CHAPTER 1

Macro Models without Frictions

This is the first of two chapters that examine the exchange-rate implications of macroeconomic models. My aim is not to survey all the macro models of exchange-rate determination, but rather to provide a theoretical overview of how exchange rates are linked to macro variables in environments that are familiar to students of macroeconomics. This overview serves two purposes. First, it highlights the degree to which the exchange-rate implications of widely used macro models accord with the empirical characteristics of exchange-rate behavior. Second, it establishes a theoretical benchmark for judging the success of the new micro-based exchange-rate models presented in subsequent chapters.

The macro models we study have standard features. There are two countries, each populated by a large number of identical utility-maximizing households. In this chapter we study models where households have access to a rich array of financial assets. More specifically, we assume that they all have access to markets for a complete set of contingent claims. As a result, households are able to share risk completely. This feature has important implications for the behavior of exchange rates in both endowment economies (where output is exogenous) and in production economies (where output is determined optimally by firms). Another standard feature of the models we study concerns the role of money. Here we assume that households derive utility from holding real balances and that central banks have complete control of their national money supplies. This framework has a long tradition in the international macro literature. When combined with the implications of complete risk-sharing, it allows us to characterize the differences between the behavior of real and nominal exchange rates in a straightforward manner.

The final noteworthy feature of the models presented herein concerns the behavior of prices. Although "price-stickiness" plays an important role in many international macro models, in this chapter we focus on models in which all prices are fully flexible. In so doing, our analysis abstracts from the complications caused by the presence of frictions in both financial and product markets. The exchange-rate implications of these frictions are examined in Chapter 2.

1.1 Preliminaries

1.1.1 Definitions

The focus of our analysis is on the behavior of the nominal spot exchange rate, which we refer to as the spot rate. The spot rate, denoted by S, is defined as the home price of foreign currency. Throughout, we take the United States as the home country, so S identifies the price of foreign currency in U.S. dollars. According to this definition, an appreciation in the value of the dollar is represented by a *fall* in S because it corresponds to a fall in the dollar price of foreign currency. Conversely, a *rise* in S represents a depreciation in the value of the dollar. Defining spot rates in this way can be a source of confusion at first, but it turns out to be very convenient when considering the determination of spot rates from an asset-pricing perspective. For this reason it is the standard definition used in the international finance literature.

The spot exchange rate identifies the price at which currencies can be traded immediately. Forward rates, by contrast, identify the price at which currencies can be traded at some future date. The k-period forward rate at time t, $\mathcal{F}_t^k$, denotes the dollar price of foreign currency in a contract between two agents at time t for the exchange of dollars and foreign currency at time $t+k$. Foreign currency is said to be selling forward at a discount (premium) relative to the current spot rate when $S_t - \mathcal{F}_t^k$ is positive (negative). Obviously, spot and forward rates are equal when the maturity of the forward contract, k, equals zero.

Two relative prices play prominent roles in international macro models. The first is the terms of trade. In international finance the convention is to define the terms of trade, $\mathcal{T}$, as the relative price of imports in terms of exports:

$$\mathcal{T} = \frac{P^{\mathrm{M}}}{S\hat{P}^{\mathrm{X}}},$$

where P^{M} is the price U.S. consumers pay for imports and $\hat{P}^{\mathrm{X}}$ is the price foreign consumers pay for U.S. exports. (Hereafter, we use a hat, i.e., "ˆ", to denote foreign variables.) Since S is defined as the price of foreign currency in dollars, $S\hat{P}^{\mathrm{X}}$ identifies the dollar price foreign consumers pay for U.S. exports. Note that a rise (fall) in the relative price of U.S. exports, representing an improvement (deterioration) in the U.S. terms of trade, implies a fall (rise) in $\mathcal{T}$. Once again, this may seem unnecessarily confusing, but the international finance literature adopts this definition to simplify the relationship between the real exchange rate and the terms of trade.

The second important relative price is the real exchange rate. This is defined as the relative price of the basket of all the goods consumed by foreign households in terms of the price of the basket of all the goods consumed by U.S. households:

$$\mathcal{E} = \frac{S\hat{P}}{P},$$

where $\hat{P}$ is the foreign currency price of the foreign basket and P is the dollar price of the U.S. basket. Hence, P and $\hat{P}$ are the U.S. and foreign consumer price indices in local currency terms, and $S\hat{P}$ identifies the foreign price index in terms of dollars.

According to this definition, a depreciation (appreciation) in the real value of the U.S. dollar corresponds to a rise (fall) in $\mathcal{E}$ and represents an increase (decrease) in the price of foreign goods relative to U.S. goods.

1.1.2 Price Indices

The behavior of the real exchange rate plays a central role in macro models, so it is important to relate the behavior of the price indices, P and $\hat{P}$, to the prices of individual goods. For this purpose, macro models identify price indices relative to a particular form for the consumption basket based on either the Constant Elasticity of Substitution (CES) or Cobb-Douglas functions.

To illustrate, suppose there are only two goods available to U.S. consumers. Under the CES formulation, the consumption basket defined over the consumption of goods a and b is given by

$$C = \mathcal{C}(a, b) = \left(\lambda^{\frac{1}{\theta}} a^{\frac{\theta-1}{\theta}} + (1-\lambda)^{\frac{1}{\theta}}\, b^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}}, \tag{1.1}$$

where $\lambda \in (0, 1)$ and $\theta > 0$. This function aggregates the consumption of the two goods into a single index, C, from which households derive instantaneous utility, $U(C)$, for some concave utility function $U(.)$. The consumption-based price index, P, is identified as the minimum expenditure that buys one unit of the consumption index, C. In other words, P minimizes the expenditure $Z = aP^a + bP^b$, such that $\mathcal{C}(a, b) = 1$, given the prices of goods a and b, P^a and P^b.

The mechanics of solving this problem illustrate several properties of the consumption basket and price index, so they are worth reviewing. Choosing a and b to minimize Z such that $\mathcal{C}(a, b) = 1$ gives

$$\frac{b}{a} = \frac{1-\lambda}{\lambda}\left(\frac{P^b}{P^a}\right)^{-\theta}. \tag{1.2}$$

Thus, the relative demand for good b depends on its relative price, P^b/P^a, and the ratio of shares in the basket, $\frac{1-\lambda}{\lambda}$. Note, also, that θ identifies the elasticity of substitution between goods a and b.

Combining equation (1.2) with the definition of total expenditure, $Z = aP^a + bP^b$, gives us the total demand for each good:

$$a = \frac{\lambda(P^a)^{-\theta}}{\left(\lambda(P^a)^{1-\theta} + (1-\lambda)\left(P^b\right)^{1-\theta}\right)} Z$$

and (1.3)

$$b = \frac{(1-\lambda)(P^b)^{-\theta}}{\left(\lambda(P^a)^{1-\theta} + (1-\lambda)\left(P^b\right)^{1-\theta}\right)} Z.$$

Substituting these expressions into (1.1) and setting the result equal to one gives us the minimum necessary expenditure:

$$1=\left(\lambda^{\frac{1}{\theta}}\left[\frac{\lambda(P^a)^{-\theta}Z}{\lambda(P^a)^{1-\theta}+(1-\lambda)\left(P^b\right)^{1-\theta}}\right]^{\frac{\theta-1}{\theta}}\right.$$

$$\left.+\ (1-\lambda)^{\frac{1}{\theta}}\left[\frac{(1-\lambda)(P^b)^{-\theta}Z}{\lambda(P^a)^{1-\theta}+(1-\lambda)\left(P^b\right)^{1-\theta}}\right]^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}}.$$

Simplifying this equation and solving for Z produces the equation for the price index:

$$P=\left(\lambda(P^a)^{1-\theta}+(1-\lambda)\ (P^b)^{1-\theta}\right)^{\frac{1}{1-\theta}}. \tag{1.4}$$

By definition, an expenditure of Z purchases Z/P units of the consumption index C, so we can use (1.4) to rewrite the expressions in (1.3) as

$$a=\lambda\left(\frac{P^a}{P}\right)^{-\theta}C \quad \text{and} \quad b=(1-\lambda)\left(\frac{P^b}{P}\right)^{-\theta}C. \tag{1.5}$$

These equations identify the demand for the individual goods as a function of the share parameter, λ, relative prices, and the consumption index.

The specification of the consumption basket in (1.1), its implications for the price index in (1.4), and individual demand functions in (1.5) prove very useful in the analyses below. In essence, if instantaneous utility over individual goods, $U(a, b)$, can be written in terms of the consumption basket, that is, $U(C)$ with $C=\mathcal{C}(a, b)$, household consumption decisions can be separated into two parts. The first is an intertemporal decision concerning the size of the current basket, C. The second is an intratemporal decision about the consumption of individual goods that make up the current basket. As we shall see, understanding how exchange rates affect both intertemporal and intratemporal consumption decisions lies at the heart of macro exchange-rate models.

Finally, it is worth noting that in the limit as the elasticity parameter θ approaches 1, the CES function in (1.1) becomes

$$\mathcal{C}(a, b)=\frac{a^{\lambda}b^{1-\lambda}}{\lambda^{\lambda}(1-\lambda)^{1-\lambda}}, \tag{1.6}$$

with the associated price index of

$$P=\left(P^a\right)^{\lambda}\left(P^b\right)^{1-\lambda}. \tag{1.7}$$

1.1.3 Purchasing Power Parity and the Law of One Price

We can now use the price indices to link the real exchange rate to the behavior of individual prices. This allows us to consider three related concepts: the Law of One Price, Absolute Purchasing Power Parity, and Relative Purchasing Power Parity.

The Law of One Price (LOOP) states that identical goods sell in two locations for the same price. This means that when the local currency price of the good for sale abroad is converted into dollars with the spot exchange rate, it will match the price of the same good available in the United States. Thus, if the LOOP applies to good a, $P^a = S\hat{P}^a$, where $\hat{P}^a$ is the local currency price of good a in the foreign country.

The LOOP has implications for the behavior of the real exchange rate. Suppose there are two goods, a and b, that make up the U.S. consumption basket with share parameters λ and $1-\lambda$, and the foreign basket with shares $\hat{\lambda}$ and $1-\hat{\lambda}$. The real exchange rate will then be given by

$$\mathcal{E} = \frac{S\hat{P}}{P} = \left(\frac{\hat{\lambda}(S\hat{P}^a)^{1-\theta} + (1-\hat{\lambda})(S\hat{P}^b)^{1-\theta}}{\lambda(P^a)^{1-\theta} + (1-\lambda)\,(P^b)^{1-\theta}} \right)^{\frac{1}{1-\theta}} .$$

If the LOOP applies to both goods, we can rewrite this expression as

$$\mathcal{E} = \left(\frac{\hat{\lambda} + (1-\hat{\lambda})(P^b/P^a)^{1-\theta}}{\lambda + (1-\lambda)\,(P^b/P^a)^{1-\theta}} \right)^{\frac{1}{1-\theta}} . \tag{1.8}$$

Clearly, $\mathcal{E}$ will equal one when $\lambda = \hat{\lambda}$. Thus, if the consumption baskets have the same composition in each country and the LOOP applies to each good, then the price of the consumption basket will be the same across countries: $P = S\hat{P}$. This condition is known as Absolute Purchasing Power Parity (PPP).

Now suppose that $\lambda \neq \hat{\lambda}$ so that good a has a different weight in the U.S. consumption basket than in the foreign consumption basket. In this case, (1.8) implies that $\mathcal{E}$ is a function of P^b/P^a. If this relative price is constant, so too is the real exchange rate, but its value can differ from one. This condition, known as Relative Purchasing Power Parity, implies that the depreciation in the spot rate is equal to the difference between the foreign and domestic rates of inflation,

$$\Delta s = \Delta\hat{p} - \Delta p,$$

where Δ denotes the first-difference and lowercase letters denote natural logs, for example, $p = \ln P$.

If $\lambda \neq \hat{\lambda}$ and P^b/P^a varies, both forms of PPP break down. This case is most easily illustrated by taking a log linear approximation of equation (1.8) around the point where $P^b = P^a$. This form of approximation will prove very useful throughout the book and is fully described in Appendix 1.A.2. Here it produces

$$\ln \mathcal{E} = \varepsilon = (\lambda - \hat{\lambda})(p^b - p^a).$$

Clearly, the correlation between the log real exchange rate and log relative prices depends on the relative weights in the two consumption baskets. One important application of this relationship arises when there is bias in consumption toward domestically produced goods. In particular, suppose that good a is produced in the United States and good b in Europe. Home bias in consumption would then be characterized by $\lambda > 1/2$ and $\hat{\lambda} < 1/2$, so clearly $\lambda - \hat{\lambda}$ would be positive. Furthermore, $p^b - p^a$ now represents the log of the U.S. terms of trade, $\tau = p^{\mathrm{M}} - (s + \hat{p}^{\mathrm{X}})$: p^b is the log price of imports and p^a is the dollar price of exports, which is equal to $s + \hat{p}^{\mathrm{X}}$ under the LOOP. Thus,

$$\varepsilon = (\lambda - \hat{\lambda})\tau.$$

Home bias in consumption implies that an improvement in the U.S. terms of trade (i.e., a fall in τ) is associated with a real appreciation of the dollar (i.e., a fall in ε).

1.2 Empirical Characteristics of Real Exchange Rates

Any successful model of exchange-rate determination must account for the behavior of both nominal and real exchange rates. This section documents two key empirical characteristics of real exchange rates that the model has to explain. First, we examine how the cross-country relative price variations across different types of goods contribute to real exchange variability. Second, we consider the volatility and persistence of real-exchange-rate variations.

1.2.1 Real Exchange Rates and Relative Prices

Variations in real exchange rates can come from many sources because national price indices are composed of the prices of many different types of goods. Historically, researchers have placed goods into two categories; nontraded and traded. The nontraded good category includes any goods that are produced solely for domestic consumption, whereas the traded category includes goods that can be consumed in any country regardless of where they are produced. With this classification, variations in the real exchange rate can be decomposed into changes in the relative price of nontraded goods across countries and changes in the relative price of traded goods across countries.

This decomposition of real-exchange-rate variations is most easily constructed using log approximations to the consumption-based price indices. In particular, assume that the U.S. consumption basket, $\mathcal{C}(\mathrm{T,N})$, is defined in terms of traded goods, T, and nontraded goods, N, with price indices P^{T} and P^{N}. In this case, equation (1.4) implies that the U.S. price level in period t is

$$P_t = \left(\lambda (P_t^{\mathrm{T}})^{1-\theta} + (1-\lambda)\,(P_t^{\mathrm{N}})^{1-\theta}\right)^{\frac{1}{1-\theta}}.$$

Log linearizing this expression around the point where $P_t^{\mathrm{N}} = P_t^{\mathrm{T}}$ gives

$$p_t = p_t^{\mathrm{T}} + (1-\lambda)(p_t^{\mathrm{N}} - p_t^{\mathrm{T}}). \tag{1.9}$$

Similarly, we can approximate the log price level in the foreign country by

$$\hat{p}_t = \hat{p}_t^{\mathrm{T}} + (1-\hat{\lambda})(\hat{p}_t^{\mathrm{N}} - \hat{p}_t^{\mathrm{T}}), \tag{1.10}$$

where $\hat{p}_t^{\mathrm{T}}$ and $\hat{p}_t^{\mathrm{N}}$ denote the log foreign currency price indices for traded and nontraded goods, respectively, and $\hat{\lambda}$ is the share parameter for traded goods in the foreign consumption basket.

