks are high enough to provide an interesting channel to explain some of the fluctuations that seem puzzling in the data. As emphasized in the introduction and in the discussion of the model, the key to generate fluctuations is the heterogeneity between countries.[20] In considering this feature, it becomes evident that the level of aggregation used in this paper implies abstracting from many dimensions of heterogeneity in the data. For instance, in considering only three commodities, we aggregated sectors that are themselves quite heterogeneous across countries.

5.1 Increasing the heterogeneity in the commodity sector

To make this statement precise, note that in calibrating the size of each sector, we started from 26 different commodities, and aggregated them into our three sectors.[21] To show how this affects the true heterogeneity in the model, let α_j^{US} and α_j^{JPN} be the share of value added generated by sector $j = 1, ..., 26$, for each country. One can define an index of heterogeneity using

$$H(US, JPN) = \frac{1}{2}\sum_{j=1}^{36} \mid \alpha_j^{US} - \alpha_j^{JPN} \mid$$

where the brackets $\mid x \mid$ define the absolute value of x.

It is straightforward to show that the index has a minimum of zero (when $\alpha_j^{US} = \alpha_j^{JPN}$ for all j) when countries are identical and a maximum of 1, when countries produce a totally different set of commodities. If we compute this index using the 26 different commodities in our IO Table, the result is 0.5. On the other hand, when we aggregate the sectors in the three groups for the calibration of the model, we reduce that index to 0.3. This is one specific dimension in which aggregation reduces the true heterogeneity in the data.[22]

To provide a quantitative measure of the relevance of this heterogeneity, we simulated

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

[21]There are 26 sectors within the commodity sector in the Input-Output Table. See the Online Appendix.

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

Figure 10: More heterogeneity in the distribution of endowments E

(a) Standard deviation of RER

(b) Correlation of relative consumption and RER

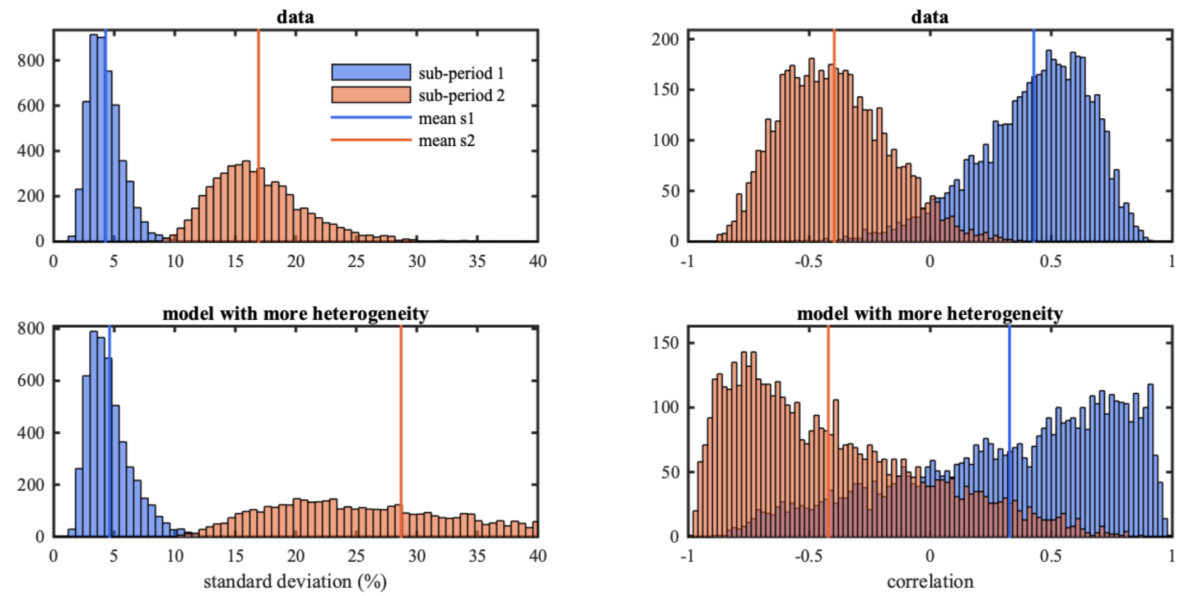

the model with the following variation. We set the relative value added of commodities 2 and 3 in the USA to be zero, and the relative value added of commodities 1 and 2 in Japan to be zero. This is a case in which the countries would be as heterogeneous as possible in the production of commodities—the index above would be equal to 1. The results are illustrated in Figure 10. The figure shows that by making the countries more heterogeneous in commodity production, the mean standard deviation of RER increases to a value closer to 30 percent in the second sub-period, substantially higher than in the data.[23] In doing so, the model generates a mean correlation between RER and consumption ratios that matches the data both in the first and second sub-periods almost exactly.

