

Can International Macroeconomic Models Explain Low-Frequency Movements of Real Exchange Rates?*

Pau Rabanal[†]

Juan F. Rubio-Ramírez[‡]

October 17, 2011

Abstract

Real exchange rates exhibit important low-frequency fluctuations. This makes the analysis of real exchange rates at all frequencies a more sound exercise than the typical business cycle one, which compares actual and simulated data after the Hodrick-Prescott filter is applied to both. A simple two-country, two-good model, as described in Heathcote and Perri (2002), can explain the volatility of the real exchange rate when all frequencies are studied. The puzzle is that the model generates too much persistence of the real exchange rate instead of too little, as the business cycle analysis asserts. Finally, we show that the introduction of adjustment costs in production and in portfolio holdings allows us to reconcile theory and this feature of the data.

JEL Classification: E32, F32, F33, F41.

Keywords: International Business Cycles, Spectrum, Real Exchange Rates, Cointegration.

*We thank George Alessandria, Boragan Aruoba, Sanjay Chugh, John Haltiwanger, Federico Mandelman, Enrique Mendoza, Emi Nakamura, Jorge Roldós, John Shea, Pedro Silos, Jón Steinsson, Carlos Végh and seminar audiences at the University of Maryland, Banco de España CEPR/ESSIM meeting in Tarragona, CEMFI, and the Federal Reserve Banks of Atlanta, Dallas, and Philadelphia for useful comments. We also thank Hernán Seoane and Béla Személy for their research support. Beyond the usual disclaimer, we must note that any views expressed herein are those of the authors and not necessarily those of the International Monetary Fund, the Federal Reserve Bank of Atlanta, or the Federal Reserve System. Finally, Juan F. Rubio-Ramírez also thanks the NSF for financial support.

[†]IMF Institute, International Monetary Fund, <prabanal@imf.org>.

[‡]Duke University, Federal Reserve Bank of Atlanta, CEPR, FEDEA, and BBVA Research. <juan.rubio-ramirez@duke.edu>.

1. Introduction

This paper challenges the conventional wisdom that a baseline international real business cycle (IRBC) two-country, two-good model, such as the one described in Heathcote and Perri (2002), cannot generate either enough volatility or enough persistence in the real exchange rate (RER) when compared to the data. When the object of interest is RER fluctuations at all frequencies, instead of business cycle (BC) frequencies only, this model can explain the standard deviation of the U.S. dollar RER. However, the model implies a higher persistence of the RER than in the data.

We advocate that analyzing RER fluctuations at all frequencies is a more compelling exercise than just studying the BC ones. Spectral analysis shows that most of the variance of the RER in the data can be assigned to low-frequency movements (about 70 percent), while movements at BC frequencies account for only a small share of the RER fluctuations (just 25 percent). The baseline IRBC model accounts for the area below the spectrum of the RER, i.e., its standard deviation, but not for its shape, since it places a larger share of fluctuations of the RER in low-frequency movements than in the data. We call this shortcoming of the model the “excess persistence of RER” puzzle. We show that extending the model to consider adjustment costs in the composition of domestic and imported intermediate input and portfolio adjustment costs helps to solve this puzzle (i.e., replicating the shape of the spectrum) while still explaining the standard deviation of the RER (i.e., the area below the spectrum).

Since the seminal works of Backus, Kehoe, and Kydland (1992) and Baxter and Crucini (1995), the IRBC literature has been preoccupied with explaining the international transmission of shocks, the cyclical comovement of variables across countries, and the behavior of international relative prices. As in the real business cycle (RBC) literature, the IRBC literature mainly concentrates on explaining the BC fluctuations of the data. The success of the model is measured by its ability to reproduce selected second moments of Hodrick-Prescott (HP) filtered data, which removes trends and low-frequency movements. Other papers use instead the band-pass filter, as described in Baxter and King (1999) or Christiano and Fitzgerald (2003). The researcher compares the second moments of actual data with those implied by artificial data generated by the model after the same detrending procedure has been applied to both. One of the most relevant facts in the HP-filtered data is that international relative prices are more volatile than output and highly persistent. IRBC models with reasonable calibrations have a hard time reproducing these

features. In earlier work Backus, Kehoe and Kydland (1994) and Stockman and Tesar (1995) showed that IRBC models cannot match the volatility of the HP-filtered terms of trade, while, in a more recent contribution, Heathcote and Perri (2002) have pointed out the standard IRBC model’s inability to explain the volatility and persistence of the HP-filtered RER.

In this paper, we first argue that analyzing only the BC fluctuations of the RER leads researchers to miss a large part of the story. The reason is as follows. The top panel of Figure 1 plots the (log) U.S. dollar RER along with its implied HP-filtered “trend” using a bandwidth of 1600. Just from eyeballing, it is evident that most of the fluctuations in the U.S. dollar RER have been low-frequency movements. This observation is confirmed by the spectral analysis that we perform in Section 2: most of the variation of the RER in the data is at frequencies lower than BC fluctuations (it is 70 percent for the U.S. dollar, and between 60 to 75 percent depending on the currency we examine). These low-frequency movements are removed by HP-filtering.¹

Second, motivated by the argument above, we propose to analyze the fluctuations of the RER at all frequencies instead. Therefore, we need to consider a model able to generate low-frequency fluctuations in the RER. Our baseline model is an extension of the two-country, two-good model of Heathcote and Perri (2002) in which stochastic processes for total factor productivity (TFP) are non-stationary but cointegrated across countries.² We show that the model can explain about 80 percent of the standard deviation of the RER in the data while closely matching the volatility of output growth when we use a benchmark calibration of the model, including a value of 0.85 for the elasticity of substitution between intermediate inputs in the production of the final good. However, in the model, the RER is too persistent and the spectrum places too much weight on low-frequency fluctuations (in the model 85 percent of the variance is caused by low-frequency fluctuations while it is 70 percent in the data). In order to solve this shortcoming, we extend the model with adjustment costs in the use of intermediate imported inputs for the production of the final good (see Erceg, Guerrieri, and Gust, 2006). The presence of these costs allows us to combine a low short-run elasticity of substitution between imported and domestic intermediate

¹The RER in emerging markets can have a trend, in particular in those emerging economies that experience higher productivity growth rates than advanced economies. In that case, the use of a trend/cycle decomposition would be justified. However, the focus of most of the IRBC literature is to explain the RER of the U.S. dollar vis-a-vis other industrialized countries. In that case RERs are highly persistent series, but they do not have a trend.

²In related work, Rabanal, Rubio-Ramírez and Tuesta (2011) show that cointegrated TFP shocks improve the model’s ability to explain certain features of the HP-filtered data, including RER volatility.

goods, which is needed to increase the volatility of the RER at BC frequencies, with a higher long-run elasticity, which is needed to reduce the excessive volatility of the RER at low frequencies. We show how these input adjustment costs, together with portfolio adjustment costs, help to solve the puzzle by increasing the impact response of the RER in the short run while reducing it at long-run horizons in the model.

The paper is organized as follows: Section 2 presents the spectral analysis of the U.S. dollar RER as well as that of other main currencies. Section 3 discusses the related literature, while Section 4 presents a baseline IRBC model. Section 5 presents the calibration and the results of the baseline model. In Section 6, we present the extensions to the model and show how they help reconcile theory and evidence. Section 7 concludes.

2. Spectral Analysis of the RER

In this section we study the spectrum of the RER of six main currencies: the U.S. dollar, the euro, the UK pound sterling, the Japanese yen, and the Canadian and Australian dollars. In order to find the longest possible time series for each currency, we choose between the IMF's International Financial Statistics (IFS) database or the measure constructed from national central banks. We verify that for the period during which both measures overlap the correlation is very high, denoting that both sources use similar methodologies to construct the RER series.

Our data sources are as follows: for the U.S. dollar we obtain the real effective exchange rate (REER) series from the Federal Reserve's Real Broad Trade-Weighted Value of the U.S. dollar. The sample period is 1973:Q1-2010:Q3. For the euro area, we use the REER series coming from the European Central Bank's Area Wide Model (sample period 1973:Q1-2008:Q4), which we extend up to 2010:Q3 using IFS data. For the Canadian dollar and the U.K. pound sterling we use the IFS measure (sample period 1975:Q2-2010:Q3). For the Australian dollar, we use the REER measure constructed by the Reserve Bank of Australia (sample period 1973:Q1-2010Q4). For the Japanese yen, we use the REER measure constructed by the Bank of Japan using the BIS methodology (sample period 1973:Q1-2010:Q3).

The spectrum contains the same information as auto-correlations and it allows us to decompose the variance of the RER across different frequencies. In order to estimate the spectrum we use the modified Bartlett kernel methodology described in Section 6.4 of Hamilton (1994). In Figure 1 we present the time series for the (log) U.S. dollar RER along with its implied HP-filtered

“trend,” its autocorrelogram, and the estimated spectrum density. From the first two panels of Figure 1, we can observe that the U.S. dollar RER does not have an evident time trend. At the same time, it is a highly persistent series: the autocorrelogram decays monotonically as the lag length is increased, but it decays slowly. As a result, the correlation between the RER and its own 15th lag is basically zero. In the bottom panel of Figure 1 we present the estimated spectrum, where we have shaded the area corresponding to BC frequencies: most fluctuations occur at low frequencies. The facts presented in Figure 1 are common to all the other major currencies we studied.³ In all cases, the low-frequency movements implied by the HP-filtered “trend” are quite sizable, the autocorrelogram decays at a slow rate (but fast enough to suggest there is not a unit root), and the estimated spectrum suggests that most fluctuations occur at low frequencies.

We put some numbers to this last claim by decomposing the variance of each RER into BC frequencies (8 to 32 quarters), lower than BC frequencies (more than 32 quarters) and higher than BC frequencies (less than 8 quarters) in Table 1. We also report the results coming from constructing our own U.S. dollar RER series by recomputing the RER against the following four countries: Japan, Canada, the U.K., and Australia, and the euro area. These four countries and the euro area are used later in the paper to calibrate the “rest of the world” TFP process; hence, for consistency it makes sense to compute the RER vis-a-vis this group. We compute bilateral RERs and aggregate them by using the currency weights from the Broad Index of the Foreign Exchange Value of the dollar computed by the U.S. Federal Reserve.⁴

As shown in Table 1, most of the variance of the U.S. dollar RER (74 percent) is concentrated at low frequencies (less than 32 quarters), while 21.1 percent of the variance is attributed to BC frequencies and only 5 percent occurs at high frequencies. Our measure vis-à-vis main industrialized countries behaves similarly. Taking an international comparison, the fraction of the variance concentrated at low-frequency movements ranges from 62.2 percent for the U.K. pound sterling to 75.7 percent for the Australian dollar. Therefore, the literature that tries to explain BC-frequency fluctuations of RERs misses a large part of the picture that resides in the low-frequency end of the spectrum.

³To save space, we do not repeat Figure 1 for the rest of the major currencies, but they are available upon request

⁴For a description see <http://www.federalreserve.gov/releases/H10/Weights/>.