Combining (1.9) and (1.10) with the definition of the log real exchange rate gives

$$\varepsilon_t = (s_t + \hat{p}_t^{\mathrm{T}} - p_t^{\mathrm{T}}) + [(1-\hat{\lambda})(\hat{p}_t^{\mathrm{N}} - \hat{p}_t^{\mathrm{T}}) - (1-\lambda)(p_t^{\mathrm{N}} - p_t^{\mathrm{T}})]. \tag{1.11}$$

The first term on the right is the log relative price of foreign traded goods in terms of U.S. traded goods. The second term is a weighted difference between the relative price of nontraded to traded goods across countries. It is convenient to identify these terms respectively by $\varepsilon_t^{\mathrm{T}}$ and $\varepsilon_t^{\mathrm{NT}}$, so that the depreciation of the real exchange rate becomes $\Delta\varepsilon_t = \Delta\varepsilon_t^{\mathrm{T}} + \Delta\varepsilon_t^{\mathrm{NT}}$. The variance of the real depreciation rate is therefore

$$\mathbb{V}(\Delta\varepsilon_t) = \mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}}) + \mathbb{V}(\Delta\varepsilon_t^{\mathrm{NT}}) + 2\mathbb{CV}(\Delta\varepsilon_t^{\mathrm{T}}, \Delta\varepsilon_t^{\mathrm{NT}}), \tag{1.12}$$

where $\mathbb{V}(\cdot)$ and $\mathbb{CV}(\cdot,\cdot)$ denote the variance and covariance operators.

Engel (1999) documents the relative contribution of $\varepsilon_t^{\mathrm{T}}$ and $\varepsilon_t^{\mathrm{NT}}$ to the variation in U.S. real-exchange-rate movements. Using monthly data on consumer prices and spot rates, he computed measures of $\varepsilon_t^{\mathrm{T}}$ and $\varepsilon_t^{\mathrm{NT}}$ for Canada, France, Germany, Italy, Japan, and the United States from January 1962 to December 1995. He then calculated the ratio, $\mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}})/\mathbb{V}(\Delta\varepsilon_t)$, over horizons ranging from 1 month to 30 years. The results of these calculations, which were striking, are reproduced in Table 1.1. Aside from the U.S.-Canada rate, his estimates of $\mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}})/\mathbb{V}(\Delta\varepsilon_t)$ in panel A are just below one at all horizons for all the other real exchange rates. He also finds that the estimates of $\mathbb{CV}(\Delta\varepsilon_t^{\mathrm{T}}, \Delta\varepsilon_t^{\mathrm{NT}})$ are very close to zero, as shown in panel B.[1] Taken together, these results imply that variations in the relative price of tradables at the consumer level account for the lion's share of the variations in real exchange rates over a wide range of horizons.

Engel's results present a challenge to traditional thinking concerning international price dynamics. To see why, let us first suppose that only one good is traded internationally. In this case $\varepsilon_t^{\mathrm{T}}$ represents the period-t deviation from the LOOP, so international arbitrage should limit the variations in $\varepsilon_t^{\mathrm{T}}$. For example, when $\varepsilon_t^{\mathrm{T}}$ falls below a certain lower bound it becomes profitable to incur the costs of importing the good to the home country, so $\varepsilon_t^{\mathrm{T}}$ should fall no further. Likewise, the profitability of exporting the good will limit the rise in $\varepsilon_t^{\mathrm{T}}$ above a certain upper bound. Under these circumstances, changes in $\varepsilon_t^{\mathrm{T}}$ are more likely to be affected by the presence of the bounds when they are computed over longer horizons. As a result, we should expect to see the variance ratio $\mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}})/\mathbb{V}(\Delta\varepsilon_t)$ fall as the horizon increases, but Engel only finds this pattern for the U.S.-Canada data.

Of course, Engel's calculations are based on price indices rather than the price of a single traded good, so variations in $\varepsilon_t^{\mathrm{T}}$ could reflect variations in the terms of trade. Recall that changes in the relative price of different traded goods can affect the real exchange rate computed from two price indices for traded goods if the weights the goods receive in each index differ. For example, if there are only two traded goods

1. Engel (1999) actually reports results for the ratios of mean square errors [i.e., $\mathrm{MSE}(\Delta\varepsilon^{\mathrm{T}})/\mathrm{MSE}(\Delta\varepsilon)$] rather than the variance ratios, but as he notes, this does not materially affect the results. In fact, he cannot reject the hypothesis that $\varepsilon_t^{\mathrm{T}}$ and $\varepsilon_t^{\mathrm{NT}}$ follow independent random walks.

TABLE 1.1
Sources of Real Exchange Variation

Horizon (months)	Canada	France	Germany	Italy	Japan
A. $\mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}})/\mathbb{V}(\Delta\varepsilon_t)$					
1	1.165	1.000	1.010	1.006	1.069
6	0.977	1.000	0.982	0.996	1.018
12	0.928	1.001	0.969	0.996	0.994
36	0.880	1.000	0.934	0.990	0.956
60	0.829	0.997	0.901	0.987	0.942
B. $\mathbb{CV}(\Delta\varepsilon_t^{\mathrm{T}}, \Delta\varepsilon_t^{\mathrm{NT}})$ $(\times 100)$					
1	−0.001	−0.006	−0.001	0.000	−0.003
6	−0.001	0.001	0.003	−0.001	−0.009
12	0.003	−0.042	0.015	−0.002	−0.002
36	0.032	−0.037	0.134	0.009	0.073
60	0.066	0.044	0.361	0.027	0.134
C. $\mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}})/[\mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}}) + \mathbb{V}(\Delta\varepsilon_t^{\mathrm{NT}})]$					
1	0.943	1.000	0.992	0.987	0.981
6	0.954	1.000	0.994	0.993	0.989
12	0.955	1.000	0.993	0.993	0.992
36	0.948	0.999	0.991	0.995	0.990
60	0.919	0.999	0.990	0.997	0.987

Source: Engel (1999).

and the shares of the U.S.-produced good are λ and $\hat{\lambda}$ in the U.S. and foreign price indices, respectively, then

$$\Delta\varepsilon_t^{\mathrm{T}} = (\lambda - \hat{\lambda})\Delta\tau_t,$$

where τ_t is the log U.S. terms of the trade. Thus, variations in $\Delta\varepsilon_t^{\mathrm{T}}$ could reflect volatility in the terms of trade if $\lambda \neq \hat{\lambda}$. Engel finds little support for this explanation in the data. The difference between λ and $\hat{\lambda}$ is just too small to account for the volatility of $\Delta\varepsilon_t^{\mathrm{T}}$ given the observed variation in the terms of trade.

Another possible explanation for Engel's findings involves the construction of the price indices. The index for traded goods is constructed from the food and all goods less food CPI subindexes, whereas the index for nontraded-goods' prices is computed from the shelter and all services less shelter CPI subindexes. These measures are imperfect. They include goods in the traded price index that are not traded (e.g., restaurant meals) and goods in the nontraded index that are traded (e.g., financial services). Nevertheless, it is very unlikely that Engel's findings are solely attributable to misclassifications like these because they appear robust to the use of non-CPI price data, such as output prices on consumption deflators.

Burstein, Eichenbaum, and Rebelo (2006) propose a related, but more far-reaching explanation for Engel's results. They argue that traded consumer goods are really composite goods with both traded and nontraded components, where the latter include distribution costs such as wholesale and retail services, marketing and advertising, and local transportation services. This means that the consumer price of each traded good comprises the price of its pure traded component and the price of the distribution services. As a result, variations in $\varepsilon_t^{\mathrm{T}}$ computed from consumer prices could reflect changes in the relative price of distribution costs across countries rather than a failure in the LOOP for pure traded goods.

To illustrate this argument, suppose there is little difference in the size of the distribution component across different traded goods, so the log consumer price indices for home- and foreign-traded goods can be approximated by

$$p_t^{\mathrm{T}} = \gamma p_t^{\mathrm{R}} + (1-\gamma) p_t^{\mathrm{D}} \quad \text{and} \quad \hat{p}_t^{\mathrm{T}} = \gamma \hat{p}_t^{\mathrm{R}} + (1-\gamma) \hat{p}_t^{\mathrm{D}},$$

where p_t^{R} and $\hat{p}_t^{\mathrm{R}}$ denote the log prices of the pure traded goods and p_t^{D} and $\hat{p}_t^{\mathrm{D}}$ are the log prices of the distribution components. The parameter γ identifies the share of the pure traded good in the composite consumption good. If the LOOP applies to pure traded goods, $p_t^{\mathrm{R}} = s_t + \hat{p}_t^{\mathrm{R}}$, the change in the log relative consumer price of traded goods becomes

$$\Delta\varepsilon_t^{\mathrm{T}} = \Delta s_t + \Delta\hat{p}_t^{\mathrm{T}} - \Delta p_t^{\mathrm{T}} = (1-\gamma)\Delta\varepsilon_t^{\mathrm{D}}, \tag{1.13}$$

with $\varepsilon_t^{\mathrm{D}} = s_t + \hat{p}_t^{\mathrm{D}} - p_t^{\mathrm{D}}$. Similarly, if traded consumer goods have the same share parameter in the U.S. and foreign price indices [i.e., $\lambda = \hat{\lambda}$ in equations (1.9) and (1.10)], then

$$\begin{aligned}\Delta\varepsilon_t^{\mathrm{NT}} &= (1-\lambda)\left[(\Delta\hat{p}_t^{\mathrm{N}} - \Delta\hat{p}_t^{\mathrm{T}}) - (\Delta p_t^{\mathrm{N}} - \Delta p_t^{\mathrm{T}})\right] \\ &= (1-\lambda)\Delta\varepsilon_t^{\mathrm{N}} - (1-\lambda)(1-\gamma)\Delta\varepsilon_t^{\mathrm{D}},\end{aligned} \tag{1.14}$$

with $\varepsilon_t^{\mathrm{N}} = s_t + \hat{p}_t^{\mathrm{N}} - p_t^{\mathrm{N}}$.

Equations (1.13) and (1.14) show that changes in the relative price of distribution services, $\Delta\varepsilon_t^{\mathrm{D}}$, contribute to variations in both components of the real depreciation rate. Consequently, it is possible that variations in $\Delta\varepsilon_t^{\mathrm{D}}$ and $\Delta\varepsilon_t^{\mathrm{N}}$ could account for Engel's findings. To see how, let $\hat{\psi}$ denote Engel's estimate of the variance ratio, $\mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}})/\mathbb{V}(\Delta\varepsilon_t)$, which is close to one. Now if $\mathbb{CV}(\Delta\varepsilon_t^{\mathrm{T}}, \Delta\varepsilon_t^{\mathrm{NT}}) = 0$, as Engel finds, $\hat{\psi} = \mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}})/(\mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}}) + \mathbb{V}(\Delta\varepsilon_t^{\mathrm{NT}}))$. Panel C of Table 1.1 shows that this ratio is very close to one. Rearranging this expression gives

$$\frac{\mathbb{V}(\Delta\varepsilon_t^{\mathrm{NT}})}{\mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}})} = \frac{1-\hat{\psi}}{\hat{\psi}}. \tag{1.15}$$

Thus, Engel's results imply that the relative variance of $\Delta\varepsilon_t^{\mathrm{NT}}$ must be very small compared to the variance of $\Delta\varepsilon_t^{\mathrm{T}}$. This is possible in the presence of distribution costs provided that the variance of $\Delta\varepsilon_t^{\mathrm{D}}$ is sufficiently large. Specifically, equations (1.13) and (1.14) imply that

$$\mathbb{CV}(\Delta\varepsilon_t^{\mathrm{NT}}, \Delta\varepsilon_t^{\mathrm{T}}) = (1-\lambda)(1-\gamma)\left[\mathbb{CV}(\Delta\varepsilon_t^{\mathrm{N}}, \Delta\varepsilon_t^{\mathrm{D}}) - (1-\gamma)\mathbb{V}(\Delta\varepsilon_t^{\mathrm{D}})\right] \quad (1.16)$$

and

$$\frac{\mathbb{V}(\Delta\varepsilon_t^{\mathrm{NT}})}{\mathbb{V}(\Delta\varepsilon_t^{\mathrm{T}})} = \left(\frac{\mathbb{V}(\Delta\varepsilon_t^{\mathrm{N}})}{(1-\gamma)^2} + \mathbb{V}(\Delta\varepsilon_t^{\mathrm{D}}) - \frac{2\mathbb{CV}(\Delta\varepsilon_t^{\mathrm{N}}, \Delta\varepsilon_t^{\mathrm{D}})}{1-\gamma}\right)\frac{(1-\lambda)^2}{\mathbb{V}(\Delta\varepsilon_t^{\mathrm{D}})}. \quad (1.17)$$

Equation (1.16) implies that $\mathbb{CV}(\Delta\varepsilon_t^{\mathrm{N}}, \Delta\varepsilon_t^{\mathrm{D}}) = (1-\gamma)\mathbb{V}(\Delta\varepsilon_t^{\mathrm{D}})$ if $\mathbb{CV}(\Delta\varepsilon_t^{\mathrm{T}}, \Delta\varepsilon_t^{\mathrm{NT}}) = 0$. Substituting this restriction into (1.17), combining the result with (1.15), and rearranging yields

$$\frac{\mathbb{V}(\Delta\varepsilon_t^{\mathrm{D}})}{\mathbb{V}(\Delta\varepsilon_t^{\mathrm{N}})} = \frac{\hat{\psi}\,(1-\lambda)^2}{(1-\gamma)^2\left(1-\hat{\psi}+\hat{\psi}\,(1-\lambda)^2\right)}. \quad (1.18)$$

Equation (1.18) shows how large the variance of $\Delta\varepsilon_t^{\mathrm{D}}$ must be compared to the variance of $\Delta\varepsilon_t^{\mathrm{N}}$ in order to account for the values of $\hat{\psi}$ estimated by Engel. Note that when the estimates are close to one, the right-hand side is approximately equal to the inverse of the squared share of distribution costs in the consumer prices of traded goods, $(1-\gamma)^{-2}$. Burstein, Neves, and Rebelo (2003) argue that this share is between 0.4 and 0.5 for the United States, which implies that the variance of $\Delta\varepsilon_t^{\mathrm{D}}$ must be four to six times larger than the variance of $\Delta\varepsilon_t^{\mathrm{N}}$ to account for Engel's results. A priori, it is hard to understand why changes in the relative price distribution services, $\Delta\varepsilon_t^{\mathrm{D}}$, should be so much more variable than changes in the relative price of nontraded goods. Indeed, theoretical models that allow for distribution costs, such as that of Corsetti, Dedola, and Leduc (2008a), do not distinguish between distribution and nontraded sectors, so that $\Delta\varepsilon_t^{\mathrm{D}} = \Delta\varepsilon_t^{\mathrm{N}}$.