These results show that small changes to the model can have substantial effects in the standard deviation of RERs and in the correlation between RERs 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. And this requires assigning values to parameters that are hard to estimate in the data, opening the possibility to a wide range of results through this simple mechanism we explore. Given that, we conjecture that these expressive changes in the behavior of RERs and relative consumption have small welfare effects, since they are

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

Figure 11: Lower share of primary commodities in the production of intermediate goods

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

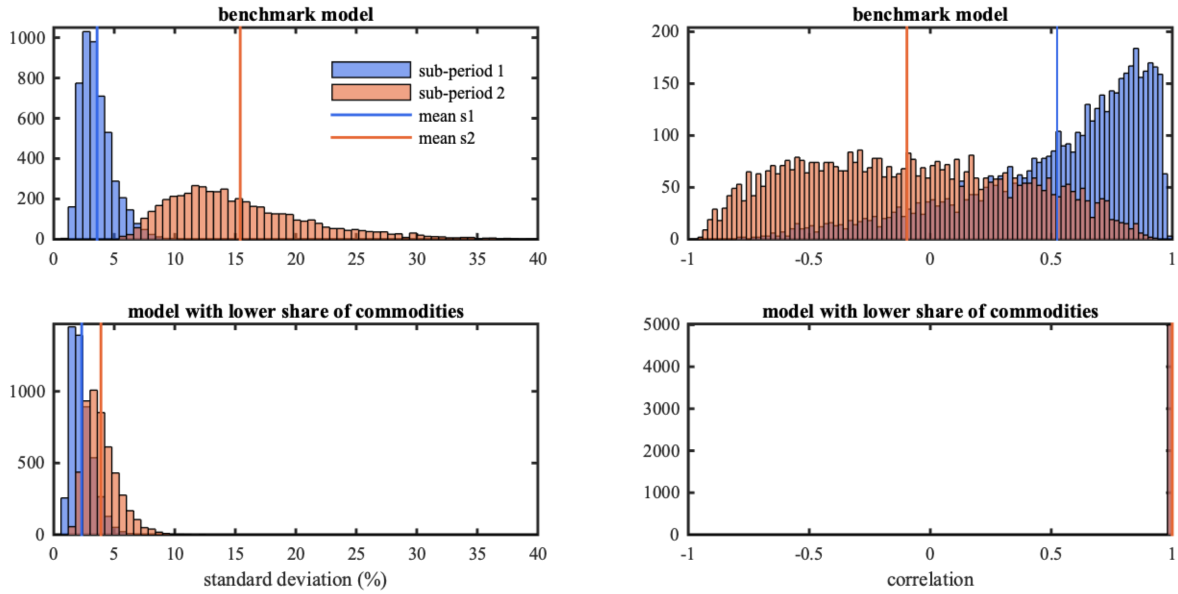

derived from small changes in the structure of these economies.

5.2 Reducing the role of primary commodities

In the next exercise, we show the effect of reducing the value added in the primary commodity sectors. We compare the benchmark model to the case in which we divide the share of primary commodities in the production of intermediate goods by 100 in each country, therefore bringing the shares 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 11 shows the results.

Without commodities, our model is unable to account for the puzzles. In particular, the model without commodities generate low volatility of RERs and a distribution of correlation between RER and consumption ratios with a point mass concentrated at one. That is, the model without commodities generates a correlation between RER and consumption ratios identical to those 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.[24]

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

References

Arezki, Rabah, Valerie A. Ramey, and Liugang Sheng. 2017. "News Shocks in Open Economies: Evidence from Giant Oil Discoveries." *The Quarterly Journal of Economics*, 132(1): 103–155.

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.

Caldara, Dario, Michele Cavallo, and Matteo Iacoviello. 2019. "Oil price elasticities and oil price fluctuations." *Journal of Monetary Economics*, 103: 1 – 20.

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.

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.

Itskhoki, Oleg, and Dimitry Mukhin. 2020. "Exchange Rate Disconnect in General Equilibrium." UCLA Working Paper.

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

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.

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.