Table 1: Variance Decomposition
of the RER (in percent)

	Low	BC	High
U.S.-Federal Reserve	73.9	21.1	5.0
U.S.-Our measure	70.0	24.3	5.7
Euro Area	64.5	29.8	5.7
U.K.	62.2	31.0	6.8
Japan	71.3	23.8	4.9
Australia	75.7	19.2	5.1
Canada	74.1	20.0	5.9

The finding that most of the variance of the RER is concentrated at low frequencies can be related to two well-documented facts. First, the large half-life of estimated IRFs of the RER (Rogoff, 1996; Murray and Papell, 2002; and Steinsson, 2008) and second, its hump-shaped dynamics (Huizinga, 1987; Eichenbaum and Evans, 1995; Cheung and Lai, 2000; and Steinsson, 2008). Both the large half-life and the dynamic non-monotonic response pattern are closely related to the high persistence of RERs in the data and to the importance of low-frequency fluctuations.

3. Relationship to the Literature

This paper bridges the gap between empirical models and dynamic stochastic general equilibrium (DSGE) models in explaining RER fluctuations. The empirical literature since the seminal work of Meese and Rogoff (1983) has mostly used univariate and multivariate time series methods to model exchange rates (nominal or real). This analysis is mostly performed at all frequencies. In a recent paper, Steinsson (2008) follows a large literature that models the linear univariate empirical properties of the RER. Other univariate nonlinear time series approaches are reviewed in Sarno (2003). In the multivariate setup, Clarida and Galí (1994) and Faust and Rogers (2003), among many others, have used VAR models to explain the response of exchange rates (both real and nominal) to several shocks. Another branch of the literature studies the role of world and country-specific factors in explaining the comovement of main macroeconomic variables across countries within the context of dynamic factor models (see, for instance, Mumtaz and Surico, 2009). Other authors examine the relationship between exchange rates (both real and nominal)

and fundamentals derived from open economy macro models, such as Engel and West (2005), and Cheung, Chinn and Garcia-Pascual (2005).

However, most calibrated DSGE models are typically concerned with explaining the BC fluctuations of the RER and hence analyze HP-filtered data. Since Heathcote and Perri (2002), the literature has been energetically trying to reconcile the discrepancy between theory and HP-filtered RER data, with some success. For example, Chari, Kehoe and McGrattan (2002) show that a monetary economy with monopolistic competition and sticky prices can explain HP-filtered RER volatility if a high degree of risk aversion is assumed and Corsetti, Dedola and Leduc (2008) show that introducing nontraded goods also helps reconcile theory with data. Rabanal, Rubio-Ramírez and Tuesta (2011) show that introducing cointegrated total factor productivity (TFP) processes across countries helps to explain the volatility of the HP-filtered RER. Although such models do a better job explaining the volatility of the HP-filtered RER, they still cannot match its persistence. A number of related papers have tried to tackle the lack of persistence of RER in the model in the context of monetary models (for example, see Bergin and Feenstra, 2001, Benigno, 2004, or Bouakez, 2005) without completely addressing it.

In this paper we combine the two approaches by comparing the properties of the RER in the DSGE model and in the data, without applying any filtering method. It is also worth noting that a few recent exceptions to this filtering practice arise in the literature that estimates open economy DSGE models with Bayesian methods. Adolfson et al. (2007) and Rabanal and Tuesta (2010) include the log of the RER in the set of observable variables, while Nason and Rogers (2008) use the log of the nominal exchange rate between the U.S. and Canadian dollars in their estimated model. Also, there are two recent exceptions to the practice of focusing only on BC fluctuations of the data and comparing them to the model. Baxter (2011) finds that there is evidence in favor of risk sharing across countries at medium and low frequencies. Corsetti, Dedola, and Viani (2011) study the correlation between the RER and the ratio of consumption levels across countries (which is known as the “Backus-Smith puzzle”) at both BC and low frequencies.

4. The Baseline Model

As a baseline we use a two-country, two-good model similar to the one described in Backus, Kehoe and Kydland (1994) and Heathcote and Perri (2002) with a main important difference: (the log of) TFP processes are assumed to be non-stationary but cointegrated across countries. In other

words, they follow a VECM process.⁵

To keep exposition to a minimum, we present only the problem of home-country households, home-country firms, and market clearing. Then we will describe the equilibrium conditions. In terms of notation, we use an asterisk superscript when we refer to the foreign-country variable analogous to a home-country variable (i.e., if C_t is consumption in the home country, then C_t^* is consumption in the foreign country). In each country, a single final good is produced by a representative competitive firm that uses intermediate goods from both countries in the production process. These intermediate goods are imperfect substitutes for each other and can be purchased from representative competitive producers of intermediate goods in both countries. Intermediate goods producers use domestic capital and domestic labor in the production process and face a domestic TFP shock. The final good can only be domestically consumed or domestically invested in by domestic households. Thus, all trade of goods between countries occurs at the intermediate goods level. In addition, households trade across countries an uncontingent international riskless bond denominated in units of the home-country intermediate good. No other financial asset is available.

4.1. Households

The representative household of the home country solves:

$$\max_{\{C_t, L_t, X_t, K_t, D_t\}} E_0 \sum_{t=0}^{\infty} \beta^t \frac{[C_t^\tau (1 - L_t)^{1-\tau}]^{1-\sigma}}{1 - \sigma}$$

subject to the following budget constraint:

$$P_t (C_t + X_t) + P_{H,t} \overline{Q}_t D_t \leq P_t (W_t L_t + R_t K_{t-1}) + P_{H,t} [D_{t-1} - \Phi(D_t, A_{t-1})] \quad (1)$$

and the law of motion for capital:

$$K_t = (1 - \delta) K_{t-1} + X_t.$$

⁵Rabanal, Rubio-Ramírez and Tuesta (2011) show that TFP processes between the U.S. and a sample of main industrialized countries are cointegrated and that the low estimated speed of convergence to the cointegrating relationship is a key ingredient for the model to explain the volatility of the RER at BC frequencies. Here, we examine how the same model performs in explaining movements of the RER at all frequencies. Since the model is the same as in the above-mentioned reference, we just show the main functional forms and optimality conditions and refer the reader to the original paper for a detailed derivation.

The following notation is used: β is the discount factor, L_t is the fraction of time allocated to work in the home country, C_t are units of consumption of the final good, X_t are units of investment, and K_t is the capital stock in the home country at the beginning of period $t + 1$. P_t is the price of the home country final good, which will be defined below; W_t is the hourly wage in the home country, and R_t is the home country rental rate of capital, where the prices of both factor inputs are measured in units of the final good. $P_{H,t}$ is the price of the home-country intermediate good, D_t denotes the holdings of the internationally traded riskless bond that pays one unit of the home-country intermediate good (minus a small cost of holding bonds, $\Phi(\cdot)$) in period $t + 1$ regardless of the state of nature, and \overline{Q}_t is its price, measured in units of the home-country intermediate good. The function $\Phi(\cdot)$ measures the cost of holding bonds measured in units of the home-country intermediate good.⁶

Following the existing literature, $\Phi(\cdot)$ takes the functional form:

$$\Phi(D_t, A_{t-1}) = \frac{\phi}{2} A_{t-1} \left(\frac{D_t}{A_{t-1}} \right)^2$$

where we have modified the adjustment cost function to include the home-country TFP level, A_t , which is characterized below, to ensure balanced growth.

4.2. Firms

We now describe the production function and profit maximization problems of the final and intermediate goods producers. Then, we portray technology.

4.2.1. Final goods producers

The final good in the home country, Y_t , is produced using home-country intermediate goods, $Y_{H,t}$, and foreign-country intermediate goods, $Y_{F,t}$, with the following technology:

$$Y_t = \left[\omega^{\frac{1}{\theta}} Y_{H,t}^{\frac{\theta-1}{\theta}} + (1 - \omega)^{\frac{1}{\theta}} Y_{F,t}^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}} \quad (2)$$

⁶The $\Phi(\cdot)$ cost is introduced to ensure stationarity of the level of D_t in IRBC models with incomplete markets, as discussed by Heathcote and Perri (2002). In this baseline model we choose the cost to be numerically small, so it does not affect the dynamics of the rest of the variables. This will not be the case when we analyze some of the extensions.

where ω denotes the fraction of home-country intermediate goods that are used for the production of the home-country final good and θ is the elasticity of substitution between home-country and foreign-country intermediate goods. Therefore, the representative final good producer in the home country solves the following problem:

$$\max_{Y_t, Y_{H,t}, Y_{F,t}} P_t Y_t - P_{H,t} Y_{H,t} - P_{F,t}^* Y_{F,t}$$

subject to the production function (2), where $P_{F,t}^*$ is the price of the foreign-country intermediate good in the home country.

4.2.2. Intermediate goods producers

The representative intermediate goods producer in the home country uses domestic labor and domestic capital in order to produce home-country intermediate goods and sells her product to both the home-country and foreign-country final good producers. Taking prices of all goods and factor inputs as given, she maximizes profits by solving:

$$\max_{L_t, K_{t-1}} P_{H,t} Y_{H,t} + P_{H,t}^* Y_{H,t}^* - P_t (W_t L_t + R_t K_{t-1})$$

subject to the production function:

$$Y_{H,t} + Y_{H,t}^* = A_t^{1-\alpha} K_{t-1}^\alpha L_t^{1-\alpha} \quad (3)$$

where $Y_{H,t}^*$ is the amount of home-country intermediate goods sold to the foreign-country final good producers and $P_{H,t}^*$ is the price of the home-country intermediate good in the foreign country.

4.2.3. TFP processes

We assume that $\log A_t$ and $\log A_t^*$ are cointegrated of order $C(1, 1)$. This assumption involves specifying the following VECM for the law of motion driving the log first difference of TFP processes for both the home and the foreign country:

$$\begin{pmatrix} \Delta \log A_t \\ \Delta \log A_t^* \end{pmatrix} = \begin{pmatrix} c \\ c^* \end{pmatrix} + \begin{pmatrix} \kappa \\ \kappa^* \end{pmatrix} [\log A_{t-1} - \gamma \log A_{t-1}^* - \log \xi] + \begin{pmatrix} \varepsilon_t \\ \varepsilon_t^* \end{pmatrix} \quad (4)$$

where $(1, -\gamma)$ is the cointegrating vector, ξ is the constant in the cointegrating relationship, $\varepsilon_t \sim N(0, \sigma)$ and $\varepsilon_t^* \sim N(0, \sigma^*)$, ε_t and ε_t^* can be correlated, and Δ is the first-difference operator.

4.3. Market Clearing

The model is closed with the following market clearing conditions in the final good markets:

$$C_t + X_t = Y_t \quad (5)$$

and in the international bond market:

$$D_t + D_t^* = 0. \quad (6)$$

4.4. Equilibrium Conditions

At this point, it is useful to define the following relative prices: $\tilde{P}_{H,t} = \frac{P_{H,t}}{P_t}$, $\tilde{P}_{F,t}^* = \frac{P_{F,t}^*}{P_t^*}$ and $RER_t = \frac{P_t^*}{P_t}$ where P_t^* is the price of the foreign-country final good. Note that $\tilde{P}_{H,t}$ is the price of home-country intermediate goods in terms of the home-country final good, $\tilde{P}_{F,t}^*$ is the price of foreign-country intermediate goods in terms of the foreign-country final good, which appears in the foreign-country's budget constraint, and RER_t is the RER between the home and foreign countries. The law of one price (LOP) holds: $P_{H,t} = P_{H,t}^*$ and $P_{F,t} = P_{F,t}^*$.