The preceding analysis shows that the existence of distribution costs *alone* cannot account for Engel's results unless the volatility of those costs is implausibly large. That is not to say that distribution costs are unimportant. Burstein, Eichenbaum, and Rebelo (2006) find that variations in $\varepsilon_t^{\mathrm{T}}$ account for a smaller share of the variance of ε_t when $\varepsilon_t^{\mathrm{T}}$ is computed with the prices of goods at the dock rather than consumer prices.[2] It appears, therefore, that some of the variations in the relative price of tradables at the consumer level are attributable to distribution costs. The remainder must come from failures in the LOOP for pure traded goods.

2. These results are not directly comparable to Engel's because Burstein, Eichenbaum, and Rebelow compute $\mathbb{V}(\tilde{\varepsilon}_t^{\mathrm{T}})/\mathbb{V}(\tilde{\varepsilon}_t)$, where $\tilde{\varepsilon}_t$ and $\tilde{\varepsilon}_t^{\mathrm{T}}$ are the cyclical components of ε_t and $\varepsilon_t^{\mathrm{T}}$ estimated from the Hodrick-Prescott filter.

1.2.2 Volatility and Autocorrelation

Table 1.2 reports the variance and correlations of monthly real and nominal depreciation rates for the United States versus the United Kingdom and Japan between January 1975 and December 2007; for the United States versus Germany between January 1975 and December 1998; and the United States versus the Euro-area between January 1999 and December 2007. The upper panel shows that real depreciation rates are as volatile as nominal depreciation rates and that the two are very highly correlated. These two features imply that very little of the variation in the real depreciation rate is attributable to changes in the inflation differential. In particular, since $\Delta\varepsilon_t = \Delta s_t + \Delta\hat{p}_t - \Delta p_t$ by definition, the variance of the real depreciation rate can be decomposed as

$$\mathbb{V}(\Delta\varepsilon_t) = \mathbb{CV}(\Delta s_t, \Delta\varepsilon_t) + \mathbb{CV}(\Delta\hat{p}_t - \Delta p_t, \Delta\varepsilon_t).$$

Rewriting this expression in terms of the correlation between the nominal and real depreciation rates, $\mathbb{CR}(\Delta\varepsilon_t, \Delta s_t)$, and dividing by the variance of $\Delta\varepsilon_t$ gives

$$1 = \mathbb{CR}(\Delta\varepsilon_t, \Delta s_t)\sqrt{\frac{\mathbb{V}(\Delta s_t)}{\mathbb{V}(\Delta\varepsilon_t)}} + \frac{\mathbb{CV}(\Delta\hat{p}_t - \Delta p_t, \Delta\varepsilon_t)}{\mathbb{V}(\Delta\varepsilon_t)}.$$

According to the statistics in the first three rows of Table 1.2, the first term on the right-hand side has an average value of 0.97 across all the currency pairs. Consequently, changing inflation differentials account for roughly 3 percent of the variance of the real depreciation rate.

The center panel of Table 1.2 reports the variance and correlations for the monthly change in the log bilateral U.S. terms of trade, $\Delta\tau_t$. These data are computed as $\tau_t = s_t + \hat{p}_t^{\mathrm{X}} - p_t^{\mathrm{X}}$ from the log of the export price indices for the United States and the foreign countries, p_t^{X} and $\hat{p}_t^{\mathrm{X}}$. A rise in τ_t therefore represents a fall in the relative price of U.S. exports compared to dollar export prices of the foreign country. As the table shows, with the exception of the JPY/USD data, the variance of $\Delta\tau_t$ is higher than the variance of the nominal and real depreciation rates. Changes in the terms of trade are also very strongly correlated with the real depreciation rates. These high correlations are not peculiar to the United States. Obstfeld and Rogoff (2000) show that the correlations between bilateral terms of trade and nominal depreciation rates are strongly positive across 15 country pairings.

One possible explanation for these findings is that household preferences are biased toward the consumption of domestically produced traded goods. To see why, let us approximate the log traded price indices in the United States and foreign country as

$$p_t^{\mathrm{T}} = \mu p_t^{\mathrm{X}} + (1-\mu)(s_t + \hat{p}_t^{\mathrm{X}}) \quad \text{and} \quad \hat{p}_t^{\mathrm{T}} = \hat{\mu}\hat{p}_t^{\mathrm{X}} + (1-\hat{\mu})(p_t^{\mathrm{X}} - s_t),$$

where the share parameters μ and $\hat{\mu}$ identify U.S. and foreign household preferences between domestic- and foreign-produced traded goods. (Note that these equations incorporate the assumption that the LOOP applies to the exports of traded goods.)

TABLE 1.2
Real and Nominal Exchange-Rate Statistics

	EUR/USD	DM/USD	GBP/USD	JPY/USD
$\mathbb{V}(\Delta s_t)$	6.89	11.05	8.92	10.40
$\mathbb{V}(\Delta \varepsilon_t)$	6.85	11.20	9.44	10.90
$\mathbb{CR}(\Delta s_t, \Delta \varepsilon_t)$	0.99	0.99	0.98	0.99
$\mathbb{V}(\Delta \tau_t)$	7.59	11.31	10.02	7.54
$\mathbb{CR}(\Delta \varepsilon_t, \Delta \tau_t)$	0.99	0.97	0.90	0.84
$\mathbb{CR}(\varepsilon_t, \varepsilon_{t-1})$	0.97	0.99	0.97	0.98
$\mathbb{CR}(\varepsilon_t, \varepsilon_{t-2})$	0.94	0.98	0.94	0.96
$\mathbb{CR}(\varepsilon_t, \varepsilon_{t-3})$	0.91	0.96	0.91	0.94
$\mathbb{CR}(\Delta \varepsilon_t, \Delta \varepsilon_{t-1})$	0.15	0.01	0.05	0.09
$\mathbb{CR}(\Delta \varepsilon_t, \Delta \varepsilon_{t-2})$	0.03	0.09	0.00	0.05
$\mathbb{CR}(\Delta \varepsilon_t, \Delta \varepsilon_{t-3})$	−0.08	0.03	−0.013	0.09

Notes: The log real exchange rate in month t, ε_t, is computed as $s_t + \hat{p}_t - p_t$, where s_t is the log spot rate (FX/USD), $\hat{p}_t$ is the log foreign consumer price index, and p_t is the log U.S. consumer price index in month t. The bilateral terms of trade, τ_t, are computed as $s_t + \hat{p}^{\mathrm{X}}_t - p^{\mathrm{X}}_t$, where p^{X}_t and $\hat{p}^{\mathrm{X}}_t$ denote the U.S. and foreign price indices for exports. Depreciation rates are calculated as the monthly difference in the log level, i.e., $\Delta s_t \equiv s_t - s_{t-1}$, $\Delta\varepsilon_t \equiv \varepsilon_t - \varepsilon_{t-1}$, and $\Delta\tau_t = \tau_t - \tau_{t-1}$ multiplied by 100.

Combining these expressions with the equation for the log real exchange rate in (1.11) gives

$$\Delta\varepsilon_t = \left(\hat{\mu} + \mu - 1\right) \Delta\tau_t + \Delta\varepsilon^{\mathrm{NT}}_t.$$

Thus the real rate of depreciation reflects variations in the terms of trade when $\hat{\mu} + \mu \neq 1$ and variations in the relative price of nontraded to traded goods across countries, $\Delta\varepsilon^{\mathrm{NT}}_t$. Using this equation to compute the variance of the real depreciation rate and rearranging the result yields

$$\frac{\mathbb{CV}(\Delta\varepsilon^{\mathrm{T}}_t, \Delta\varepsilon_t)}{\mathbb{V}(\Delta\varepsilon_t)} = \left(\hat{\mu} + \mu - 1\right) \left\{ \mathbb{CR}(\Delta\varepsilon_t, \Delta\tau_t) \sqrt{\frac{\mathbb{V}(\Delta\tau_t)}{\mathbb{V}(\Delta\varepsilon_t)}} \right\}.$$

According to the results in Engel (1999), the ratio on the left is slightly less than one, whereas the statistics in Table 1.2 imply that the term in parentheses on the right is slightly larger than one. Consequently, the sum of the share parameters, $\hat{\mu} + \mu$, must be a little less than two. Although this represents a strong degree of consumption bias, these calculations oversimplify matters in two important respects: they ignore the presence of multiple trading partners and LOOP deviations for traded goods.

The lower two panels of Table 1.2 report the first-, second-, and third-order autocorrelation coefficients for both the log level and the change in the real exchange rate, ε_t and $\Delta\varepsilon_t$. There is strong autocorrelation in ε_t across all the currency pairs and the autocorrelation coefficients are close to one. At the same time, there is little autocorrelation in $\Delta\varepsilon_t$; the autocorrelations are very close to zero. Taken together, these statistics indicate that variations in real exchange rates are quite persistent.

1.2.3 Unit Roots and Half-Lives

There is a large literature examining the persistence of real-exchange-rate movements. One strand looks at whether the time series process for the real exchange rate contains a unit root. Papers in this strand of the literature first considered the behavior of individual rates over single currency regimes and over longer time spans that cover multiple regimes. More recently, attention has centered on the joint behavior of multiple rates under the post–Bretton Woods floating regime. These studies focus on whether shocks to real exchange rates have *any* permanent effect on their level. The second strand of the literature quantifies the rate at which the effects of a shock on the level of the real-exchange-rate decay. A popular metric for this rate is the half-life of the shock, that is, the time it takes for the effect of the shock to be one-half of its initial impact. Recent research in this strand of the literature investigates whether the rate of real-exchange-rate decay can be reconciled with the characteristics of price-setting found in micro data.

Wold and Beveridge-Nelson Decompositions

We can illustrate the issues in both strands of the literature with the aid of the Wold Decomposition theorem. In particular, if the rate of real depreciation follows a mean-zero covariance-stationary process, the Wold theorem tells us that the time series for $\Delta\varepsilon_t$ may be represented by

$$\begin{aligned}\Delta\varepsilon_t &= n_t + b_1 n_{t-1} + b_2 n_{t-2} + \cdots \\ &= b(L)n_t,\end{aligned} \tag{1.19}$$

where $b(L) = 1 + b_1 L + b_2 L^2 \cdots$ is a polynomial in the lag operator, L, (i.e., $Lx_t = x_{t-1}$), and n_t are the errors in forecasting $\Delta\varepsilon_t$ based on a projection (i.e., a linear regression) of $\Delta\varepsilon_t$ on its past values, $n_t = \Delta\varepsilon_t - \mathbb{P}(\Delta\varepsilon_t|\Delta\varepsilon_{t-1}, \Delta\varepsilon_{t-2}, \ldots)$. These errors are also referred to as time series innovations. The Wold theorem also tells us that the moving average representation in (1.19) is unique (i.e., both $\{b_i\}$ and $\{n_t\}$ are unique) and that $E(n_t) = 0$, $E(n_t \Delta\varepsilon_{t-j}) = 0$ and $E(n_t n_{t-j}) = 0$ for $j > 0$ and $E(n_t^2) = \sigma_n^2$.

The Wold Decomposition in (1.19) allows us to identify the presence of a unit root in a very straightforward manner. Let $b(1) = 1 + \sum_i b_i$ denote the sum of the coefficients in $b(L)$. Then we can rewrite (1.19) as $(1 - L)\varepsilon_t = b(1)n_t + (b(L) - b(1))n_t$. Multiplying both sides of this expression by $(1 - L)^{-1}$ gives

$$\varepsilon_t = \mu_t + y_t, \tag{1.20}$$

where

$$\mu_t = \mu_{t-1} + b(1)n_t \quad \text{and} \quad y_t = b^*(L)n_t, \tag{1.21}$$

with $b^*(L) = (1 - L)^{-1} (b(L) - b(1))$. This representation of a time series is known as the Beveridge-Nelson Decomposition. In this context, it decomposes the log level of the real exchange rate into a trend component, μ_t, that follows a random walk and a cycle component, y_t.

Equations (1.20) and (1.21) imply that the real exchange rate contains a unit root when the variance of shocks to the trend, μ_t, is greater than zero. Note though, that this variance is equal to $b(1)^2\sigma_n^2$, so ε_t contains a unit root if and only if $b(1) \neq 0$. At the same time, the Wold Decomposition in (1.19) identifies the long-run effect of an n_t shock on the level of ε_t by $b(1)$. Thus, although the presence of a unit root implies that n_t shocks have *some* long-run effect on the level of the real exchange rate, it says nothing about the size of the effect [i.e., the size of $b(1)$]. This is an important point. If $b(1)$ is very close to zero, the process for ε_t will contain a unit root, but the behavior of the real exchange rate over any finite sample period will be indistinguishable from the behavior implied by a process for ε_t without a unit root. This form of near observational equivalence makes it impossible to reliably test for the presence of a unit root in a finite data sample even though the test has desirable statistical properties asymptotically (i.e., as the sample size grows without limit). To address this problem, the econometrics literature on unit roots considers tests on the joint null hypothesis of $b(1) = 0$ plus auxiliary restrictions on $b(L)$ that disappear asymptotically.

Half-Lives

The Wold Decomposition theorem also allows us to identify the half-life of a shock. First, we note from (1.19) that the impact of an n_t shock on the log level of the real exchange rate h periods later is $1 + \sum_{i=1}^{h-1} b_i$. The half-life, $\mathcal{H}$, is then the shortest time period, h, such that $1 + \sum_{i=1}^{h-1} b_i$ is less than one-half the original impact of the n_t shock, which in this case is just one:

$$\mathcal{H} = \min_{h \geq 1} \left(1 + \sum_{i=1}^{h-1} b_i\right) \leq 1/2. \tag{1.22}$$

The length of the half-life depends on the size of any unit root in the real exchange rate. To see why, we first note that $1 + \sum_{i=1}^{h-1} b_i = b(1) - \sum_{i=h}^{\infty} b_i$ for all $h \geq 1$. Substituting this identity in (1.22) implies that

$$\mathcal{H} = \min_{h \geq 1} \left(b(1) + b_h^*\right) \leq 1/2, \tag{1.23}$$

where $b_h^* = -\sum_{i=h}^{\infty} b_i$. It is easy to show that b_h^* are the coefficients in the moving average polynomial $b^*(L)$ governing the dynamics of the cycle component in the Beveridge-Nelson Decomposition. As such, we know that $\lim_{h\to\infty} b_h^* = 0$, because y_t is mean-reverting. Thus, (1.23) implies that the half-life, $\mathcal{H}$, must be finite in the absence of a unit root, that is, when $b(1) = 0$. When a unit root is present, the length of the half-life depends on the pattern of the b_i coefficients. For example, in the case where the real exchange rate follows a random walk, all the b_i coefficients are zero, so the half-life is infinite. Alternatively, if the b_i coefficients are such that $b_h^* = \rho^h$ and $b(1) < 1/2$, then the half-life is $\{\ln(1 - 2b(1)) - \ln(2)\}/\ln(\rho)$. Thus, the presence of a unit root in the real exchange rate does not necessarily imply that the half-life of a shock is infinite. Nor does the length of the half-life indicate anything about the size of the unit root component, $b(1)$.

Much of the literature on real-exchange-rate persistence has focused on estimates of the half-life implied by a stationary AR(1) process. In this case the log level of the real exchange rate is assumed to follow $\varepsilon_t = \rho\varepsilon_{t-1} + v_t$, with $|\rho| < 1$. This process implies that $b(1) = 0$ and $b_h^* = \rho^h$, so the half-life is $-\ln(2)/(\ln \rho)$.[3] According to Rogoff (1996), consensus estimates of ρ in the literature imply a half-life for the real exchange rate of 3 to 5 years. Although such a slow rate of decay could plausibly reflect the effects of productivity and/or taste shocks on the relative price of nontraded to traded goods across countries, $\varepsilon_t^{\mathrm{NT}}$, Engel's findings suggest that variations in $\varepsilon_t^{\mathrm{NT}}$ contribute very little to the volatility of real exchange rates over short- and medium-term horizons. Thus, we are left to ponder what could account for both the high volatility and slow rate of decay in response to shocks in the relative price of traded goods, $\varepsilon_t^{\mathrm{T}}$. This is often termed the Purchasing Power Parity, or PPP, Puzzle.

1.2.4 Aggregation Bias and the PPP Puzzle

At first sight, stickiness in consumer prices appears to offer a simple explanation for the PPP Puzzle. If firms enjoy some monopoly power in both home and foreign markets and adjust the local currency prices for their products infrequently, changes in nominal exchange rates could induce persistent deviations from the LOOP for individual goods, which in turn are reflected in the behavior of $\varepsilon_t^{\mathrm{T}}$. The key issue here is whether the degree of price-stickiness we observe for individual traded goods is sufficient to account for a half-life for $\varepsilon_t^{\mathrm{T}}$ in the 3- to 5-year range. Rogoff (1996) argued that this was unlikely because the first-order effects should occur while prices are sticky—a time frame far shorter than 3 to 5 years. Chari, Kehoe, and McGrattan (2002) provide formal support for this intuition. They are unable to generate the volatility and persistence of real-exchange-rate variations seen in the data from a calibrated general equilibrium model with sticky prices. Imbs, Mumtaz, Ravn, and Rey (2005), hereafter IMRR, present an alternative view. They argue that the estimated persistence of real-exchange-rate variations can be reconciled with the degree of price-stickiness seen in individual prices once proper account is taken of aggregation and heterogeneity. This argument is far from intuitive, so it is worth considering more closely.

Suppose, for illustrative purposes, that households consume just two traded goods, A and B, with dollar prices P^{A} and P^{B}, and euro prices $\hat{P}^{\mathrm{A}}$ and $\hat{P}^{\mathrm{B}}$. To eliminate any terms-of-trade effects, we also assume that U.S. and E.U. households have the same share parameters in their consumption baskets. Thus, the log real exchange rate equals $\varepsilon_t^{\mathrm{T}}$ and can be approximated by

$$\varepsilon_t = \lambda\varepsilon_t^{\mathrm{A}} + (1-\lambda)\varepsilon_t^{\mathrm{B}}, \tag{1.24}$$

where $\varepsilon_t^{\mathrm{X}} = s_t + \hat{p}_t^{\mathrm{X}} - p_t^{\mathrm{X}}$ denotes the LOOP deviation for good $\mathrm{X} = \{\mathrm{A,B}\}$. At this point we do not consider the pricing decisions of the two firms in any detail. Instead,

3. Following the literature, this calculation ignores the fact that the half-life should be an integer. With this restriction the half-life is the smallest integer at least as large as $-\ln(2)/(\ln \rho)$.

we assume that market segmentation allows firms to choose local currency prices, $\hat{p}_t^{\text{X}}$ and p_t^{X}, given the behavior of the equilibrium spot rate, s_t, so that ε_t^{X} follows an AR(1) process,

$$\varepsilon_t^{\text{X}} = \rho_{\text{X}} \varepsilon_{t-1}^{\text{X}} + u_t + v_t^{\text{X}}, \tag{1.25}$$

where u_t and v_t^{X} are mean zero, mutually and serially independent shocks, with variances σ_u^2 and σ_{X}^2. Realizations of v_t^{X} represent firm-specific shocks that arise from the pricing decisions of firm $\text{X} = \{\text{A},\text{B}\}$. By contrast, realizations of u_t are common across firms. As such, they could represent the effects of a change in the nominal exchange rate when both firms keep their prices fixed. Equation (1.25) allows for two sources of heterogeneity across firms. Pricing policies can generate firm-specific shocks with different variances, σ_{X}^2, and different rates of mean-reversion via the AR(1) coefficients, ρ_{X}.

Equations (1.24) and (1.25) allow us to investigate how heterogeneity in the pricing decisions of firms shows up in the dynamics of the real exchange rate. As a first step, we rewrite (1.24) using (1.25) to substitute for ε_t^{A} and ε_t^{B}:

$$\varepsilon_t = \left(\rho_{\text{A}} + \rho_{\text{B}}\right) \varepsilon_{t-1} - \rho_{\text{A}} \rho_{\text{B}} \varepsilon_{t-2} + \omega_t, \tag{1.26}$$

where $\omega_t = u_t + \lambda v_t^{\text{A}} + (1-\lambda) v_t^{\text{B}} - \left(\rho_{\text{B}} \lambda + \rho_{\text{A}}(1-\lambda)\right) u_{t-1}$

$$- \lambda \rho_{\text{B}} v_{t-1}^{\text{A}} - (1-\lambda) \rho_{\text{A}} v_{t-1}^{\text{B}}.$$

Note that two lags of the log real exchange rate appear on the right-hand side of (1.26) and that ω_t contains both the current and lagged values of the pricing shocks. Taken together, these features imply that the log real exchange rate follows an ARMA(2,1) process. This is an example of a general result concerning the aggregation of individual time series: An average of N variables that all follow an AR(1) process will generally follow an ARMA(N, $N-1$) process. Consequently, the half-life of the average process will usually differ from the half-lives of the individual variables. The one exception to this result occurs when the individual variables have the same AR(1) parameter. If $\rho_{\text{A}} = \rho_{\text{B}} = \rho$, equation (1.26) simplifies to $\varepsilon_t = \rho \varepsilon_{t-1} + u_t + \lambda v_t^{\text{A}} + (1-\lambda) v_t^{\text{B}}$ so the half-life for the real exchange rate provides an accurate measure of the half-life of the relative price for individual traded goods.

Up to this point we have seen how the aggregation of individual heterogeneous time series for LOOP deviations can produce a process for the real exchange rate with different time series properties. We now turn to the question of whether this form of aggregation can reconcile the micro evidence on price-stickiness with a half-life for the real exchange rate of between 3 and 5 years.

The first step is to calibrate the processes for the LOOP deviations. In a recent study of a large cross-section of good prices, Crucini and Shintani (2008) find that LOOP deviations for traded goods have half-lives ranging from 9 to 12 months. These results imply values for the ρ_{X} coefficients in monthly data between 0.925 and 0.943, so we calibrate ρ_{A} and ρ_{B} at these values. The behavior of the real exchange rate also depends on the share parameter, λ, which we set equal to 1/2. Finally, we need values for the three variances: σ_u^2, σ_{A}^2, and σ_{B}^2. Since there is no direct micro evidence concerning the appropriate calibration of these parameters, we choose values for σ_{A}^2,

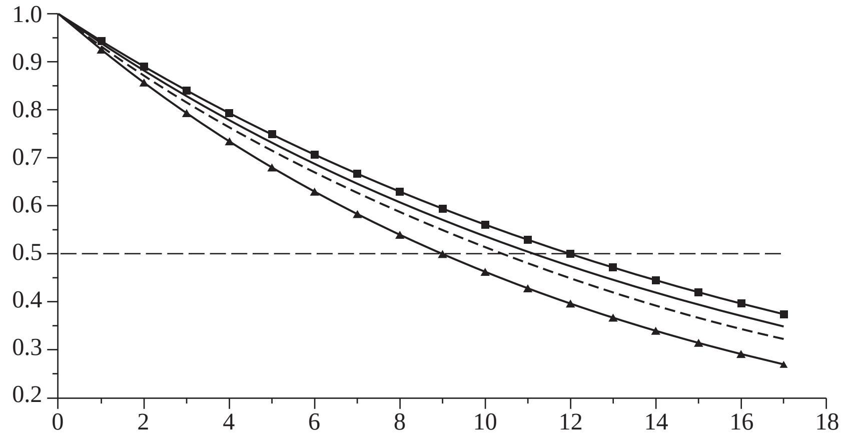

FIGURE 1.1 Impulse response functions (v_t^{A} shocks, solid triangles; v_t^{B} shocks, solid squares; u_t shocks, dashed; and n_t innovations, solid).

$\sigma_{\text{B}}^2/\sigma_{\text{A}}^2$, and $\mathbb{CR}(u_t + v_t^{\text{A}}, u_t + v_t^{\text{B}})$ such that the implied value for $\mathbb{V}(\Delta\varepsilon_t)$ equals 0.11, the estimated variance of the EUR/USD real depreciation rate reported in Table 1.2.

The second step is to compute the parameters of the ARMA(2,1) process implied by (1.26). Without loss of generality, we can write this process as

$$\varepsilon_t = \alpha_1\varepsilon_{t-1} + \alpha_2\varepsilon_{t-2} + n_t + \beta_1 n_{t-1}, \tag{1.27}$$

where n_t are serially uncorrelated mean-zero innovations with variance σ_n^2. We find the values for α_1, α_2, β_1, and σ_n^2 that produce the same autocorrelation pattern [i.e., the values of $\mathbb{CV}(\varepsilon_t, \varepsilon_{t-k})$ for all $k \geq 0$] as the calibrated process in (1.26). With these parameter values in hand, we can compare the impulse response functions from the ARMA(2,1) process against those implied by LOOP deviations and examine the half-life of the real exchange rate with respect to v_t^{A}, v_t^{B}, u_t, and n_t.

Figure 1.1 shows the impulse responses of the real exchange rate for the case where $\sigma_{\text{B}}^2/\sigma_{\text{A}}^2 = 2$ and $\mathbb{CR}(u_t + v_t^{\text{A}}, u_t + v_t^{\text{B}}) = 0.2$. The square and triangle plots show the impulse responses to a one unit v_t^{A} and v_t^{B} shock, respectively. Note that these plots intersect the 0.5 line at 9 and 12 months, representing the half-lives of the calibrated LOOP deviations. The dashed plot shows the impulse response with respect to a one-unit u_t shock, the shock that is common to both goods. As one would expect, the half-life of this shock lies between 9 and 12 months. Its exact value, $\mathcal{H}$, solves $1/2 = \lambda\rho_{\text{A}}^{\mathcal{H}} + (1-\lambda)\rho_{\text{B}}^{\mathcal{H}}$, which in this case turns out to be 10.36 months. The impulse response to a unit innovation, n_t, is shown by the solid plot. Once again, this response lies between those for the v_t^{A} and v_t^{B} shocks, but it also differs slightly from the response to the u_t shock. It is this difference that represents the effects of aggregation. More specifically, the figure shows that the half-life of real-exchange-rate innovations is slightly longer than the half-life of a common shock to the LOOP deviations. In this sense, there is positive aggregation bias in the half-life.

The distinction between economic shocks and innovations provides the key to understanding the source of aggregation bias. In this example there are three economic shocks, $\{v_t^{\mathrm{A}}, v_t^{\mathrm{B}},$ and $u_t\}$, that represent the effects of nominal-exchange-rate variations and firms' pricing decisions. As Figure 1.1 shows, each of these shocks has a different dynamic impact on the real exchange rate. By contrast, the innovations, n_t, are the errors from forecasting the real exchange rate based on a projection of ε_t on its past values: $n_t = \varepsilon_t - \mathbb{P}(\varepsilon_t|\varepsilon_{t-1}, \varepsilon_{t-2}, \ldots)$. As such, they do not generally represent the effects of any one economic shock or even a linear combination of shocks. Consequently, it should not be a surprise to find that the exchange rate responds differently to an innovation than to a particular economic shock.

With this perspective, we can now think about the possible factors affecting the size of the aggregation bias. In particular, aggregation bias will be greater when the innovations are poorer approximations to the common shock, u_t. This can occur in two main ways. First, aggregation bias increases with the variance ratio, $\sigma_{\mathrm{B}}^2/\sigma_{\mathrm{A}}^2$. In this case the innovations take on more of the characteristics of the v_t^{B} shocks, so their half-life rises. Second, aggregation bias increases as the difference between ρ_{B} and ρ_{A} gets larger. For example, if ρ_{B} equals 0.955 so that the half-life for v_t^{B} shocks is 15 months, aggregation bias increases to 2 months. On the other hand, if ρ_{B} equals ρ_{A}, aggregation bias disappears.

This discussion suggests that aggregation bias could offer a potential explanation for the PPP Puzzle if the half-lives for the LOOP deviations of *some* goods are 2 to 3 years and the innovations to the real exchange rate well approximate the shocks specific to these goods because their variance is much larger than those for other goods. For example, the half-life of the real exchange rate rises to 28 months if we increase $\sigma_{\mathrm{B}}^2/\sigma_{\mathrm{A}}^2$ to 5 and ρ_{B} to 0.981. Note though, that even in this case, the half-life of the real exchange rate is still shorter than the 36-month half-life implied by the new value for ρ_{B}. Thus, if LOOP deviations for traded goods have half-lives ranging from 9 to 12 months, as reported by Crucini and Shintani (2008), aggregation bias alone cannot resolve the PPP Puzzle.

The discussion of aggregation bias to this point omits a key element in the argument presented by IMRR. They claim that aggregation effects can account for the *estimated* half-life of the real exchange rate of 2 to 3 years when the estimates are derived from an AR(1) specification for ε_t. It is important to recognize that this is not a pure aggregation bias story. It also involves the estimation of the real exchange rate's half-life. This aspect of IMRR's analysis would not be important if the data on the real exchange rate spanned a very long time period. Under these circumstances it would be possible to precisely estimate the parameters of the ARMA(N,$N-1$) process for ε_t implied by the aggregation of N LOOP deviations. These estimates could then be used to compute an accurate half-life for the real exchange rate. Unfortunately, the span of the available data is far too short for IMRR to implement this method. Instead, they compute the half-life implied by their estimates of ρ from an AR(1) model: $\varepsilon_t = \rho\varepsilon_{t-1} + \omega_t$. As a consequence, their estimates of ρ have to be corrected for two problems: (1) finite sample bias induced by the short span of the data, and (2) misspecification bias that arises when the implied ARMA(N, $N-1$) process for ε_t does not simplify to an AR(1). As one can imagine, quantifying these biases is far from straightforward. Indeed, although IMRR claim that they are large enough to resolve the PPP Puzzle when combined with pure aggregation bias, Chen and Engel (2005) argue the opposite. At this point, the most we can say is that the question of whether aggregation and estimation biases resolve the PPP Puzzle remains open.