The equilibrium conditions include the first order conditions of households, intermediate and final good producers in both countries, as well as the relevant laws of motion, production functions, and market clearing conditions. Here, we detail the home-country equilibrium conditions only. The foreign-country conditions are very similar, with the appropriate change of notation. The marginal utility of consumption and the labor supply are given by:

$$U_{C_t} = \lambda_t, \quad (7)$$

$$\frac{U_{L_t}}{U_{C_t}} = W_t, \quad (8)$$

where U_x denotes the partial derivative of the utility function U with respect to variable x . The first order condition with respect to capital delivers an intertemporal condition that relates the

marginal rate of consumption to the rental rate of capital and the depreciation rate:

$$\lambda_t = \beta E_t [\lambda_{t+1} (R_{t+1} + 1 - \delta)].$$

The law of motion of capital is:

$$K_t = (1 - \delta) K_{t-1} + X_t. \quad (9)$$

The optimal savings choice delivers the following expression for the price of the riskless bond:

$$\overline{Q}_t = \beta E_t \left(\frac{\lambda_{t+1}}{\lambda_t} \frac{\tilde{P}_{H,t+1}}{\tilde{P}_{H,t}} \right) - \frac{\Phi'(D_t)}{\beta}. \quad (10)$$

The next condition uses the expression for the price of the bond in both countries to derive the expression for optimal risk sharing across countries:

$$E_t \left[\frac{\lambda_{t+1}^*}{\lambda_t^*} \frac{\tilde{P}_{H,t+1}}{\tilde{P}_{H,t}} \frac{RER_t}{RER_{t+1}} - \frac{\lambda_{t+1}}{\lambda_t} \frac{\tilde{P}_{H,t+1}}{\tilde{P}_{H,t}} \right] = -\frac{\Phi'(D_t)}{\beta}. \quad (11)$$

From the intermediate goods producers' maximization problems, labor and capital are paid their marginal product, where the rental rate of capital and the real wage are expressed in terms of the final good in each country:

$$W_t = (1 - \alpha) \tilde{P}_{H,t} A_t^{1-\alpha} K_{t-1}^\alpha L_t^{-\alpha} \quad (12)$$

and

$$R_t = \alpha \tilde{P}_{H,t} A_t^{1-\alpha} K_{t-1}^{\alpha-1} L_t^{1-\alpha}. \quad (13)$$

From the final good producers' maximization problem, the demand for home and foreign country intermediate goods depends on their relative price:

$$Y_{H,t} = \omega \tilde{P}_{H,t}^{-\theta} Y_t, \quad (14)$$

$$Y_{F,t} = (1 - \omega) \left(\tilde{P}_{F,t}^* RER_t \right)^{-\theta} Y_t. \quad (15)$$

Using the production functions of the final good:

$$Y_t = \left[\omega^{\frac{1}{\theta}} Y_{H,t}^{\frac{\theta-1}{\theta}} + (1-\omega)^{\frac{1}{\theta}} Y_{F,t}^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}}, \quad (16)$$

(14) and (15), the final good deflator in the home-country is:

$$P_t = \left[\omega P_{H,t}^{1-\theta} + (1-\omega) P_{F,t}^{1-\theta} \right]^{\frac{1}{1-\theta}}.$$

Hence, given that the LOP holds, the RER is equal to:

$$RER_t = \frac{P_t^*}{P_t} = \frac{\left[\omega P_{F,t}^{1-\theta} + (1-\omega) P_{H,t}^{1-\theta} \right]^{\frac{1}{1-\theta}}}{\left[\omega P_{H,t}^{1-\theta} + (1-\omega) P_{F,t}^{1-\theta} \right]^{\frac{1}{1-\theta}}}.$$

Note that the only source of RER fluctuations is the presence of home bias ($\omega > 1/2$). Also, intermediate goods, final good, and bond markets clear as in equations (3), (5), and (6). Finally, the law of motion of the level of bonds:

$$\tilde{P}_{H,t} \bar{Q}_t D_t = \tilde{P}_{H,t} Y_{H,t}^* - RER_t \tilde{P}_{F,t} Y_{F,t} + \tilde{P}_{H,t} D_{t-1} - \tilde{P}_{H,t} \Phi(D_t, A_{t-1}) \quad (17)$$

is obtained using (1) and the fact that intermediate and final good producers make zero profits. Finally, the TFP shocks follow the VECMs described above. Since the model is non-stationary, we need to normalize it and check for the existence of a balanced growth path. Rabanal, Rubio-Ramírez and Tuesta (2011) find that the estimated γ is one, which is a sufficient condition for balanced growth to exist in this economy (in addition to the standard restrictions on technology and preferences, as in King, Plosser and Rebelo, 1988). Hence, along the balanced growth path, real variables in each country grow at the same rate as its TFP. To solve and simulate the model, we normalize real variables in each country by the lagged level of TFP in that country to obtain a stationary system. Then, we take a log-linear approximation to the normalized equilibrium conditions.

5. Results of the Baseline Model

In this section we describe the results of the baseline model. First, we describe the benchmark calibration for the baseline model. Then, we show that the baseline model with the benchmark calibration can closely replicate the standard deviation of the RER when all frequencies are considered. In other words, it reproduces the area below the RER spectrum. However, we also show that the model cannot replicate the shape of the spectrum. It assigns too much variance of the RER to fluctuations with frequencies below BC ones when compared to the data. This is what we call the “excess persistence of RER” puzzle. Finally, we show that these findings are robust to some standard changes in the literature such as assuming stationary TFP shocks or cointegrated investment-specific technology (IST) shocks.

5.1. Benchmark Calibration for the Baseline Model

Our benchmark calibration follows that in Heathcote and Perri (2002) closely. The model is quarterly. The discount factor β is set equal to 0.99, which implies an annual real rate of 4 percent. In the utility function, we set the consumption share τ to 0.34 and the coefficient of risk aversion σ to 2. Parameters on technology are fairly standard in the literature. Thus, the depreciation rate δ is set to 0.025, the capital share of output α is set to 0.36, and the ratio of intermediate inputs in the production of the final good ω is set to 0.9, which matches the actual import/output ratio in the steady state. We calibrate the elasticity of substitution between intermediate goods to $\theta = 0.85$. We will also consider other values of θ to check the robustness of our results. We assume a cost of bond holdings, ϕ , of 1 basis point (0.01).

The calibration of the VECM process follows the estimates in Rabanal, Rubio-Ramírez and Tuesta (2011). Their paper constructed series of TFP for the United States and for a “rest of the world” aggregate of main industrialized trade partners of the U.S. (Australia, Canada, Euro Area, Japan, and the U.K.) using data on output, employment, hours and capital stock. They tested for and confirmed the presence of unit roots in each series and cointegration between the two TFP series using Johansen’s (1991) test. Finally, they estimate a process like (4). In addition to not rejecting that $\gamma = 1$, they find that (i) zero lags are necessary and (ii) cannot reject that $\kappa = -\kappa^*$ (i.e., that the speed of convergence to the cointegrating relationship is the same for both countries). Following their estimates, we set $\gamma = 1$, $\kappa = -0.007$, $c = 0.001$, $c^* = 0.006$,

$\sigma = 0.0108$ and $\sigma^* = 0.0088$.

5.2. Matching the RER Spectrum

Figure 2 presents the spectrum of the RER implied by our baseline model under the benchmark calibration and compares it with the estimated spectrum for our constructed measure of the U.S. dollar RER. Our measure includes the same countries that we considered when constructing the “rest of the world” TFP. Since we can compute the theoretical moments of the growth rates of variables and of the RER implied by the model, then it is possible to compute the theoretical spectrum of the RER.

Table 2 displays some key statistics of the RER implied by the baseline model under the benchmark calibration and compares them to the data. The same table also shows results for alternative values for θ . The baseline model with the benchmark calibration can closely replicate the standard deviation of the RER when compared to the data (8.33 in the model versus 10.56 in the data), and also gets the standard deviation of output growth about right (0.75 in the model versus 0.8 in the data). However, Figure 2 and Table 2 highlight the model’s main problem. It assigns too large of a share of the variance of the RER to low-frequency fluctuations: almost 89 percent in the model versus 70 percent in the data. This result is related to the usual finding that the model cannot explain the volatility of the HP-filtered RER because it is precisely the low-frequency component that is removed with the HP filter.⁷ As mentioned above, we call this discrepancy between the model and the data the “excess persistence of RER” puzzle.

Next, we present results for $\theta = 0.62$. This is a relevant value because Rabanal, Rubio-Ramírez and Tuesta (2011) found that it allowed the model to match the relative volatility of the HP-filtered RER with respect to HP-filtered output. The model now implies a larger standard deviation of the RER than in the data (16.2 versus 10.56). The shape of the RER spectrum does not change much and most of the volatility (88 percent) is again assigned to low-frequency movements. Hence, in order to match the standard deviation of the HP-filtered RER, the model generates too much volatility of the RER at all frequencies. Finally, we also analyze the implications of the value of $\theta = 1.5$ (which is used by Chari, Kehoe and McGrattan, 2002, and Erceg, Guerrieri and Gust, 2006). As expected, the model explains less of the volatility

⁷Rabanal, Rubio-Ramírez and Tuesta (2011) found that when $\theta = 0.85$, this exact same model can explain only about half of the volatility of the HP-filtered RER.

of the RER (3.55 versus 10.56) and the shape of the spectrum is basically the same. Hence, while the standard deviation of the RER at all frequencies is inversely related to the elasticity of substitution, θ , the shape of the spectrum seems to be invariant to it. Low values of θ help to explain RER variance (the area under the spectrum) but do not solve the “excess persistence of RER” puzzle (the shape of the spectrum).

Table 2: Implications of the Model with Only TFP					
	Standard Deviation		Frequency of RER		
	RER	Output Growth	Low	BC	High
U.S. Data	10.56	0.8	70.0	24.3	5.7
$\theta = 0.85$	8.33	0.75	88.4	8.8	2.8
$\theta = 0.62$	16.02	0.68	88.2	8.9	2.9
$\theta = 1.5$	3.55	0.85	89.0	8.4	2.6

Note: RER denotes the log of the RER. Output is real GDP. Growth rates are computed taking the first differences of the logs.

5.3. Some Robustness

We have found that the model’s main failure is the “excess persistence of RER” puzzle. In this subsection, we perform some robustness analysis to determine whether the puzzle survives after simple modifications of the model. In particular, we analyze two variations that involve different assumptions on the shocks that drive the model. First, we use the Heathcote and Perri (2002) estimates for the joint evolution of stationary TFP shocks. Second, we use the cointegrated TFP and IST shocks as in Mandelman et al. (2011). Results are reported in Table 3. We use the label “Stationary” to refer to the Heathcote and Perri (2002) model, and we use “TFP and IST” to refer to the model with cointegrated TFP and IST shocks.

Table 3: Robustness					
	Standard Deviation		Frequency of RER		
	RER	Output Growth	Low	BC	High
U.S. Data	10.56	0.8	70.0	24.3	5.7
Stationary, $\theta = 0.85$	4.03	0.92	85.3	11.3	3.3
Stationary, $\theta = 0.62$	7.35	0.87	85.5	11.1	3.4
Stationary, $\theta = 1.5$	1.86	0.97	85.8	11.3	2.9
TFP and IST, $\theta = 0.85$	8.58	0.76	88.5	8.7	2.8
TFP and IST, $\theta = 0.62$	16.61	0.65	88.2	8.9	2.9
TFP and IST, $\theta = 1.5$	3.77	0.82	89.1	8.3	2.6

Note: RER denotes the log of the RER. Output is real GDP. Growth rates are computed taking the first differences of the logs.

Heathcote and Perri (2002) estimate a VAR(1) in levels to model the joint behavior of TFP processes across countries (the U.S. and a “rest of the world” aggregate). When we use their estimated process, we find that their model cannot explain the volatility of the RER. With their benchmark calibration using $\theta = 0.85$ the model explains less than 40 percent of the standard deviation of the RER. Even reducing the value of θ to 0.62 is not enough. As explained in Rabanal, Rubio-Ramírez and Tuesta (2011), the presence of a common unit root and slow transmission of shocks across countries is a crucial ingredient to explain large RER volatility, and this feature is missing in Heathcote and Perri (2002). Note that the model with stationary TFP shocks assigns somehow less volatility to low-frequency fluctuations than the baseline model, but the differences are not relevant and the results are still far away from matching the data. Next, we look at what happens when we go back to the case of cointegrated TFP shocks but also introduce cointegrated IST shocks, as estimated by Mandelman et al. (2011). Including IST shocks results in marginal changes for explaining RER volatility and the spectrum.

The conclusion of this section is that, while the baseline model can replicate the area below the spectrum of the RER for low values of the elasticity of substitution, it has a hard time reproducing its shape because too much weight is placed on low-frequency fluctuations. In addition, none of the modifications analyzed, which involve only different assumptions on the exogenous shocks driving the model, help in solving the puzzle. In the next section, we modify the model so that it can replicate not only the area below the RER spectrum (the standard deviation) but also its

shape (the persistence), i.e., we introduce an extended model that can solve the puzzle.

6. Extensions to the Baseline Model

In this section, we will add two ingredients to the baseline model that will help us solve the puzzle while still replicating the variance of RER. First, we consider adjustment costs in the use of intermediate imported inputs for the production of the final good, and second, we analyze the role of the portfolio adjustment costs.

6.1. Adjustment Costs in the Use of Intermediate Imported Inputs

The first additional ingredient will be to assume adjustment costs in the use of intermediate imported inputs for the production of the final good. As we will see below, this feature will allow us to consider low short-run elasticities of substitution between intermediate goods with high long-run ones. The empirical literature that estimates trade elasticities argues that, due to the slow adjustment of quantities in response to prices, elasticities of substitution differ in the short run and in the long run. For instance, Hooper, Johnson and Marquez (2000) estimate import and export equations for the G-7 countries and show that the long-run elasticities are much higher than the short-run ones.

In order to include input adjustment costs, we follow Erceg, Guerrieri and Gust (2006). Hence, the production function is now:

$$Y_t = \left[\omega^{\frac{1}{\theta}} Y_{H,t}^{\frac{\theta-1}{\theta}} + (1-\omega)^{\frac{1}{\theta}} (\varphi_t Y_{F,t})^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}}.$$

As we will see below, θ is now the elasticity of substitution between home-country and foreign-country intermediate goods in the long-run. The input adjustment, φ_t , follows the following functional form:

$$\varphi_t = \left[1 - \frac{\iota}{2} \left(\frac{Y_{F,t}/Y_{H,t}}{Y_{F,t-1}/Y_{H,t-1}} - 1 \right)^2 \right]. \quad (18)$$

With this specification, changing the ratio of home-country to foreign-country intermediate goods reduces the efficiency of the imported intermediate input.⁸ There are no direct available

⁸Obstfeld and Rogoff (2000) analyze the role of transportation costs (in the form of iceberg costs) in explaining several puzzles of international macroeconomics. However, they conclude that this type of friction alone cannot solve the puzzle of high volatility of real exchange rates, which they label “the exchange rate disconnect puzzle.”

estimates of the cost function (18). Hence, how can we interpret the ι parameter and the cost function? Suppose that the ratio $\frac{Y_{F,t}/Y_{H,t}}{Y_{F,t-1}/Y_{H,t-1}}$ deviates by 1 percent from its steady-state value at time t . Then the value of $\varphi_t = 1 - \frac{\iota}{2}(0.01)^2$. If $\iota = 200$, then $\varphi_t = 0.99$ and the intermediate foreign-country input is 1 percent less efficient. Because the foreign-country intermediate input is only 10 percent of home-country final production, this means that home-country output will be 0.1 percent smaller than without the presence of this cost. By the same reasoning, if $\iota = 2000$ the cost in terms of output would be 1 percent.

The input adjustment cost function depends on variables dated at $t - 1$, and hence this introduces an intertemporal dimension to the final good producers' profit maximization problem. We use the domestic households' stochastic discount factor to discount future profits. The representative final good producer in the home country solves the following problem:

$$\max_{Y_{t+k}, Y_{H,t+k}, Y_{F,t+k}} E_t \sum_{k=0}^{\infty} \beta^k \Lambda_{t+k} (P_{t+k} Y_{t+k} - P_{H,t+k} Y_{H,t+k} - P_{F,t+k} Y_{F,t+k})$$

subject to the production function (2) and the input adjustment cost function (18). Note that $\beta^k \Lambda_{t+k} = \beta^k (\lambda_{t+k}/P_{t+k})/(\lambda_t/P_t)$ is the stochastic discount factor. The first order conditions of the problem are given by:

$$P_t \frac{\partial Y_t}{\partial Y_{H,t}} + \beta E_t \left(\Lambda_{t+1} P_{t+1} \frac{\partial Y_{t+1}}{\partial Y_{H,t}} \right) = P_{H,t} \quad (19)$$

and

$$P_t \frac{\partial Y_t}{\partial Y_{F,t}} + \beta E_t \left(\Lambda_{t+1} P_{t+1} \frac{\partial Y_{t+1}}{\partial Y_{F,t}} \right) = P_{F,t}. \quad (20)$$

Using the previous functional forms we obtain the following expressions:

$$\begin{aligned} \frac{P_{H,t}}{P_t} = & Y_t^{\frac{1}{\theta}} \left[\omega^{\frac{1}{\theta}} Y_{H,t}^{\frac{-1}{\theta}} + (1 - \omega)^{\frac{1}{\theta}} (Y_{F,t})^{\frac{\theta-1}{\theta}} (\varphi_t)^{\frac{-1}{\theta}} \iota \left(\frac{Y_{F,t}/Y_{H,t}}{Y_{F,t-1}/Y_{H,t-1}} - 1 \right) \frac{Y_{F,t}/(Y_{H,t})^2}{Y_{F,t-1}/Y_{H,t-1}} \right] \\ & - \beta E_t \left\{ \left(\frac{\lambda_{t+1}}{\lambda_t} \right) Y_{t+1}^{\frac{1}{\theta}} \left[(1 - \omega)^{\frac{1}{\theta}} (Y_{F,t+1})^{\frac{\theta-1}{\theta}} (\varphi_{t+1})^{\frac{-1}{\theta}} \iota \left(\frac{Y_{F,t+1}/Y_{H,t+1}}{Y_{F,t}/Y_{H,t}} - 1 \right) \frac{Y_{F,t+1}/Y_{H,t+1}}{Y_{F,t}} \right] \right\} \end{aligned} \quad (21)$$

and

$$\begin{aligned} \frac{P_{F,t}}{P_t} = & Y_t^{\frac{1}{\theta}} \left[(1 - \omega)^{\frac{1}{\theta}} (\varphi_t Y_{F,t})^{\frac{-1}{\theta}} \left(\varphi_t - \iota Y_{F,t} \left(\frac{Y_{F,t}/Y_{H,t}}{Y_{F,t-1}/Y_{H,t-1}} - 1 \right) \frac{1/Y_{H,t}}{Y_{F,t-1}/Y_{H,t-1}} \right) \right] \\ & + \beta P_t E_t \left\{ \left(\frac{\lambda_{t+1}}{\lambda_t} \right) Y_{t+1}^{\frac{1}{\theta}} \left\{ (1 - \omega)^{\frac{1}{\theta}} (\varphi_{t+1} Y_{F,t+1})^{\frac{-1}{\theta}} \left[Y_{F,t+1} \iota \left(\frac{Y_{F,t+1}/Y_{H,t+1}}{Y_{F,t}/Y_{H,t}} - 1 \right) \frac{Y_{F,t+1}/Y_{H,t+1}}{(Y_{F,t})^2/Y_{H,t}} \right] \right\} \right\}. \end{aligned} \quad (22)$$

Foreign-country intermediate goods producers face the same problem, which we do not describe because of space considerations. We calibrate the parameters as described in section 5.1 except the long-run elasticity of substitution between intermediate goods is now set to a value of 2. This value is somewhat higher than that typically used in open economy macro models (Chari, Kehoe and McGrattan, 2002; and Erceg, Guerrieri and Gust, 2006 use $\theta = 1.5$). We now vary the degree of the cost, ι , and look at the implications for the model. The results are reported in Table 4.

Table 4: The Role of Input Adjustment Costs						
	Standard Deviation			Frequency of RER		
	RER	Output Growth		Low	BC	High
U.S. Data	10.56	0.8		70.0	24.3	5.7
$\iota = 0$	2.5	0.88		89.3	8.2	2.5
$\iota = 125$	4.98	0.86		84.8	11.7	3.5
$\iota = 250$	8.10	0.81		83.3	12.8	3.9
$\iota = 330$	10.66	0.77		82.5	13.1	4.0
$\iota = 500$	20.64	0.61		81.7	13.9	4.3

Note: RER denotes the log of the RER. Output is real GDP. Growth rates are computed taking the first differences of the logs.

Introducing an input adjustment cost has important implications for the RER. As expected, when $\iota = 0$ the model does not generate enough volatility of the RER and the fraction of volatility assigned to BC- and high-frequency fluctuations is still too small. As the cost increases, the volatility of the RER and the fraction of volatility assigned to BC- and high-frequency fluctuations increases. A value of $\iota = 330$ allows the model to match the volatility of the RER and of output growth in the data and also improves the fit to the shape of the spectrum. Yet, too much weight is still placed on the low-frequency movements (82.5 percent of fluctuations at low frequencies in the model versus 70 percent in the data for $\iota = 330$), i.e., the “excess persistence of RER” puzzle

is not fully solved. As ι grows, the model generates too much RER volatility but the share of variance assigned to low-frequency fluctuations remains higher than in the data. Hence, input adjustment costs can dramatically help to replicate RER volatility, even for large values of θ , but not to solve the puzzle completely. In what follows, we explain why input adjustment costs can help generate more RER volatility in the model. In the next section, we analyze how the interaction between input and portfolio adjustment costs can help in fully matching the spectrum of the RER.

In Figures 3-5 we plot the IRFs to a home-country TFP shock for different values of ι to understand how this parameter shapes the behavior of the RER. When $\iota = 0$, standard results in the IRBC literature apply (see Backus, Kehoe and Kydland, 1992). When a TFP shock hits the home-country economy, we get the usual effect from an IRBC model: output, consumption, investment and hours worked increase in the home country, while in the foreign country, output, investment and hours worked decline, and consumption increases. As output expands, the demand for home- and foreign-country intermediate goods increases, although it increases more for home-country intermediate goods. In the foreign country, investment declines because foreign-country households buy home-country bonds to invest in the home country, with higher productivity, instead of foreign-country capital. Hours decline because of the associated decline of the marginal product of capital. Right away, foreign-country households increase their consumption because of an income effect related to future spillovers from the home-country technological improvement and higher returns on their bond holdings in the home country. In addition, this income effect leads the foreign-country households to supply even less labor. As output decreases in the foreign country, the demand for home- and foreign-country intermediate goods also decreases.