1.3 Macro Exchange-Rate Models

1.3.1 Overview

We now study the behavior of the nominal exchange rate in a canonical macro model with two countries, the United States and Europe. Each is populated by a large number of identical utility-maximizing households that have access to markets for a complete set of contingent claims. In this setting, households are able to share risk completely, a feature that has important implications for the behavior of exchange rates. To appreciate these implications fully, we first examine the exchange-rate implications of a model where outputs of all the goods are exogenous. We then extend the model to include production. Here, risk-sharing also affects the investment decisions of firms and hence the behavior of output. In the light of Engel's findings concerning the importance of the relative prices of traded versus nontraded goods, there are no nontraded goods in the model. Instead, households consume baskets composed of two traded goods, one produced in the United States and one produced in Europe.

Households

The United States is populated by a continuum of identical households distributed on the interval [0,1/2) with preferences defined over a consumption basket of traded goods and real balances. In particular, the expected utility of a representative U.S. household in period t is given by

$$\mathbb{U}_t = \mathbb{E}_t \sum_{i=0}^{\infty} \beta^i \left\{ \tfrac{1}{1-\gamma} C_{t+i}^{1-\gamma} + \tfrac{\chi}{1-\nu} \left(M_{t+i}/P_{t+i} \right)^{1-\nu} \right\}, \tag{1.28}$$

where γ, ν, and χ are positive parameters and $1 > \beta > 0$ is the subjective discount factor. $\mathbb{E}_t$ denotes expectations conditioned on period-t information and C_t is the consumption index defined over two consumption goods, a U.S.-produced good, C_t^{US}, and a European-produced good, C_t^{EU}. We assume that the index takes the CES form, so $C_t = \mathcal{C}(C_t^{\text{US}}, C_t^{\text{EU}})$ as defined in (1.1) with elasticity parameter θ and share parameter λ for domestically produced goods. P_t is the associated consumption price index and M_t denotes the household's holdings of dollars. The budget constraint for the household is

$$FA_t + M_t = R_t^{\text{FA}} FA_{t-1} + M_{t-1} - P_t C_t, \tag{1.29}$$

where FA_t denotes the dollar value of financial assets held at the end of period t and R_t^{FA} is the (gross) nominal return on assets held between the start of period $t-1$ and t.

Europe is also populated by a continuum of households distributed on the interval [1/2,1]. Here the preferences of the representative E.U. household have an analogous form to those shown in (1.28) except that the consumption index, $\hat{C}_t$, price index, $\hat{P}_t$, and holdings of euros, $\hat{M}_t$, replace, C_t, P_t, and M_t:

$$\widehat{\mathbb{U}}_t = \mathbb{E}_t \sum_{i=0}^{\infty} \beta^i \left\{ \tfrac{1}{1-\gamma} \hat{C}_{t+i}^{1-\gamma} + \tfrac{\chi}{1-\nu} (\hat{M}_{t+i}/\hat{P}_{t+i})^{1-\nu} \right\}. \tag{1.30}$$

E.U. and U.S. households have symmetric preferences with respect to individual traded goods, so $\hat{C}_t = \mathcal{C}(\hat{C}_t^{\text{EU}}, \hat{C}_t^{\text{US}})$, where $\hat{C}_t^{\text{EU}}$ and $\hat{C}_t^{\text{US}}$ denote the consumption of E.U.- and U.S.-produced goods by E.U. households. The budget constraint for the household is

$$\widehat{FA}_t + \hat{M}_t = \hat{R}_t^{\text{FA}} \widehat{FA}_{t-1} + \hat{M}_{t-1} - \hat{P}_t \hat{C}_t, \tag{1.31}$$

where $\widehat{FA}_t$ denotes the euro value of financial assets held at the end of period t and $\hat{R}_t^{\text{FA}}$ is the corresponding period-t (gross) nominal return.

Three features of the household sector deserve comment. First, the presence of real balances in (1.28) and (1.30) generates money demand functions that are reminiscent of traditional models and allow us to examine how exogenous variations in the money supply affect exchange rates. Second, the setup omits any role for labor income; there is no disutility from work and all household income comes from asset holdings. Appendix 1.A.3 shows that household behavior is unaffected by the presence of labor income when markets are complete, so there is no cost attached to this simplifying feature. Finally, the setup leaves the composition of each household's portfolio unspecified. We study this in detail below.

Financial Markets

Both U.S. and E.U. households hold their financial wealth in contingent claims. These primitive financial assets, called Arrow-Debreu (AD) securities, allow U.S. and E.U. households to share their risks perfectly. To understand how households choose among different AD securities, let us assume that there are $\mathcal{Z}$ states of the world indexed by z. Let us also assume that AD securities pay off in dollars. Thus, a household holding one AD security for state z in period t receives \$1 if z is the period-t state and zero otherwise.

Consider the portfolio problem facing the representative U.S. household that has access to a complete set of AD securities. Let $\mathcal{P}_t(z)$ be the price at time t of an AD security that pays \$1 if the state of the world in period $t+1$ is z and zero otherwise. If $\mathcal{A}_t(z)$ denotes the number of type-z AD securities held at the end of period t, the return on last period's portfolio is $R_t^{\text{FA}} FA_{t-1} = \mathcal{A}_{t-1}(z_t)$ and the value of the U.S. household's portfolio at the end of period t is $FA_t = \sum_{\mathcal{Z}} \mathcal{P}_t(z) \mathcal{A}_t(z)$.[4] This portfolio can be decomposed into risky assets and a riskless bond. A riskless bond can be constructed from a portfolio comprising one AD security of each type. This portfolio has a payoff of \$1 next period in all states of the world and has price $\mathcal{P}_t = \sum_{\mathcal{Z}} \mathcal{P}_t(z)$. Thus, the one-period nominal interest rate implied by the prices of AD securities is $r_t = -\ln \mathcal{P}_t$. We can now represent the U.S. household's portfolio as $B_t \exp\left(-r_t\right) + \sum_{\mathcal{Z}} \mathcal{P}_t(s) \left(\mathcal{A}_t(s) - B_t\right)$, where B_t denotes the number of bond-equivalent portfolios of AD assets. Substituting these definitions into the U.S. household's budget constraint in (1.29) gives

4. To be clear, $\mathcal{A}_{t-1}(z_t)$ denotes the number of AD securities for the state in period t, z_t, held at the end of period $t-1$, that is, $\mathcal{A}_{t-1}(z_t) = \mathcal{A}_{t-1}(z = z_t)$.

$$\sum\nolimits_{\mathcal{Z}} \mathcal{P}_t(z)\left(\mathcal{A}_t(z) - B_t\right) + B_t\mathcal{P}_t + M_t = \mathcal{A}_{t-1}(z_t) + M_{t-1} - P_tC_t. \quad (1.32)$$

The budget constraint facing the E.U. household can be derived in a similar manner. The only difference arises from the fact that the available set of AD securities payoff is in dollars rather than euros. To account for this fact, we construct a set of synthetic AD securities that pay off in euros. In other words, we build a portfolio that pays off €1 in state z and zero otherwise. This is straightforward using the contingent spot rate, $S_t(z)$, which is the dollar price of euros in period t when the state is z. Recall that a type-z AD security pays off \$1 in period $t+1$ if the state is z, so in terms of euros, the payoff is €$\left(1/S_{t+1}(z)\right)$ if the state is z and zero otherwise. This means that the dollar cost in period t of purchasing a claim to €1 in state z next period is $\mathcal{P}_t(z)S_{t+1}(z)$ because we have to purchase $S_{t+1}(z)$ of type-z AD securities at price $\mathcal{P}_t(z)$. The cost in euros of putting together this portfolio is $\widehat{\mathcal{P}}_t(z) = \mathcal{P}_t(z)S_{t+1}(z)/S_t$. It is important to recognize that $\widehat{\mathcal{P}}_t(z)$ does not depend upon the actual rate of depreciation between t and $t+1$. Rather it is a function of the *contingent* rate of depreciation, $S_{t+1}(z)/S_t$ (i.e., the rate if state z is realized in period $t+1$). Consequently, the value of $\widehat{\mathcal{P}}_t(z)$ for all $z \in \mathcal{Z}$ is known to households in period t.

We are now in a position to write the E.U. household's budget constraint in terms of risky and riskless euro-denominated assets. Recall that $\widehat{FA}_t$ represents the euro value for assets held at the end of period t. Since these assets must be held in the form of AD securities, $\widehat{FA}_t$ must equal $\sum_{\mathcal{Z}} \mathcal{P}_t(z)\mathcal{N}_t(z)/S_t$, where $\mathcal{N}_t(z)$ represents the number of type-z assets held by the E.U. household at the end of period t. Combining this expression for $\widehat{FA}_t$ with the definition of $\widehat{\mathcal{P}}_t(z)$, we can write

$$\widehat{FA}_t = \sum\nolimits_{\mathcal{Z}} \widehat{\mathcal{P}}_t(z)\frac{\mathcal{N}_t(z)}{S_{t+1}(z)} = \sum\nolimits_{\mathcal{Z}} \widehat{\mathcal{P}}_t(z)(\widehat{\mathcal{A}}_t(z) - \hat{B}_t) + \widehat{\mathcal{P}}_t\hat{B}_t,$$

where $\widehat{\mathcal{A}}_t(z) = \mathcal{N}_t(z)/S_{t+1}(z)$ and $\widehat{\mathcal{P}}_t = \sum_{\mathcal{Z}} \widehat{\mathcal{P}}_t(z)$. $\widehat{\mathcal{A}}_t(z)$ represents the number of synthetic type-z AD securities that pay off in euros. $\widehat{\mathcal{P}}_t$ is the euro price of a portfolio that pays €1 in all states in the next period, so the one-period foreign nominal interest rate is $\hat{r}_t = -\ln \widehat{\mathcal{P}}_t$. $\hat{B}_t$ identifies the number of foreign bond-equivalent portfolios of AD assets held at the end of period t. Finally, note that $\hat{R}_t^{\text{FA}}\widehat{FA}_{t-1} = \widehat{\mathcal{A}}_{t-1}(z_t)$. Thus, using the expression for $\widehat{FA}_t$ above, we can rewrite the E.U. household's budget constraint as

$$\sum\nolimits_{\mathcal{Z}} \widehat{\mathcal{P}}_t(z)(\widehat{\mathcal{A}}_t(z) - \hat{B}_t) + \widehat{\mathcal{P}}_t\hat{B}_t + \hat{M}_t = \widehat{\mathcal{A}}_{t-1}(z_t) + \hat{M}_{t-1} - \hat{P}_t\hat{C}_t. \quad (1.33)$$

Goods and Money Markets

Recall that in this economy there are just two traded goods, one produced in the United States and the other in Europe. The aggregate demand for each good comprises the sum of the individual demands of U.S. and E.U. households. Let P_t^{US} ($\hat{P}_t^{\text{US}}$) and P_t^{EU} ($\hat{P}_t^{\text{EU}}$) denote the prices of the U.S. and E.U. good in dollars (euros). Since all households within each country have the same preferences, we can use the demand functions in (1.5) to write the aggregate demand for the two goods as

$$X_t = \frac{\lambda}{2}\left(\frac{P_t^{\mathrm{US}}}{P_t}\right)^{-\theta} C_t + \frac{1-\lambda}{2}\left(\frac{\hat{P}_t^{\mathrm{US}}}{\hat{P}_t}\right)^{-\theta} \hat{C}_t \tag{1.34a}$$

and

$$\hat{X}_t = \frac{1-\lambda}{2}\left(\frac{P_t^{\mathrm{EU}}}{P_t}\right)^{-\theta} C_t + \frac{\lambda}{2}\left(\frac{\hat{P}_t^{\mathrm{EU}}}{\hat{P}_t}\right)^{-\theta} \hat{C}_t. \tag{1.34b}$$

The first term on the right-hand side of each equation identifies the aggregate demand from U.S. households and the second the aggregate demand from E.U. households. (Recall that national populations are represented by one-half the unit interval, so aggregate demand from each country is one-half the demand of the U.S. and E.U. representative households.) Note, also, that U.S. demand depends on relative prices in dollars, whereas E.U. demand depends on relative prices in euros. Market clearing requires that the aggregate demand for each good, X_t and $\hat{X}_t$, match the available supply. In the case of an endowment economy, these supplies are exogenous. In the case of a production economy, the supply of each good is endogenously determined by the production decisions of U.S. and E.U. firms.

The aggregate demand for dollars and euros is determined analogously by adding the aggregate national demands. However, in this case, things are further simplified because no household will find it optimal to hold foreign nominal balances when it has access to a complete set of AD securities. As a result, the aggregate demand for dollar balances will be one half the demand of the representative U.S. household. The aggregate demand for euro balances will be similarly just one half the demand of the representative E.U. household. We assume that the aggregate supplies of dollars and euros are completely under the control of the Federal Reserve (FED) and European Central Bank (ECB), respectively.

1.3.2 Equilibrium

Equilibrium in an endowment economy constitutes a sequence for goods prices, $\{P_t^{\mathrm{US}}, P_t^{\mathrm{EU}}, \hat{P}_t^{\mathrm{US}}, \hat{P}_t^{\mathrm{EU}}\}$, interest rates, $\{r_t, \hat{r}_t\}$, AD security prices $\{\mathcal{P}_t(z)\}$, and the nominal exchange rate, $\{S_t\}$, consistent with market clearing in the goods, money, and contingent-claims markets given the optimal consumption and portfolio decisions of households, money supply decisions of central banks, and exogenous output. In a production economy, the output of each good is also determined optimally by firms. Our focus is on the process for the equilibrium exchange rate. Below we lay the groundwork for this analysis in three steps: First, we study households' decisionmaking. Second, we derive two important implications of complete risk-sharing. Third, we examine the implications of market clearing for the links among aggregate consumption, relative prices, and the supplies of traded goods.

Household Decisions

The problem facing the representative U.S. household may be expressed as follows: At the start of each period, the household observes the prices of the $\mathcal{Z}$ AD securities, $\mathcal{P}_t(z)$, and the price index, P_t. On the basis of these prices and real wealth, $W_t = \left(\mathcal{A}_{t-1}(z_t) + M_{t-1}\right)/P_t$, the household then chooses the consumption basket, C_t, the share of wealth in real balances, $\alpha_t^m = M_t/P_t W_t$, and the share of wealth held in each of the AD securities, $\alpha_t(z) = \mathcal{P}_t(z)\left(\mathcal{A}_{z,t}(z) - B_t\right)/P_t W_t$ for all $z \in \mathcal{Z}$, to maximize expected utility (1.28). We can write this problem in the form of the following dynamic programming problem (see Appendix 1.A.1 for details):

$$\mathcal{J}(W_t) = \max_{C_t, \alpha_t^m, \alpha_t(z)} \left\{ \tfrac{1}{1-\gamma} C_t^{1-\gamma} + \tfrac{\chi}{1-\nu} \left(\alpha_t^m W_t\right)^{1-\nu} + \beta \mathbb{E}_t \mathcal{J}(W_{t+1}) \right\},$$
$$\text{s.t.} \qquad (1.35)$$
$$W_{t+1} = \exp(r_t - \Delta p_{t+1})(ER_{t+1} W_t - C_t),$$

where $\mathcal{J}(W_t)$ is the U.S. value function (i.e., the maximized value of expected utility for the representative U.S. household written as a function of real wealth). The second equation rewrites the budget constraint in (1.32) in terms of real wealth and the excess return on wealth relative to the return on dollar bonds:

$$ER_{t+1} = 1 + \frac{\alpha_t(z_{t+1}) \exp(-r_t)}{\mathcal{P}_t(z_{t+1})} - \sum\nolimits_{\mathcal{Z}} \alpha_t(z) + \alpha_t^m \left(\exp(-r_t) - 1\right).$$

Note that this excess return depends on the portfolio shares chosen in period t, $\alpha_t(z)$ and α_t^m, and the state in period $t+1$, z_{t+1}.