As the literature has pointed out, the reaction of the RER is not too large but very persistent. The peak of the IRF happens after 20 quarters and the half-life is reached after more than 50 quarters. This highly persistent response of the RER is related to the “excess persistence of RER” puzzle: regardless of the value of θ , far too much weight is placed at low-frequency movements. As a result of the decline in the price of home-country intermediate goods, and the increase in both the price and the quantity of foreign-country intermediate goods, a trade deficit for the home country emerges. This implies that variable D_t , which denotes the holding of bonds by the home-country household, becomes negative (see equation 17). The variable D_t also denotes the net foreign asset position (NFA) of the home country. Thus, when a TFP shock hits the home

country, its NFA position becomes negative in order to finance higher investment.

Introducing input adjustment costs leads to important changes in the behavior of some variables. The larger ι , the closer Y_H and Y_F need to move in order to avoid reducing the efficiency of the foreign-country intermediate input. Without input adjustment costs Y_H increases more than Y_F , but the presence of the costs leads to a reduction in this difference. Something similar happens to Y_H^* and Y_F^* . As a result, the home-country demand for home-country intermediate goods increases less and the demand for foreign-country intermediate inputs increases more (when compared with the case of $\iota = 0$). This implies that, the larger ι , the larger is the trade deficit that the home country runs (or the worse is its NFA position). This is key to inducing more RER volatility. Why is this the case? An inspection of the risk-sharing condition across countries gives us the answer. The linearized risk-sharing equation of the model reads as follows:

$$\begin{aligned}\widehat{rer}_t &= E_t \left[\widehat{rer}_{t+1} + (\hat{\lambda}_{t+1} - \hat{\lambda}_t) - (\hat{\lambda}_{t+1}^* - \hat{\lambda}_t^*) \right] - \frac{\phi}{\beta} d_t \\ &= E_t \sum_{i=0}^{\infty} \left\{ [(\hat{\lambda}_{t+i+1} - \hat{\lambda}_{t+i}) - (\hat{\lambda}_{t+i+1}^* - \hat{\lambda}_{t+i}^*)] - \frac{\phi}{\beta} d_{t+i} \right\}\end{aligned}\quad (23)$$

where lower case variables with a hat (such as \widehat{rer}_t) denote log-deviations from steady-state values and lower case variables (in this case, just d_t) denote deviations from steady-state values (this is the case because in the steady state, $D = 0$). Leaving aside changes in the relative marginal utilities of consumption, equation (23) links movements in the RER with the expected discounted sum of movements in the NFA position. Hence, the larger the input adjustment costs, the larger the NFA deterioration and the larger the depreciation of the RER. In fact, the NFA movements will mostly drive the behavior of the RER because households dislike changes in the marginal utility of consumption.

Therefore, there are two channels through which the introduction of import adjustment costs increases the volatility of the RER in the short run in the model. First, the adjustment costs make relative quantities less sensitive to changes in relative prices, and this increases the volatility of the terms of trade and the RER. But at the same time, the volatility of net exports and net foreign assets increases, which feeds back into higher exchange rate volatility through equation (23). The large effects of input adjustment costs on RER fluctuations are important in the short run, when the costs plays a role. In the long run, these adjustment costs dissipate and because of a large θ , RER fluctuations are dramatically reduced. Hence, the adjustment costs of imported inputs

and the large long-run elasticity of substitution allow us to increase the size of RER fluctuations in the short run (because of large movements of the NFA in the short-run due to the cost) and reduce them in the long-run (because of large θ).

An alternative way to understand the mechanism is to analyze how the relationship between relative quantities of intermediate inputs and their relative prices changes across time once input adjustment costs are introduced. In Figure 6, we compute a “pseudo-elasticity” of substitution when input adjustment costs are introduced as a function of time. In the baseline model, the elasticity of substitution between home and foreign goods is constant and equal to:

$$\frac{\partial \log(Y_{H,t}/Y_{F,t})}{\partial \log(P_{H,t}/P_{F,t})} = -\theta.$$

Computing the elasticities of substitution is not straightforward in the model with input adjustment costs (see equations 21 and 22). As a short cut, we compute the ratio:

$$\theta_k^{pseudo} = -\frac{\hat{y}_{H,t+k} - \hat{y}_{F,t+k}}{\hat{p}_{H,t+k} - \hat{p}_{F,t+k}}$$

at several time horizons k based on the IRFs to a home-country TFP shock presented in Figure 6. The $\iota = 0$ case trivially delivers a constant elasticity of substitution of $\theta = 2$. The introduction of input adjustment costs delivers a short-term elasticity that is very low and close to zero (the limiting case of zero would be a Leontieff production function for the final good). Over time, the elasticity slowly increases to its long-run value of 2. Thus, introducing input adjustment costs allows us to have low short-run elasticities (that increase RER volatility at BC frequencies) with higher long-run elasticities (that lower RER volatility at lower frequencies). This mechanism goes a long way towards getting the shape of the spectrum right, but it does not fully solve the “excess persistence of RER” puzzle.

At this point, it is relevant to highlight two issues. First, this feedback channel (between larger NFA volatility and larger RER short-run depreciation because of input adjustment costs) would not operate under complete markets. Hence, incomplete markets is a crucial part of the story. Second, portfolio adjustment costs, ϕ , interact with intermediate input adjustment costs and NFA movements to generate RER fluctuations. The next subsection shows how a slightly larger value of ϕ can help to solve the “excess persistence of RER” puzzle without affecting the capability of the model to generate enough RER volatility.

6.2. The Role of Portfolio Adjustment Costs

As we have shown in the previous subsection, there are limits to how much input adjustment costs help to solve the “excess persistence of RER” puzzle. Here, we show how combining those with slightly larger portfolio adjustment costs than reported in the benchmark calibration helps solve the puzzle. Introducing input adjustment costs increases the volatility of the NFA position and, through the risk-sharing condition, of the RER. Higher portfolio adjustment costs help the model match the persistence of the RER in the data by reducing the persistence of the long-run response of the RER to shocks. Thus, the larger portfolio adjustment costs, the lower the fraction of the variance of the RER that the model assigns to low-frequency fluctuations. This modified model is successful at replicating the shape of the spectrum without affecting the model’s capability to explain the area below it.

In the baseline version of the model, we have chosen a value of $\phi = 0.01$, which is small enough to ensure that the model goes back to the initial steady state and the NFA position is stationary in the long run. In Table 5, we show the results of setting $\phi = 0.05$, while keeping $\theta = 2$, for different values of ι .

Table 5: The Role of Input and Portfolio Adjustment Costs					
	Standard Deviation		Frequency of RER		
	RER	Output Growth	Low	BC	High
U.S. Data	10.56	0.8	70.0	24.3	5.7
$\iota = 0$	2.43	0.84	88.7	8.7	2.6
$\iota = 62.5$	5.28	0.81	76.0	18.8	5.2
$\iota = 125$	11.13	0.70	72.9	21.1	6.0

Note: RER denotes the log of the RER. Output is real GDP. Growth rates are computed taking the first differences of the logs.

Comparing Tables 4 and 5 shows that increasing portfolio adjustment costs in the baseline model (without input adjustment costs, i.e., $\iota = 0$) does not change the predictions of the model. Interestingly, the interaction of (i) a large long-run elasticity of substitution, $\theta = 2$, (ii) input adjustment costs, and (iii) larger portfolio adjustment costs leads the model to replicate both the standard deviation and the persistence of the RER (the area and the shape of the spectrum). In particular, when we set $\iota = 125$ and $\phi = 0.05$, the model explains almost perfectly the volatility

and the persistence of the RER. We plot the spectrum of the data and the model in Figure 7. The fit is remarkably good.

In order to understand why the combination helps replicate both the area under the spectrum and its shape, we present the IRFs to a domestic TFP shock when $\theta = 2$ and $\iota = 125$, while varying ϕ from 0.01 to 0.05 in Figures 8-10. We have already discussed what happens when $\theta = 2$, $\iota = 125$, and $\phi = 0.01$. Hence, we focus on the effects of moving ϕ from 0.01 to 0.05. If there were no input adjustment costs, increasing ϕ mainly affects the dynamics of net exports and bonds, without much spillover to the rest of the variables (this is a standard result of the Heathcote-Perri-BKK models and therefore not shown but available upon request).

As we discussed in the previous section, increasing ι leads to a bigger short-term response of the NFA and the RER to a shock. Slightly larger portfolio adjustment costs ($\phi = 0.05$) amplify the RER depreciation due to the portfolio adjustment cost term in the risk-sharing condition. Hence, the initial response of the RER when $\iota = 125$ and $\phi = 0.05$ is larger than in Figure 5, where we used $\iota = 330$ and $\phi = 0.01$. At the same time, larger portfolio adjustment costs imply a faster return to the balanced growth path. The faster decline of D_t to balanced growth dampens the long-term response of the NFA and the RER to a shock. This explains why slightly larger portfolio adjustment costs allow us to match the RER volatility with smaller input adjustment costs and, at the same time, solve the “excess persistence of RER” puzzle. Why? As explained above, the RER and the NFA are positively related through the parameter PHI. Hence, we only need a value of $\iota = 125$ to replicate the variance of the RER. If $\iota = 125$, a deviation in the ratio of inputs of 1 percent implies that output will be 0.0625 percent smaller than without the presence of the cost. A value of $\phi = 0.05$ implies that an increase of NFA of 1 percent costs the home-country household 0.0025 percent units of consumption. We view these numbers as being fairly small, and yet they have important implications for the dynamics of the model.

7. Concluding Remarks

In this paper, we have shown that most of the volatility of the RER can be assigned to low frequencies (below BC frequencies). Therefore, it makes sense to ask if IRBC models can replicate the spectrum of the RER when no filter is applied to either the actual data or the simulated data coming from the model. Filtering the RER implies removing low-frequency movements and eliminating most of the fluctuations of the RER. When matching the spectrum of the RER the

challenge is twofold. First, we need to match its area (the volatility of the RER) and, second, its shape (the share of variance assigned to different frequencies). In Section 4, we have presented a standard version of a two-country, two-good IRBC model in the spirit of Heathcote and Perri (2002), which includes cointegrated TFP shocks across countries as in Rabanal, Rubio-Ramírez and Tuesta (2011). This baseline model is capable of explaining the volatility of the RER (the area below the spectrum), but places too much weight on low-frequency movements (it cannot explain the shape of the spectrum). We call this shortcoming of the model the “excess persistence of RER” puzzle.

In Section 5 we study if modeling TFP shocks as stationary processes or adding IST shocks to our baseline model help to solve the puzzle. We conclude that they do not. In Section 6 we try a new venue. We extend the baseline model to allow for adjustment costs in the use of intermediate inputs as in Erceg, Guerrieri and Gust (2006). We conclude that what is needed to solve the puzzle while still explaining the volatility of the RER is the interaction of three ingredients: The first ingredient is a large steady-state elasticity of substitution ($\theta = 2$). The second ingredient is the introduction of adjustment costs in intermediate inputs, which help lower the implied elasticity of substitution in the short run and hence increase RER volatility. Our calibrated value is $\iota = 125$, which means that a deviation in the ratio of inputs of 1 percent implies that output will be 0.0625 percent smaller. And the third ingredient is higher portfolio adjustment costs, in the range of 5 basis points, which increase the volatility of the RER in the short run but lower the persistence at long horizons.