The problem facing the representative E.U. household is analogously described by

$$\widehat{\mathcal{J}}(\hat{W}_t) = \max_{\hat{\alpha}_t(z), \hat{\alpha}_t^m, \hat{C}_t} \left\{ \tfrac{1}{1-\gamma} \hat{C}_t^{1-\gamma} + \tfrac{\chi}{1-\nu} (\hat{\alpha}_t^m \hat{W}_t)^{1-\nu} + \beta \mathbb{E}_t \widehat{\mathcal{J}}(\hat{W}_{t+1}) \right\},$$
$$\text{s.t.} \qquad (1.36)$$
$$\hat{W}_{t+1} = \exp(\hat{r}_t - \Delta \hat{p}_{t+1})(\widehat{ER}_{t+1} \hat{W}_t - \hat{C}_t),$$

where $\widehat{\mathcal{J}}(\hat{W}_t)$ is the E.U. value function defined over real wealth, $\hat{W}_t = (\widehat{\mathcal{A}}_{t-1}(z_t) + \hat{M}_{t-1})/\hat{P}_t$. The excess return on wealth relative to euro bonds is defined by

$$\widehat{ER}_{t+1} = 1 + \frac{\hat{\alpha}_t(z_{t+1}) \exp\left(-\hat{r}_t\right)}{\widehat{\mathcal{P}}_t(z_{t+1})} - \sum\nolimits_{\mathcal{Z}} \hat{\alpha}_t(z) + \hat{\alpha}_t^m \left(\exp\left(-\hat{r}_t\right) - 1\right),$$

with portfolio shares $\hat{\alpha}_t(z) = \widehat{\mathcal{P}}_t(z)(\widehat{\mathcal{A}}_t(z) - \hat{B}_t)/\hat{P}_t \hat{W}_t$ and $\hat{\alpha}_t^m = \hat{M}_t/\hat{P}_t \hat{W}_t$.

Let $\pi_t(z)$ denote the probability that the state in period $t+1$ is z, conditioned on the current state z_t. Appendix 1.A.1 shows that the first-order conditions from the representative U.S.-household's problem in (1.35) are

$$\alpha_t(z): \ \mathcal{P}_t(z) = \beta\pi_t(z)\left(\frac{C_{t+1}(z)}{C_t}\right)^{-\gamma}\frac{P_t}{P_{t+1}(z)}, \tag{1.37}$$

$$C_t: \ C_t^{-\gamma} = \beta\mathbb{E}_t\Big[C_{t+1}^{-\gamma}\exp\left(r_t - \Delta p_{t+1}\right)\Big], \tag{1.38}$$

and

$$\alpha_t^m: \ M_t/P_t = \chi^{1/\nu}\left(1-\exp(-r_t)\right)^{-1/\nu} C_t^{\gamma/\nu}, \tag{1.39}$$

where $C_{t+1}(z)$ and $P_{t+1}(z)$ denote the contingent consumption and U.S. price level in period $t+1$ if the state is z. The corresponding first-order conditions from the representative E.U. household's problem in (1.36) are

$$\hat{\alpha}_t(z): \ \widehat{\mathcal{P}}_t(z) = \beta\pi_t(z)\left(\frac{\hat{C}_{t+1}(z)}{\hat{C}_t}\right)^{-\gamma}\frac{\hat{P}_t}{\hat{P}_{t+1}(z)}, \tag{1.40}$$

$$\hat{C}_t: \ \hat{C}_t^{-\gamma} = \beta\mathbb{E}_t\Big[\hat{C}_{t+1}^{-\gamma}\exp\left(\hat{r}_t - \Delta\hat{p}_{t+1}\right)\Big], \tag{1.41}$$

and

$$\hat{\alpha}_t^m: \ \hat{M}_t/\hat{P}_t = \chi^{1/\nu}\left(1-\exp(-\hat{r}_t)\right)^{-1/\nu} \hat{C}_t^{\gamma/\nu}, \tag{1.42}$$

where $\hat{C}_{t+1}(z)$ and $\hat{P}_{t+1}(z)$ are the contingent E.U. consumption and price levels in period $t+1$.

Implications of Risk-Sharing

Access to a complete set of AD securities allows households to perfectly share risks both within and across countries. We now use the first-order conditions to characterize these implications of risk-sharing for the behavior of consumption, the real exchange rate, and the foreign exchange risk premium.

We begin by noting that first-order conditions (1.37) and (1.40) hold for all $\mathcal{Z}$ states. As such, they implicitly define the consumption plan for each household in period $t+1$. In the case of U.S. households, the plan comprises consumption contingent on each state, $C_{t+1}(z)$ for all $z \in \mathcal{Z}$. According to (1.37), these values should be chosen so that the ratio of the marginal utility of \$1 in states z and z' is proportional to the ratio of the AD security prices, with the relative likelihood of states z and z' making up the proportionality factor:

$$\frac{\mathcal{P}_t(z)}{\mathcal{P}_t(z')} = \frac{\pi_t(z)}{\pi_t(z')}\left(\frac{C_{t+1}(z)^{-\gamma}}{P_{t+1}(z)}\frac{P_{t+1}(z')}{C_{t+1}(z')^{-\gamma}}\right).$$

Since all U.S. households face the same AD security prices, $\mathcal{P}_t(z)$ and $\mathcal{P}_t(z')$, and future contingent price levels, $P_{t+1}(z)$ and $P_{t+1}(z')$, their plans for state-contingent consumption, $C_{t+1}(z)$ and $C_{t+1}(z')$, must also be the same. Thus, access to a complete set of AD securities allows U.S. households to share the risk completely within the country. E.U. households also share risks completely within Europe by choosing $\hat{C}_{t+1}(z)$ in a similar fashion based on $\widehat{\mathcal{P}}_t(z)$ and $\hat{P}_{t+1}(z)$.

Equations (1.37) and (1.40) not only identify how the consumption plans of U.S. and E.U. households should be formed across future states, but also show how the contingent consumption decisions are linked internationally. To see this clearly, we combine the identity $\widehat{\mathcal{P}}_t(z) = \mathcal{P}_t(z)S_{t+1}(z)/S_t$ with (1.37) and (1.40) to get

$$\beta\left(\frac{C_{t+1}(z)}{C_t}\right)^{-\gamma}\frac{P_t}{P_{t+1}(z)} = \frac{\mathcal{P}_t(z)}{\pi_t(z)} = \beta\left(\frac{\hat{C}_{t+1}(z)}{\hat{C}_t}\right)^{-\gamma}\frac{\hat{P}_t S_t}{\hat{P}_{t+1}(z)S_{t+1}(z)}, \quad (1.43)$$

for all $z \in \mathcal{Z}$. The left-hand side of this equation identifies the nominal intertemporal marginal rate of substitution (MRS) for the U.S. household between t and $t+1$ when the state in period $t+1$ is z. The corresponding MRS for the E.U. household measured in dollars is shown on the right-hand side. Equation (1.43) therefore shows that when households construct optimal consumption plans with access to a complete set of AD securities, the implied nominal MRSs measured in terms of a common currency are equalized internationally.

Combining the left- and right-hand sides of (1.43) yields

$$\frac{\hat{P}_{t+1}(z)S_{t+1}(z)}{P_{t+1}(z)}\frac{P_t}{\hat{P}_t S_t} = \left(\frac{C_{t+1}(z)}{C_t}\right)^{\gamma}\left(\frac{\hat{C}_{t+1}(z)}{\hat{C}_t}\right)^{-\gamma}.$$

This condition must hold for all period $t+1$ states, $z \in \mathcal{Z}$. Thus, using the definition for the real exchange rate, $\mathcal{E}_t = S_t\hat{P}_t/P_t$, we can write

$$\frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} = \left(\frac{C_{t+1}}{C_t}\right)^{\gamma}\left(\frac{\hat{C}_{t+1}}{\hat{C}_t}\right)^{-\gamma}. \quad (1.44)$$

Equation (1.44) shows that the (gross) rate of depreciation for the real exchange rate must be perfectly correlated with the international ratio of consumption growth. It should be emphasized that this is an *equilibrium* condition that arises from the international risk-sharing properties of the household consumption plans. The equation does not *determine* the real exchange rate any more than it *determines* relative consumption growth; all the variables in (1.44) are endogenous. Nevertheless, this risk-sharing condition plays an important role in determining the equilibrium dynamics of both real and nominal exchange rates.

The presence of a complete set of AD securities also affects the foreign exchange risk premium linking nominal interest rates with the depreciation of the nominal exchange rate. To see this, we start with the identities linking U.S. and E.U. nominal interest rates to the price of AD securities: $1 = \sum_{\mathcal{Z}} \mathcal{P}_t(z)\exp\left(r_t\right)$ and $1 = \sum_{\mathcal{Z}} \widehat{\mathcal{P}}_t(z)\exp\left(\hat{r}_t\right)$. Combining these identities with the definition of $\widehat{\mathcal{P}}_t(z)$ gives

$$\exp(r_t)/\exp(\hat{r}_t) = \left(\sum\nolimits_{\mathcal{Z}} \pi_t^{\mathcal{P}}(z)S_{t+1}(z)\right)\Big/ S_t, \quad (1.45)$$

where $\pi_t^{\mathcal{P}}(z) = \mathcal{P}_t(z)/\mathcal{P}_t$. The term in parentheses identifies the dollar price in period $t+1$ of a portfolio with a payoff equal to €1 that is set in period t. This is a one-period forward contract for the euro constructed from AD securities:

$\mathcal{F}_t = \sum_{\mathcal{Z}} \pi_t^{\mathcal{P}}(z) S_{t+1}(z)$. Equation (1.45) is therefore nothing other than a statement of Covered Interest Parity (CIP):

$$\exp(r_t)/\exp(\hat{r}_t) = \mathcal{F}_t/S_t.$$

The ratio of returns on one-period U.S. and E.U. bonds equals the ratio of the forward rate relative to the spot exchange rate.

Note, also, that $\sum_{\mathcal{Z}} \pi_t^{\mathcal{P}}(z) = 1$ and $\pi_t^{\mathcal{P}}(z) \geq 0$ for all $z \in \mathcal{Z}$, so the price ratios, $\pi_t^{\mathcal{P}}(z)$, can be interpreted as pseudo probabilities. With this interpretation, we can rewrite (1.45) as

$$\exp(r_t)/\exp(\hat{r}_t) = \mathbb{E}_t^{\mathcal{P}}\left[S_{t+1}/S_t\right]. \tag{1.46}$$

Equation (1.46) links the nominal interest differential to the "expected" rate of nominal depreciation, where expectations are calculated using pseudo probabilities. This expression represents the *exact* uncovered interest parity condition implied by complete markets.

The foreign exchange risk premium is found by comparing (1.46) against the standard expression for Uncovered Interest Parity (UIP). For this purpose we rewrite (1.46) as

$$\exp(r_t)/\exp(\hat{r}_t) = \exp(-\delta_t)\mathbb{E}_t\left[S_{t+1}/S_t\right], \tag{1.47}$$

where $\delta_t = \ln \mathbb{E}_t\left[S_{t+1}/S_t\right] - \ln \mathbb{E}_t^{\mathcal{P}}\left[S_{t+1}/S_t\right]$ identifies the foreign exchange risk premium. When $\delta_t = 0$, equation (1.47) simplifies to become the UIP condition in which the interest differential (on the left) equals the expected depreciation rate (on the right). This condition implies that the expected nominal return on U.S. and E.U. bonds are the same when expressed in terms of a common currency.

UIP fails to hold when the expected rate of depreciation based on the pseudo probabilities, $\mathbb{E}_t^{\mathcal{P}}\left[S_{t+1}/S_t\right]$, differs from the expected rate based on the true probabilities $\mathbb{E}_t\left[S_{t+1}/S_t\right]$. We can examine the source of this difference by returning to the first-order condition governing the portfolio choice of U.S. households. Combining the identity $\mathcal{P}_t = \sum_{\mathcal{Z}} \mathcal{P}_t(z)$ with (1.37) allows us to write the pseudo probability for state z in terms of the true probability:

$$\pi_t^{\mathcal{P}}(z) = \pi_t(z)\frac{C_{t+1}^{-\gamma}(z)/P_{t+1}(z)}{\mathbb{E}_t\left[C_{t+1}^{-\gamma}/P_{t+1}\right]}.$$

Here we see that the pseudo probability associated with state z is distinct from the true probability when the state-contingent marginal utility of \$1 in $t+1$ [i.e., $C_{t+1}^{-\gamma}(z)/P_{t+1}(z)$] differs from the expected marginal utility across all future states. This means that the pseudo expectations, $\mathbb{E}_t^{\mathcal{P}}\left[S_{t+1}/S_t\right]$, will place more (less) weight on state-contingent depreciation rates $S_{t+1}(z)/S_t$ for states z where the marginal utility of \$1 is higher (lower) than is the case when computing standard expectations. Thus, the size of the foreign exchange risk premium depends upon the correlation between $S_{t+1}(z)/S_t$ and $C_{t+1}^{-\gamma}(z)/P_{t+1}(z)$ across period $t+1$ states. For example, when the correlation is positive, $\mathbb{E}_t^{\mathcal{P}}\left[S_{t+1}/S_t\right]$ will be greater than $\mathbb{E}_t\left[S_{t+1}/S_t\right]$ and δ_t will be negative. In this situation, euro bonds provide a consumption hedge to

U.S. households because they offer high dollar returns in states where the marginal utility of \$1 is high. In equilibrium these hedging benefits must be offset by a lower expected dollar return. This is exactly what (1.47) indicates. Indeed, Appendix 1.A.2 shows that if the joint distribution of consumption, prices, and spots rates in period $t+1$ is approximately log normal, then the risk premium can be expressed as

$$\delta_t = \mathbb{CV}_t\left(s_{t+1}, \gamma c_{t+1} + p_{t+1}\right), \tag{1.48}$$

where $\mathbb{CV}_t(\cdot,\cdot)$ denotes the covariance conditioned on period-t information. This approximation to the foreign exchange risk premium under complete markets will be useful in characterizing the behavior of the spot rate in what follows.

Relative Prices and Output

We next examine the implications of goods market clearing for the behavior of output, aggregate consumption, and relative prices. Equations (1.34a) and (1.34b) show how aggregate demand for U.S. and E.U. goods depends on the aggregate consumption and relative consumer prices in both countries. Our task is to identify how relative prices must adjust to equate aggregate demand for each good with the available supply.

Since aggregate demand for each good and the price indices are nonlinear functions of individual prices, the analysis is greatly simplified if we work with log approximations. Consider the aggregate demand for U.S. goods shown in equation (1.34a). Log-linearizing this expression around the point where $p_t^{\text{US}} = p_t$, $\hat{p}_t^{\text{US}} = \hat{p}_t$ and $\hat{c}_t = c_t$ gives

$$x_t = \ln(1/2) + \lambda c_t + (1-\lambda)\hat{c}_t - \lambda\theta(p_t^{\text{US}} - p_t) - (1-\lambda)\theta(\hat{p}_t^{\text{US}} - \hat{p}_t). \tag{1.49}$$

The log aggregate demand for E.U. goods is analogously derived from (1.34b) as

$$\hat{x}_t = \ln(1/2) + \lambda\hat{c}_t + (1-\lambda)c_t - \lambda\theta(\hat{p}_t^{\text{EU}} - \hat{p}_t) - (1-\lambda)\theta(p_t^{\text{EU}} - p_t). \tag{1.50}$$

Thus, the aggregate demand for each good is increasing in both U.S. and E.U. log consumption and decreasing in relative prices.[5]

The next step is to relate relative prices to the terms of trade and LOOP deviations. By definition, $p_t^{\text{EU}} - p_t^{\text{US}} = \tau_t + \varepsilon_t^{\text{US}}$ and $\hat{p}_t^{\text{EU}} - \hat{p}_t^{\text{US}} = \tau_t + \varepsilon_t^{\text{EU}}$, where $\varepsilon_t^{\text{X}} = \hat{p}_t^{\text{X}} + s_t - p_t^{\text{X}}$ is the LOOP deviation for good X$=\{$US,EU$\}$ and $\tau_t = p_t^{\text{EU}} - s_t - \hat{p}_t^{\text{US}}$ is the log terms of trade. Since there are just two traded goods, the U.S. and E.U. log price indices are approximately

$$p_t = \lambda p_t^{\text{US}} + (1-\lambda)p_t^{\text{EU}} \quad \text{and} \quad \hat{p}_t = \lambda\hat{p}_t^{\text{EU}} + (1-\lambda)\hat{p}_t^{\text{US}}.$$

Combining these approximations with the previous definitions yields

5. The log-linearizations in (1.49) and (1.50) are derived from first-order Taylor approximations (see Appendix 1.A.2), so their accuracy depends on how closely the right-hand-side variables satisfy the restrictions at the approximation point (e.g., $p_t^{\text{US}} = p_t$, $\hat{p}_t^{\text{US}} = \hat{p}_t$ and $\hat{c}_t = c_t$). The approximations are quite accurate for the equilibria we study later because shocks to the economy only have temporary effects on $p_t^{\text{US}} - p_t$, $\hat{p}_t^{\text{US}} - \hat{p}_t$, and $\hat{c}_t - c_t$.

$$p_t^{\mathrm{US}} - p_t = -(1-\lambda)(\tau_t + \varepsilon_t^{\mathrm{US}}), \quad \hat{p}_t^{\mathrm{US}} - \hat{p}_t = -\lambda(\tau_t + \varepsilon_t^{\mathrm{EU}}),$$

and

$$\hat{p}_t^{\mathrm{EU}} - \hat{p}_t = (1-\lambda)(\tau_t + \varepsilon_t^{\mathrm{EU}}), \quad p_t^{\mathrm{EU}} - p_t = \lambda(\tau_t + \varepsilon_t^{\mathrm{US}}). \tag{1.51}$$

Thus, an improvement in the U.S. terms of trade (i.e., a fall in τ_t) is associated with an increase in the relative price of U.S.-produced goods and a fall in the relative price of E.U.-produced goods in both countries. Variations in the terms of trade also affect the real exchange rate. In particular, combining the identity $\varepsilon_t = s_t + \hat{p}_t - p_t$ with the approximations for the price indices and relative prices gives

$$\varepsilon_t = (2\lambda - 1)\tau_t + 2\lambda\bar{\varepsilon}_t, \tag{1.52}$$

with $\bar{\varepsilon}_t = \frac{1}{2}(\varepsilon_t^{\mathrm{US}} + \varepsilon_t^{\mathrm{EU}})$. Note that the real exchange rate is unaffected by the terms of trade when $\lambda = 1/2$ because goods' prices receive equal weight in both price indices. Under these circumstances, movements in the real exchange rate only reflect LOOP deviations, $\bar{\varepsilon}_t$. When household preferences are biased toward the consumption of domestically produced goods, $\lambda > 1/2$, so improvements in the U.S. terms of trade are associated with a real appreciation of the dollar.

Finally, we combine (1.49) with (1.50) and (1.51) to get

$$x_t - \hat{x}_t = 4\lambda\theta(1-\lambda)(\tau_t + \bar{\varepsilon}_t) + (2\lambda - 1)\left(c_t - \hat{c}_t\right). \tag{1.53}$$

This equation shows that the relative demand for U.S. versus E.U. goods depends on the terms of trade, LOOP deviations, and the aggregate U.S. consumption relative to E.U. consumption. More specifically, a deterioration in the U.S. terms of trade (i.e., a rise in τ_t) lowers the relative price of U.S. goods relative to E.U. goods in both countries so ceteris paribus households shift their demand toward U.S. goods via the substitution effect. Similarly, equation (1.51) shows that an increase in $\varepsilon_t^{\mathrm{US}}$ and/or $\varepsilon_t^{\mathrm{EU}}$ lowers the relative price of U.S. goods and raises the relative price of E.U. goods in both countries, which will also shift demand toward U.S. goods. Finally, the last term on the right-hand side of (1.53) shows that a relative rise in U.S. versus E.U. aggregate consumption will increase the demand for U.S. versus E.U. goods when $\lambda > 1/2$, that is, when household preferences are biased toward the consumption of domestically produced goods.

Equation (1.53) plays a key role in macro exchange-rate determination. Market clearing requires that x_t and $\hat{x}_t$ equal the log supply of U.S. and E.U. goods, respectively. These supplies vary exogenously in an endowment economy and endogenously in a production economy, but in either case equation (1.53) allows us to identify how these variations are accommodated by changes in the terms of trade and aggregate consumption.

1.3.3 Exchange Rates in an Endowment Economy

We now study the implications of complete risk-sharing for the behavior of nominal and real exchange rates in an endowment economy, where the supply of traded goods and LOOP deviations are determined exogenously and the national money supplies are controlled by central banks.

The Nominal-Exchange-Rate Equation

We first derive an equation for the equilibrium nominal exchange rate. For this purpose we use the risk-sharing condition in (1.44), the interest parity condition (1.47), and the first-order conditions for real balances in (1.39) and (1.42).

In (1.44) we established that the real depreciation rate must be proportional to relative consumption growth across countries when markets are complete. Without loss of generality, we assume that the world economy is initially in a symmetric equilibrium with $C = \mathcal{E}\hat{C}$ and $\mathcal{E} = 1$. The risk-sharing condition in (1.44) now becomes $\mathcal{E}_t = (C_t/\hat{C}_t)^{\gamma}$, which we can write in logs as

$$\varepsilon_t = \gamma(c_t - \hat{c}_t). \tag{1.54}$$

Next, we consider the interest parity condition in (1.47). Taking a log-linear approximation around the point where $r_t = \hat{r}_t = r$ and $\Delta s_{t+1} = 0$ gives

$$\mathbb{E}_t \Delta s_{t+1} = r_t - \hat{r}_t + \delta_t. \tag{1.55}$$

This equation says that the expected rate of nominal depreciation equals the interest differential between U.S. and E.U. nominal bonds plus the foreign exchange risk premium, δ_t.

Finally, we derive the log money demand function of U.S. and E.U. households by taking logs of (1.39) and (1.42) and linearizing the resulting term involving nominal interest rates:

$$m_t - p_t = \kappa + \frac{\gamma}{\nu} c_t - \sigma r_t, \tag{1.56}$$

$$\hat{m}_t - \hat{p}_t = \kappa + \frac{\gamma}{\nu} \hat{c}_t - \sigma \hat{r}_t, \tag{1.57}$$

where $\kappa = \frac{1}{\nu}\{\ln \chi + (1 - \gamma\sigma)r + \ln \gamma\sigma\}$ and $\sigma = \frac{1}{\nu}(\exp(r) - 1)^{-1} > 0$. Since national money supplies are exogenous, we can treat (1.56) and (1.57) as log approximations to the money market clearing conditions in each country.

An expression for the nominal exchange rate is easily derived from equations (1.54)–(1.57). First we subtract (1.56) from (1.57) and combine the result with (1.55), (1.54), and the definition of the real exchange rate to get

$$s_t = f_t + \sigma \mathbb{E}_t \Delta s_{t+1}, \tag{1.58}$$

where f_t denotes exchange-rate fundamentals:

$$f_t = m_t - \hat{m}_t + \left(\frac{\nu - 1}{\nu}\right) \varepsilon_t - \sigma \delta_t. \tag{1.59}$$

Next, we rewrite (1.58) as

$$s_t - f_t = b\mathbb{E}_t \left(s_{t+1} - f_{t+1}\right) + b\mathbb{E}_t \Delta f_{t+1}, \tag{1.60}$$

with $b = \frac{\sigma}{1+\sigma} < 1$. Finally, we solve forward under the assumption that $\lim_{h\to\infty} b^h \mathbb{E}_t(s_{t+h} - f_{t+h}) = 0$ to give

$$s_t = f_t + \mathbb{E}_t \sum_{i=1}^{\infty} b^i \Delta f_{t+i}. \tag{1.61}$$

Equation (1.61) links the equilibrium log spot rate to two components. The first is the long-run level implied by the current state of fundamentals, f_t. The second is proportional to the discounted present value of expected future changes in fundamentals. This component appears because expected future changes in fundamentals are reflected in the expected depreciation rate, $\mathbb{E}_t \Delta s_{t+1}$, which, as is shown in equation (1.58), directly affects the current spot rate. Note, also, that fundamentals comprise the relative money supplies and two endogenous variables: the foreign exchange risk premium, δ_t, and the real exchange rate, ε_t. The former depends on the covariance structure of fundamentals, consumption, and prices. In particular, substituting (1.61) into (1.48) gives

$$\begin{aligned}\delta_t = {} & \mathbb{CV}_t(f_{t+1}, \gamma c_{t+1} + p_{t+1}) \\ & + \sum_{i=1}^{\infty} b^i \mathbb{CV}_t \left((\mathbb{E}_{t+1} - \mathbb{E}_t)\Delta f_{t+1+i}, \gamma c_{t+1} + p_{t+1} \right). \end{aligned} \tag{1.62}$$

The risk premium comprises two elements: The first on the right is the covariance between the log marginal utility of a dollar in the next period (i.e., $-\gamma c_{t+1} - p_{t+1}$) and fundamentals, f_{t+1}. The second is the covariance between the log marginal utility and future news concerning the growth in fundamentals, $(\mathbb{E}_{t+1} - \mathbb{E}_t)\Delta f_{t+1+i} = \mathbb{E}_{t+1}\Delta f_{t+1+i} - \mathbb{E}_t \Delta f_{t+1+i}$.

Two points deserve emphasis here: First, the definition of exchange-rate fundamentals includes both exogenous and endogenous variables. This can be potentially confusing until one understands that equations (1.61) and (1.62) provide representations for the behavior of the nominal exchange rate and the risk premium that apply across a large class of models. As we shall see, the nominal rate and risk premium can always be expressed solely in terms of exogenous variables in a particular equilibrium, but (1.61) and (1.62) apply much more generally. The second point concerns equation (1.62). This is not an *explicit* expression for the risk premium because δ_t is also a component of fundamentals. Instead, the equation represents a complex restriction that links the value of the current risk premium to the second moments of future risk premia, money supplies, and the real exchange rate. These second moments are constant in the equilibria of some models, so the risk premium is also constant and (1.61) alone governs the dynamics of the nominal exchange rate. More generally, equations (1.61) and (1.62) are needed to jointly characterize the behavior of the nominal exchange rate and risk premium in terms of exogenous money supplies and the endogenous real exchange rate.

The Real-Exchange-Rate Equation

We now consider how the real exchange rate is determined. In an endowment economy, the terms of trade must adjust to equate the demand for individual traded goods with the exogenous supplies. More specifically, equation (1.53) implies that the

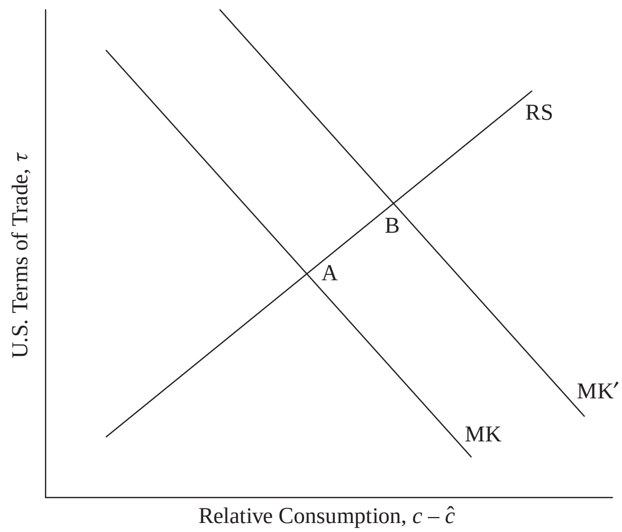

FIGURE 1.2 Goods market equilibrium.

market clearing level for the log U.S. terms of trade is

$$\tau_t = -\frac{2\lambda - 1}{4\lambda\theta(1-\lambda)}\left(c_t - \hat{c}_t\right) + \frac{1}{4\lambda\theta(1-\lambda)}(x_t - \hat{x}_t) - \bar{\varepsilon}_t, \tag{1.63}$$

where x_t and $\hat{x}_t$ now represent the exogenous log endowments of U.S. and E.U. goods. Equation (1.63) says that the U.S. terms of trade improve (i.e., τ_t falls) when U.S. consumption rises relative to E.U. consumption if there is home bias in consumption, that is, when $\lambda > 1/2$. Intuitively, the presence of home bias means that there is a shift in demand toward U.S. goods and away from E.U. goods when U.S. consumption rises relative to E.U. consumption, so the relative price of U.S. goods must rise to clear markets. The MK schedule in Figure 1.2 shows this relation between τ and $c - \hat{c}$ given the endowments of U.S. and E.U. goods, x_t and $\hat{x}_t$, and LOOP deviation, $\bar{\varepsilon}_t$.

When there is home bias in consumption, the terms of trade affect the demand for individual traded goods and cross-country differences in marginal utility. Recall that $\varepsilon_t = (2\lambda - 1)\tau_t + 2\lambda\bar{\varepsilon}_t$ from (1.52), so a deterioration in the U.S. terms of trade induces a real depreciation of the dollar when $\lambda > 1/2$. Combining this equation with the risk-sharing equation (1.54) gives the following relation between the terms of trade and relative consumption:

$$\tau_t = \frac{\gamma}{(2\lambda - 1)}(c_t - \hat{c}_t) - \frac{2\lambda}{(2\lambda - 1)}\bar{\varepsilon}_t. \tag{1.64}$$

Thus, risk-sharing requires that the U.S. terms of trade deteriorate if U.S. consumption rises relative to E.U. consumption when there is home bias. The RS schedule in Figure 1.2 shows the risk-sharing relation between τ and $c - \hat{c}$ in (1.64) for a given LOOP deviation $\bar{\varepsilon}_t$.