References

- [1] Adolfson, M., S. Laseen, J. Lindé and M. Villani, 2007. Bayesian estimation of an open economy DSGE model with incomplete pass-through. *Journal of International Economics* 72, 481-511.
- [2] Backus, D., P. Kehoe and F. Kydland, 1992, International Business Cycles, *Journal of Political Economy*, Vol. 100, No. 41, 745 -775.
- [3] Backus, D., P. Kehoe and F. Kydland, 1994, Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?, *American Economic Review*, Vol. 84, No. 1, pp. 84-103.
- [4] Baxter, M., 2011, International Risk Sharing in the Short Run and in the Long Run, NBER Working Paper 16789.
- [5] Baxter, M. and M. Crucini, 1995, Business Cycles and the Asset Structure of Foreign Trade, *International Economic Review*, Vol. 36, No. 4, 821-854.
- [6] Baxter, M. and R. King, 1999, Measuring Business Cycles: Approximate Band-Pass Filters for Macroeconomic Time Series. *Review of Economics and Statistics*, 81(4), November 1999: 575-93.
- [7] Benigno, G., 2004. Real exchange rate persistence and monetary policy rules. *Journal of Monetary Economics* 51, 473-502.
- [8] Bergin, P. and R. Feenstra, 2001, Pricing to Market, Staggered Contracts, and Real Exchange Rate Persistence. *Journal of International Economics* 54 (2), 333-359.
- [9] Bouakez, H, 2005, Nominal rigidity, desired markup variations and real exchange rate persistence. *Journal of International Economics*, Vol. 61, 49-74.
- [10] Chari, V.V. , P. Kehoe and E. McGrattan, 2002, Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates? *Review of Economic Studies*, Vol. 69, 533-563.
- [11] Cheung, Y., and K. Lai, 2000, On the Purchasing Power Parity, *Journal of International Economics*, Vol. 52, 321-330.

- [12] Cheung, Y., M. Chinn and A. Garcia-Pascual, 2005, Empirical exchange rate models of the nineties: Are any fit to survive?, *Journal of International Money and Finance*, Vol. 24(7), 1150-1175.
- [13] Christiano, L. and T. Fitzgerald 2003, The Band Pass Filter. *International Economic Review*, Vol. 44, No. 2, 435-465.
- [14] Clarida, R. and J. Galí, 1994, Sources of real exchange-rate fluctuations: How important are nominal shocks?, *Carnegie-Rochester Conference Series on Public Policy*, Vol. 41(1), 1-56.
- [15] Corsetti, G., L. Dedola and S. Leduc, 2008, International Risk Sharing and the Transmission of Productivity Shocks, *Review of Economic Studies*, Vol. 75, 443-473.
- [16] Corsetti, G., L. Dedola and F. Viani, 2011, The International Risk-Sharing Puzzle is at Business Cycle and Lower Frequency, *CEPR Discussion Papers* 8355.
- [17] Eichenbaum, M. and C. Evans. 1995, Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates. *Quarterly Journal of Economics*, 110(4): 975–1009.
- [18] Engel, C. and K. West, 2005, Exchange Rates and Fundamentals, *Journal of Political Economy*, Vol. 107, No. 3, 485-517.
- [19] Erceg, C., L. Guerrieri and C. Gust, 2006, SIGMA: A New Open Economy Model for Policy Analysis, *International Journal of Central Banking*, vol. 2(1), March.
- [20] Faust, J. and J. Rogers, 2003, Monetary Policy's Role in Exchange Rate Behavior, *Journal of Monetary Economics*, vol. 50 (November 2003), pp. 1403-1424.
- [21] Hamilton, J., 1994, *Time Series Analysis*, Princeton University Press.
- [22] Heathcote, J., and F. Perri, 2002, Financial Autarky and International Business Cycles, *Journal of Monetary Economics*, Vol. 49, 601-627.
- [23] Hooper, P., Johnson, K. and J. Marquez, 2000, Trade Elasticities for the G-7 Countries. *Princeton Studies in International Economics*, no. 87, August.
- [24] Huizinga, J., 1987, An Empirical Investigation of the Long-Run Behavior of Real Exchange Rates. *Carnegie-Rochester Conference Series on Public Policy*, Vol. 27, pp. 149-214.

- [25] Johansen, S., 1991, Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models, *Econometrica* 59, 1551-1580.
- [26] King, R., C., Plosser and S., Rebelo, 1988, Production, Growth and the Business Cycle, *Journal of Monetary Economics*, 21, 195-232.
- [27] Mandelman, F., Rabanal, P., Rubio-Ramírez, J.F. and D. Vilán, 2011, Investment-Specific Technology Shocks and International Business Cycles: An Assessment, *Review of Economic Dynamics* 14, pp. 136-155.
- [28] Meese, R. and K. Rogoff, 1983, Empirical exchange rate models of the seventies: do they fit out of sample?, *Journal of International Economics* 14, pp. 3–24.
- [29] Muntaz, H. and P. Surico, 2009, The Transmission of International Shocks: A Factor-Augmented VAR Approach, *Journal of Money, Credit & Banking* 41, pp. 71-100.
- [30] Murray, C. and D. Papell, 2002, The Purchasing Power Parity Persistence Paradigm. *Journal of International Economics*, 56(1): 1–19.
- [31] Nason J. and J. Rogers, 2008, Exchange Rates and Fundamentals: A Generalization. Federal Reserve Bank of Atlanta Working Paper 2008-16.
- [32] Obstfeld, M, and K. Rogoff, 2000, The Six Major Puzzles in International Macroeconomics: Is There A Common Cause? NBER Working Paper 7777.
- [33] Rabanal, P., J. Rubio-Ramírez, and V. Tuesta, 2011, Cointegrated TFP Processes and International Business Cycles, *Journal of Monetary Economics*, Vol. 58, 156-171.
- [34] Rabanal, P. and V. Tuesta, 2010, Euro-Dollar Real Exchange Rate Dynamics in an Estimated Two-Country Model: An Assessment, *Journal of Economic Dynamics and Control*, Vol. 34, pp. 780-797.
- [35] Rogoff, K., 1996, The Purchasing Power Parity Puzzle, *Journal of Economic Literature*, vol. 34, 647–68.
- [36] Sarno, L., 2003, Nonlinear Exchange Rate Models: A Selective Overview, IMF Working Paper 03/111.

- [37] Steinsson, J., 2008, The Dynamic Behavior of the Real Exchange Rate in Sticky Price Models, American Economic Review, Vol. 98, pp. 519-533.
- [38] Stockman, A., L. Tesar, 1995, Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Comovements, American Economic Review, Vol 85(1), 165-185.

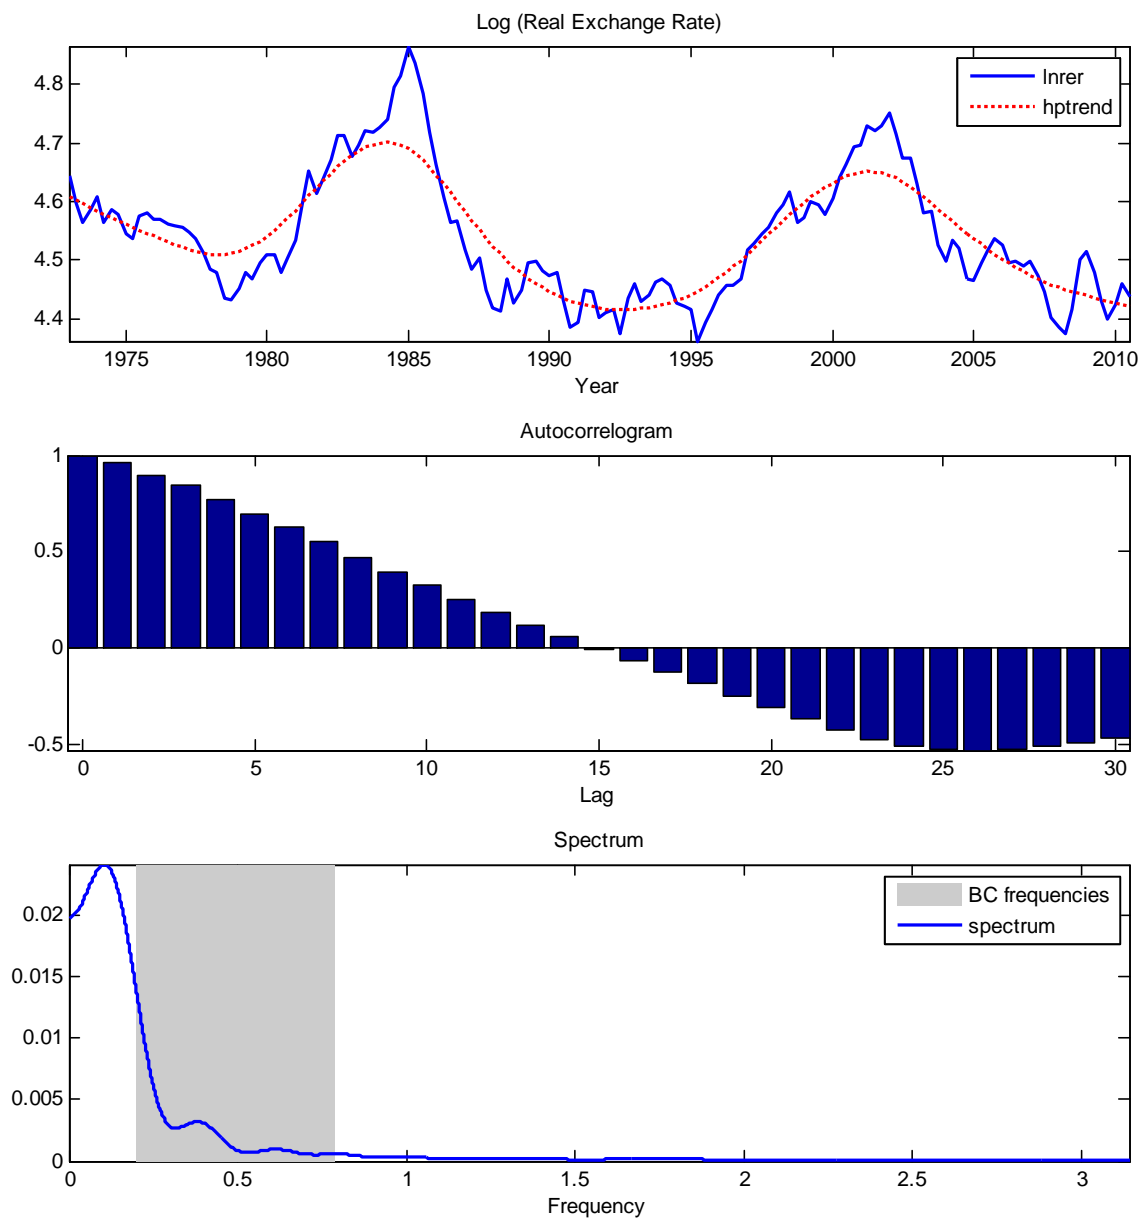


Figure 1: Log RER, autocorrelation function, and spectral density of the U.S. dollar.




Figure 2: Comparison between the model and the data.

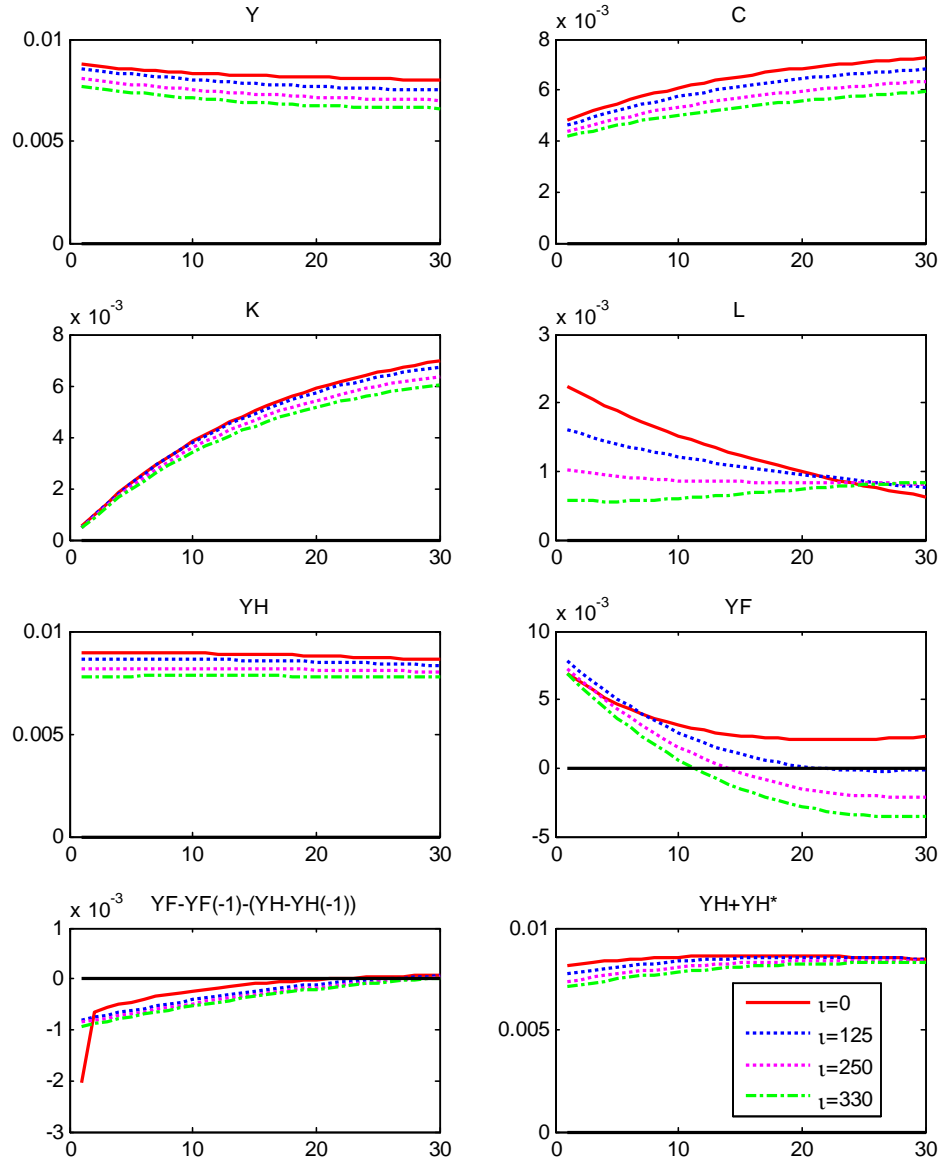


Figure 3: IRF to a home TFP shock when ι changes.

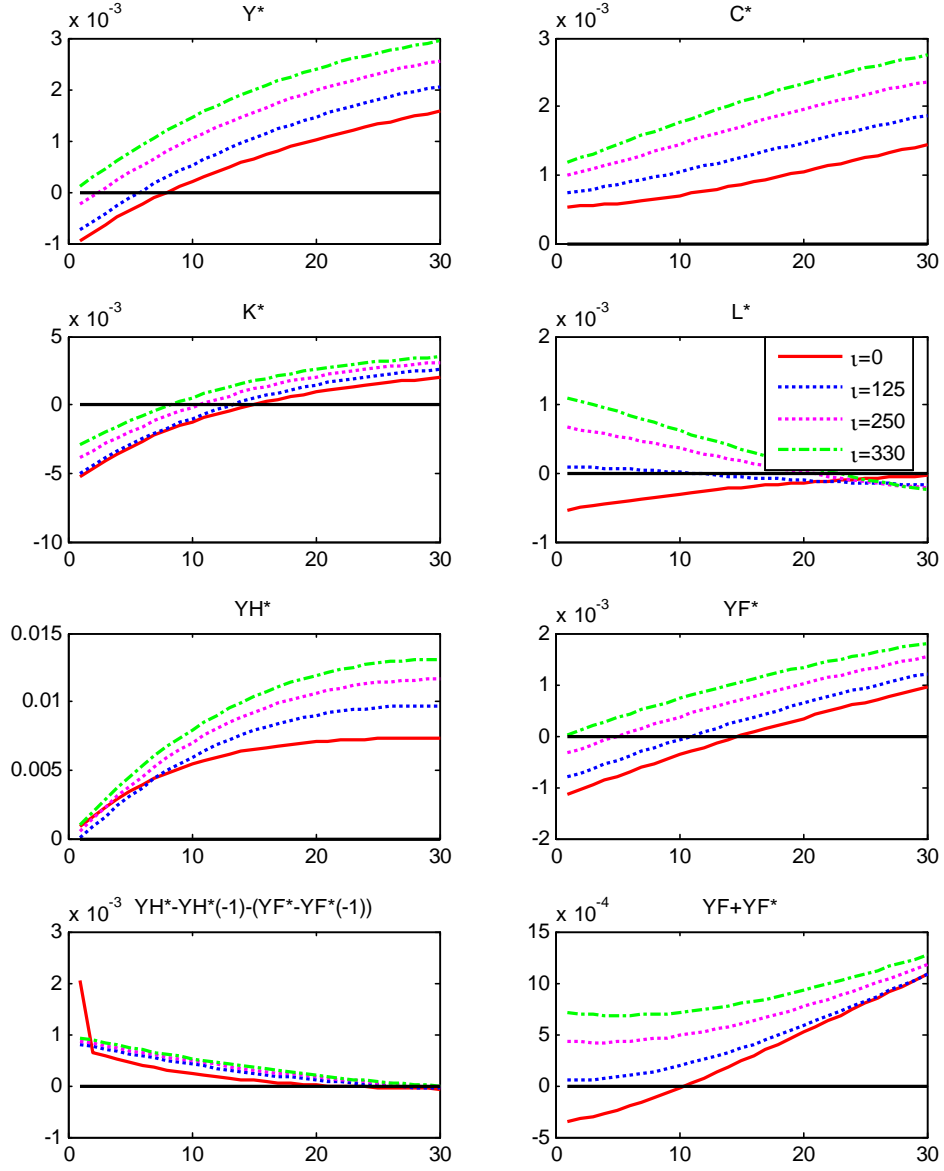


Figure 4: IRF to a home TFP shock when ι changes.

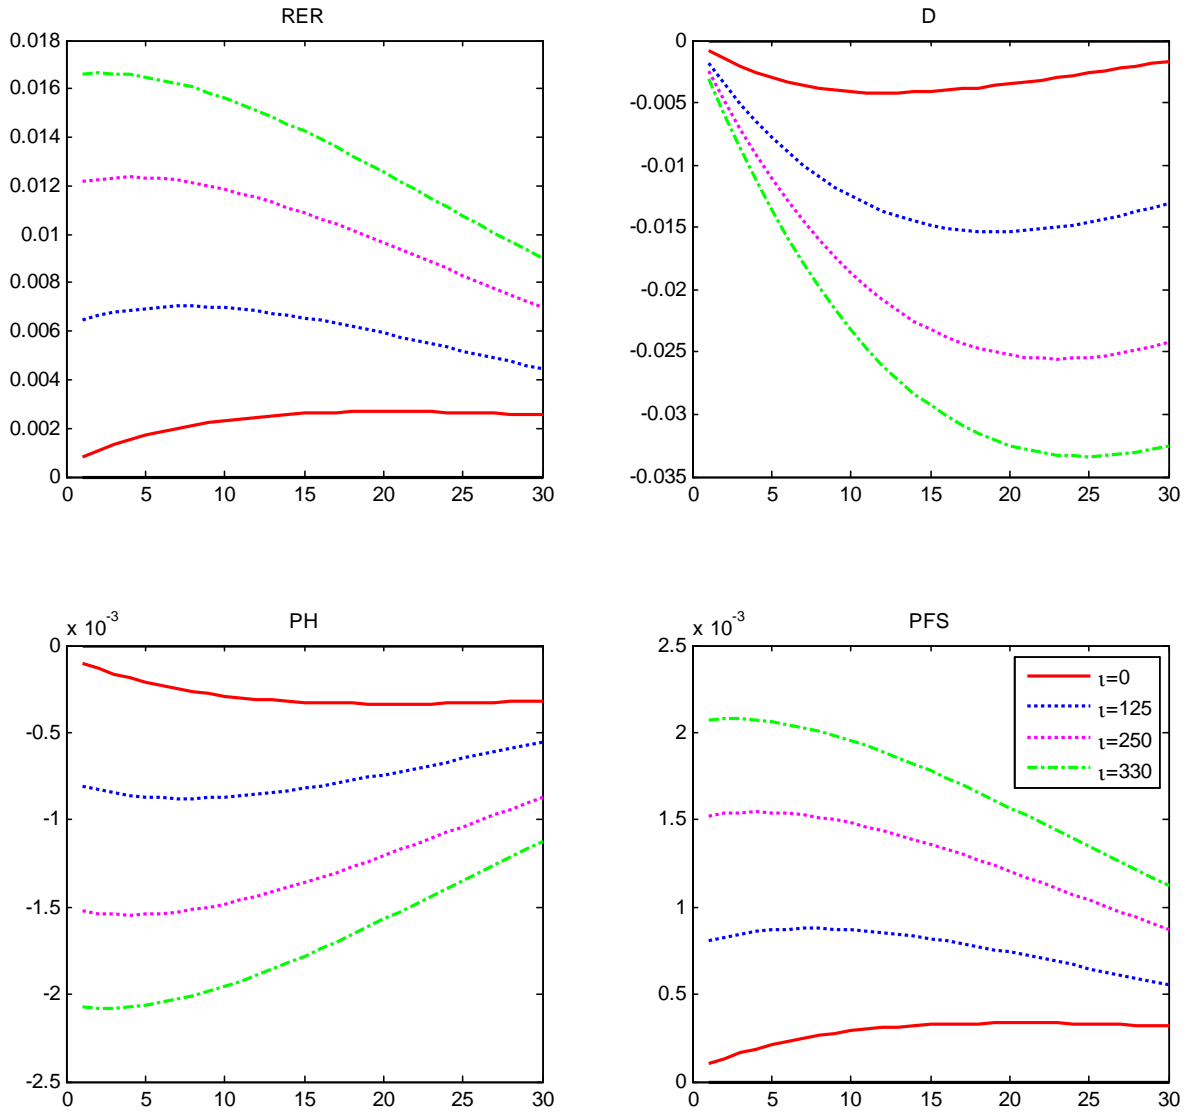


Figure 5: IRF to a home TFP shock when ι changes.




Figure 6: Implied elasticity when ι changes.




Figure 7: Comparison between the extended model and the data.

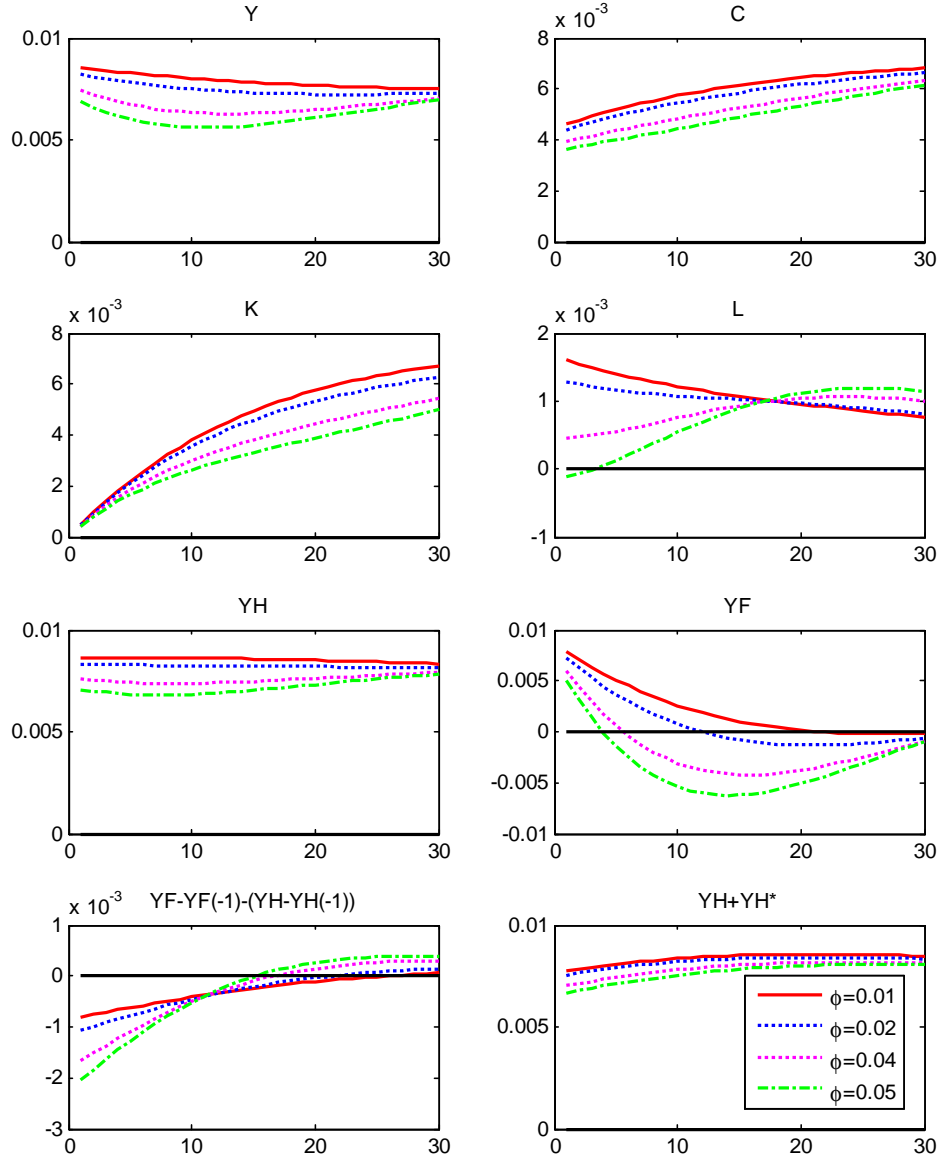


Figure 8: IRF to a home TFP shock when ϕ changes and $\iota = 125$.

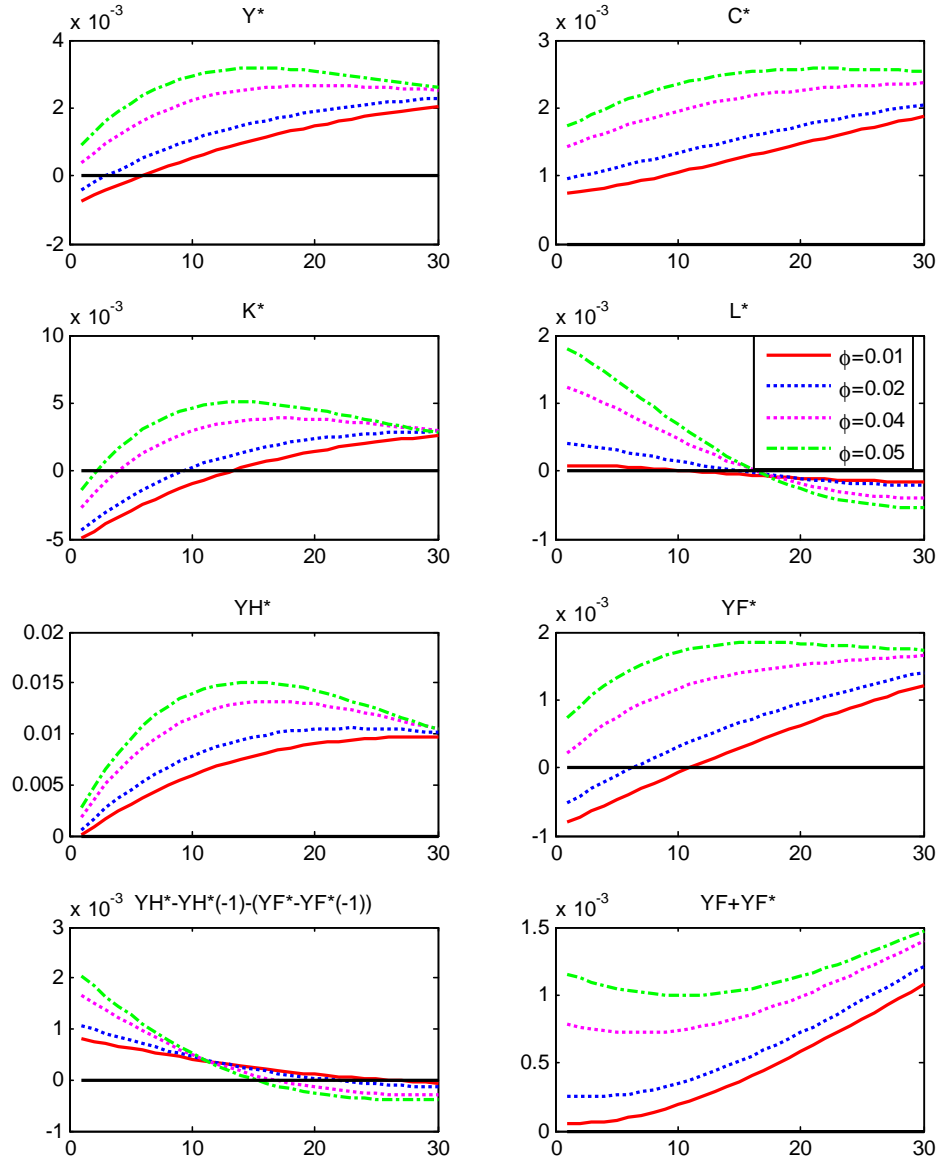


Figure 9: IRF to a home TFP shock when ϕ changes and $\iota = 125$.

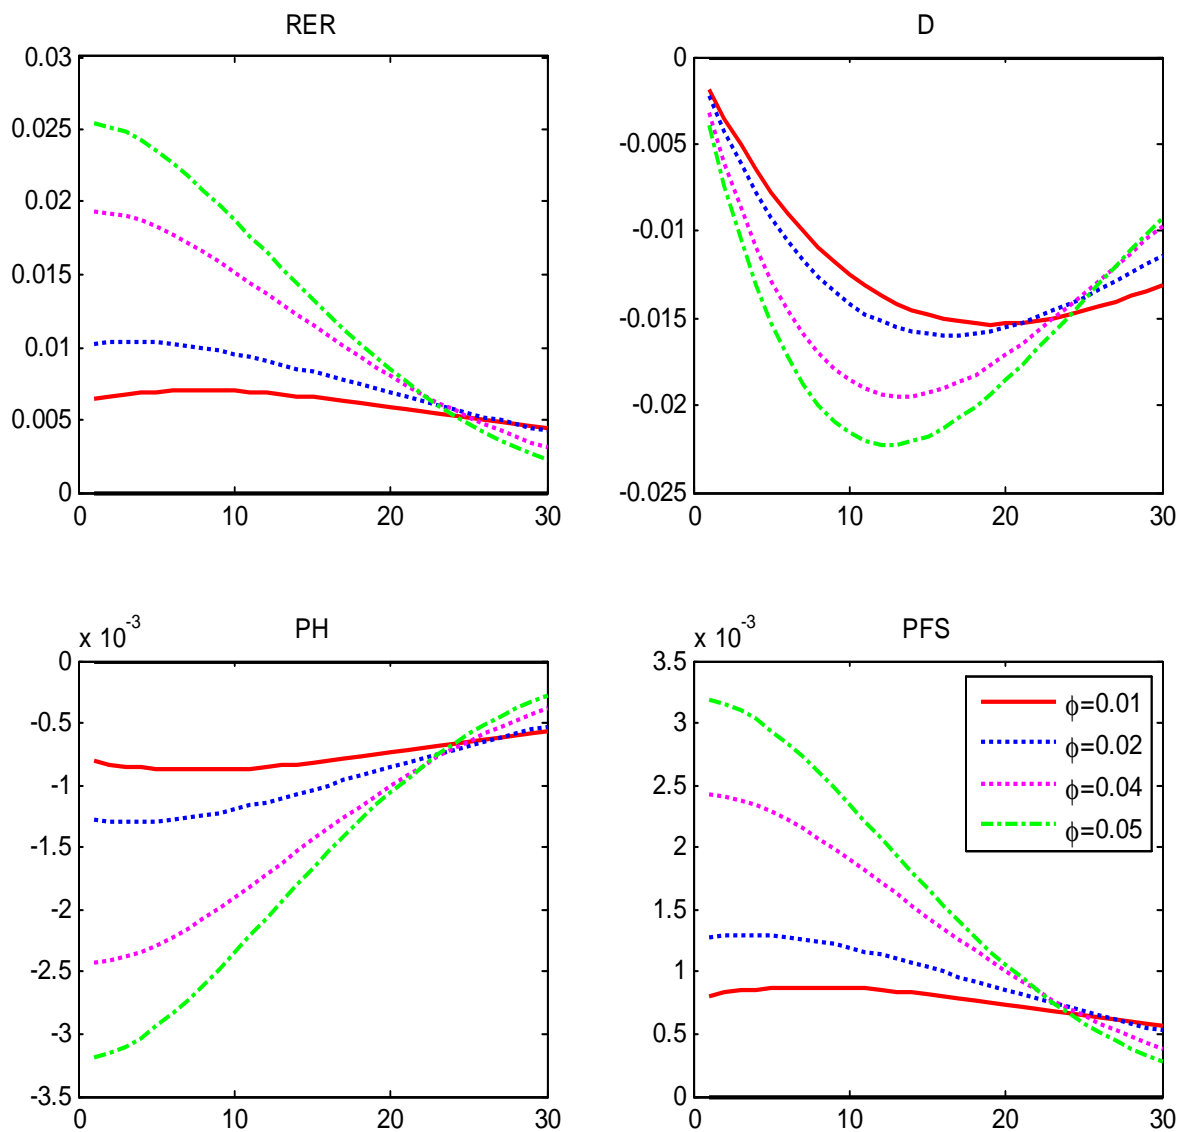


Figure 10: IRF to a home TFP shock when ϕ changes and $\iota = 125$.