Figure 1.2 identifies the equilibrium terms of trade and relative consumption by the coordinates of point A, where the MK and RS schedules intersect. The figure also illustrates what happens when there is a rise in the endowment of the U.S. good. Under these circumstances, the MK schedule shifts to MK′ because relative demand must shift toward the U.S. good via a U.S. terms-of-trade deterioration for a given pattern of aggregate consumption. As the figure shows, this results in a rise in both τ and $c - \hat{c}$ (the coordinates of point B) because the implied real depreciation of the dollar must be accompanied by a rise in $c - \hat{c}$ under complete risk-sharing.

The positive relation between the terms of trade and relative endowments implied by market clearing and complete risk-sharing carries over to the real exchange rate provided $\lambda \neq 1/2$. In particular, combining (1.53), (1.54), and (1.64) gives

$$\varepsilon_t = \xi(x_t - \hat{x}_t) + \frac{4\lambda\theta(1-\lambda)}{2\lambda - 1}\xi\bar{\varepsilon}_t, \tag{1.65}$$

where

$$\xi = \frac{\gamma(2\lambda - 1)}{4\gamma\lambda\theta(1-\lambda) + (2\lambda - 1)^2}.$$

There are two noteworthy features of this equation. First, all the effects of changing endowments on the real exchange rate come via variations in the terms of trade, so ξ is positive when there is home bias in consumption (i.e., $\lambda > 1/2$). If there is no home bias, the real exchange rate only varies with LOOP deviations so the RS schedule becomes vertical [see equation (1.64)]. In that case the terms of trade vary with changes in the endowment, but have no impact on the real exchange rate. The second feature concerns the role of risk-aversion. Greater risk-aversion steepens the slope of the RS schedule, so any variation in the relative endowments or LOOP deviations has to be accompanied by a larger change in the terms of trade. For this reason, ξ is increasing in the risk-aversion parameter, γ.

We can gain further perspective on the role of risk-sharing by studying the behavior of aggregate consumption. In particular, if we combine (1.49), (1.50), and (1.51) with (1.52) and (1.65), we find that

$$c_t = \varphi x_t + (1-\varphi)\hat{x}_t - \varsigma\bar{\varepsilon}_t \quad \text{and} \quad \hat{c}_t = \varphi\hat{x}_t + (1-\varphi)x_t + \varsigma\bar{\varepsilon}_t, \tag{1.66}$$

where

$$\varphi = \frac{\lambda\,(2\lambda - 1 + 2\theta\gamma(1-\lambda))}{(2\lambda - 1)^2 + 4\gamma\lambda\theta(1-\lambda)} \quad \text{and} \quad \varsigma = \frac{2\lambda\theta(1-\lambda)}{(2\lambda - 1)^2 + 4\gamma\lambda\theta(1-\lambda)}.$$

Equation (1.66) shows that the log of aggregate consumption in each country is a weighted average of the log endowments across countries plus a term proportional to the LOOP deviation, $\bar{\varepsilon}_t$. The weighting parameter, φ, depends on the degree of home bias and risk-aversion. Recall that LOOP deviations account for all the changes in the real exchange rate when $\lambda = 1/2$. This means that consumption must be perfectly correlated across countries when markets are complete and $\bar{\varepsilon}_t$ is a constant. To achieve this, changes in each endowment must have the same impact on log consumption in each country. Hence, as the preceding definition implies, φ must equal $1/2$ when $\lambda = 1/2$. When household preferences exhibit home bias in consumption, the weighting parameter is decreasing in the risk-aversion parameter, γ. The reason is that variations in endowments have a smaller effect on the terms of trade when risk-aversion is low, so domestic demand must absorb most of any change in the supply of the domestic good. As a result, aggregate consumption is more related to changes in the endowment of domestic goods than foreign goods. Under these circumstances the weighting parameter, φ, is greater than 1/2. When risk-aversion is high, changes in endowments have larger effects on the terms of trade, so domestic demand responds less to changes in the endowment of domestic goods and more to changes in the endowment of foreign goods. Under these circumstances, the weighting parameter, φ, is smaller, but still larger than 1/2. In fact, the limiting value of φ as γ approaches infinity is 1/2. Thus, home bias in consumption has a negligible effect on the cross-country pattern of aggregate consumption at high levels of risk-aversion.

A Volatility Bound for the Real Exchange Rate

Complete risk-sharing has important implications for the cross-country correlation of consumption growth and the relative volatility of the real depreciation rate. To see why, suppose, for simplicity, that the growth in the endowments of each good are uncorrelated [i.e., $\mathbb{CV}(\Delta x_t, \Delta\hat{x}_t) = 0$] and have the same variance. Furthermore, let us assume that the LOOP deviation is constant. Under these conditions, (1.66) implies that the correlation between U.S. and E.U. consumption growth is

$$\mathbb{CR}(\Delta c_t, \Delta\hat{c}_t) = \frac{2\varphi\,(1-\varphi)}{\varphi^2 + (1-\varphi)^2}. \tag{1.67}$$

The expression on the right is a concave function of φ with a maximum of 1 when $\varphi = 1/2$ and a minimum of zero when $\varphi = \{0, 1\}$. Since the weighting parameter, φ, is decreasing in γ, greater risk-aversion increases the correlation between consumption growth across countries via its impact on φ. Greater risk aversion also affects the variance of the real depreciation rate. In particular, the risk-sharing condition $\varepsilon_t = \gamma(c_t - \hat{c}_t)$ implies that

$$\mathbb{V}(\Delta\varepsilon_t) = \gamma^2\left(\mathbb{V}(\Delta c_t) + \mathbb{V}(\Delta\hat{c}_t) - 2\mathbb{CV}(\Delta c_t, \Delta\hat{c}_t)\right).$$

Under our preceding assumptions, the variance of consumption growth is the same in both countries, so we can rewrite this expression as

$$\frac{\mathbb{V}(\Delta\varepsilon_t)}{\mathbb{V}(\Delta c_t)} = 2\gamma^2\left(1 - \mathbb{CR}(\Delta c_t, \Delta\hat{c}_t)\right). \tag{1.68}$$

Equation (1.68) clearly shows that a higher level of risk-aversion raises the relative volatility of the real depreciation rate [i.e., the variance ratio, $\mathbb{V}(\Delta\varepsilon_t)/\mathbb{V}(\Delta c_t)$] given the correlation of consumption growth across countries. However, since higher risk-aversion also pushes the correlation toward one, the relative volatility approaches an upper bound as risk-aversion rises. We can compute this upper bound if we combine the definition for φ with (1.67) and (1.68), and take the limit as $\gamma \to \infty$:

$$\lim_{\gamma\to\infty}\left\{\frac{\mathbb{V}(\Delta\varepsilon_t)}{\mathbb{V}(\Delta c_t)}\right\} = \left(\frac{(2\lambda - 1)}{2\theta\lambda\,(1-\lambda)}\right)^2. \tag{1.69}$$

With benchmark values for λ and θ of 0.85 and 0.74, respectively, the upper bound is approximately 13.8, so the variance of the real depreciation rate must be less than 14 times the variance of consumption growth in a standard model with complete markets. By contrast, the variance of the real depreciation rate for the dollar is approximately 18 times the variance of U.S. consumption growth. This is a large discrepancy. Moreover, it cannot be reconciled by changing the correlation between endowments: It is easy to check that $\lim_{\gamma\to\infty} \mathbb{CR}(\Delta c_t, \Delta \hat{c}_t) = 1$ when $\mathbb{CV}(\Delta x_t, \Delta \hat{x}_t) \neq 0$. In fact, the only way to account for the relative volatility of the real depreciation rate in a model where the risk-sharing condition in (1.54) holds true is to introduce variations in the LOOP deviations, $\bar{\varepsilon}_t$. The upper bound disappears under these circumstances because $\lim_{\gamma\to\infty} \mathbb{CR}(\Delta c_t, \Delta \hat{c}_t) < 1$. We can therefore match the relative volatility of the real depreciation rate and consumption growth with a high level of risk-aversion.

Exchange-Rate Dynamics

Equations (1.59), (1.61), (1.62), and (1.65) provide a complete characterization of the nominal exchange rate in terms of exogenous variables. In particular, we can write the equilibrium log spot rate as

$$s_t = f_t + \mathbb{E}_t \sum_{i=1}^{\infty} b^i \Delta f_{t+i}, \tag{1.70}$$

where fundamentals are now identified by

$$f_t = m_t - \hat{m}_t + \left(\frac{\nu - 1}{\nu}\right) \xi \left(x_t - \hat{x}_t + \frac{4\lambda\theta(1-\lambda)}{2\lambda - 1}\bar{\varepsilon}_t\right) - \sigma\delta_t,$$

and δ_t is determined by (1.62).

To examine the economic implications of (1.70), we begin with a special case. If $\sigma = 0$ and $\nu = \gamma$, the money market equilibrium conditions in (1.56) and (1.57) become log versions of cash-in-advance constraints:

$$m_t - p_t \cong \kappa + c_t \quad \text{and} \quad \hat{m}_t - \hat{p}_t \cong \kappa + \hat{c}_t.$$

Under these circumstances, the model has the same form as Lucas' (1982) neoclassical exchange-rate model. In particular, since the discount parameter, b, now equals zero, the equilibrium exchange rate in (1.70) becomes

$$s_t = m_t - \hat{m}_t + \left(\frac{\nu - 1}{\nu}\right) \xi \left(x_t - \hat{x}_t + \frac{4\lambda\theta(1-\lambda)}{2\lambda - 1}\bar{\varepsilon}_t\right).$$

Here the spot rate depends only on the contemporaneous money supplies, endowments, and the LOOP deviations.[6] In the Lucas model there is no home bias in consumption so ξ is equal to zero. Under these circumstances, $\varepsilon_t = 0$, so the spot rate is solely determined by the relative money supplies: $s_t = m_t - \hat{m}_t$.

We now develop some intuition for the exchange-rate behavior implied by equation (1.70) when $b > 0$. First, consider the effects of a permanent positive shock to the U.S. money stock relative to the E.U. money stock, $m_t - \hat{m}_t$. If the real exchange rate is unaffected because the monetary shock has no effect on endowments or LOOP deviations, risk-sharing requires that relative consumption across countries remain constant. As a result, the shock must be accommodated by either a fall in the interest differential, $r_t - \hat{r}_t$, or a relative rise in U.S. prices, $p_t - \hat{p}_t$, to clear the money

6. Note that fundamentals do not include the risk premium when $\sigma = 0$.

markets. The former adjustment implies an expected future appreciation of the dollar (i.e., $\mathbb{E}_t \Delta s_{t+i} < 0$), but this is inconsistent with the fact that expected future monetary growth is not affected by the shock. Consequently, adjustment in the money markets requires a one-for one relative rise in U.S. prices, so the dollar must depreciate by the same amount to keep the real exchange rate unchanged. This is exactly what is shown in (1.70). The permanent rise in $m_t - \hat{m}_t$ has no effect on $\mathbb{E}_t \Delta f_{t+i}$ for $i \geq 1$, so the spot rate depreciates one-for-one.

Now suppose that the positive shock to $m_t - \hat{m}_t$ is expected to be temporary but again there is no effect on the real exchange rate. In this case, the dollar will be expected to appreciate in the future (as the effects of the stock dissipate), so the interest differential, $r_t - \hat{r}_t$, falls when the shock hits. As a result, the relative rise in U.S. prices needed to accommodate the immediate effects of the shock on the money markets is smaller and so too is the depreciation of the dollar necessary to maintain the level of the real exchange rate. In terms of equation (1.70), the temporary rise in $m_t - \hat{m}_t$ lowers $\mathbb{E}_t \Delta f_{t+i}$ for $i \geq 1$, so the immediate impact of an increase in fundamentals is partially offset by a fall in the present value term.

Shocks to output and LOOP deviations affect the spot rate via the real exchange rate in a similar manner. A permanent shock to U.S. output, for example, induces an immediate real depreciation of the dollar, but does not affect the expected future rate of real or nominal depreciation. Consequently, the shock increases relative U.S. consumption, $c_t - \hat{c}_t$ (via risk-sharing), but leaves the interest differential unchanged. This means that there must be a relative fall in the U.S. price level, $p_t - \hat{p}_t$, to clear the money markets. When ν is greater (less) than one, the fall in $p_t - \hat{p}_t$ is smaller (larger) than the real depreciation of the dollar, so the spot rate depreciates (appreciates). When ν is equal to one, the spot rate remains unchanged because the fall in prices $p_t - \hat{p}_t$ required to clear the money markets matches the real depreciation of the dollar necessary to clear goods markets.

Variations in the risk premium, δ_t, provide a further source of variation in the sport rate. Equation (1.48) showed that variations in δ_t must come from changes in the condition covariance between the log spot rate and the log marginal utility of a dollar. If a change in the conditional distribution of future shocks to the economy permanently increases the covariance, there will be no change in the expected future depreciation rate, so the rise in the risk premium is matched by a fall in the interest differential, $r_t - \hat{r}_t$. Under these circumstances, there must be a relative fall in the U.S. price level, $p_t - \hat{p}_t$, to clear the money markets and an immediate appreciation in the spot rate to keep the real exchange rate unchanged. In terms of (1.70), f_t falls, whereas $\mathbb{E}_t \Delta f_{t+i}$ remains unchanged for $i \geq 1$.

Finally, we consider how well the model accounts for the statistical features of exchange rates discussed in Section 1.2. In particular, are the dynamic implications of (1.70) consistent with: (1) the relative volatility of nominal and real depreciation rates, (2) the high correlation between real and nominal depreciation rates, and (3) the lack of serial correlation in real depreciation rates?

These questions are easily addressed once we specify the process for the relative money supplies, endowments, and the LOOP deviations. For example, suppose that the money supplies follow

$$m_t = m_{t-1} + \mu_m + e_t \quad \text{and} \quad \hat{m}_t = \hat{m}_{t-1} + \mu_m + \hat{e}_t,$$

where e_t and $\hat{e}_t$ are i.i.d. mean-zero shocks with a common variance σ_e^2. Further, let us assume that the LOOP holds for both goods, so $\bar{\varepsilon}_t = 0$, and let the endowments of U.S. and E.U. goods evolve according to

$$x_t = \phi x_{t-1} + (1-\phi)\mu_x + \tfrac{1}{2}v_t \quad \text{and} \quad \hat{x}_t = \phi \hat{x}_{t-1} + (1-\phi)\mu_x - \tfrac{1}{2}v_t,$$

with $\phi < 1$, where v_t are i.i.d. mean-zero shocks with variance σ_v^2. Note that the v_t shocks change the *relative* endowments. Under these assumed processes, equations (1.65) and (1.70) imply that the spot and real exchange rates follow[7]

$$\Delta s_t = e_t - \hat{e}_t + \left(\frac{\nu-1}{\nu}\right)\left(\frac{1-b}{1-b\phi}\right)\Delta\varepsilon_t \quad \text{and} \quad \varepsilon_t = \phi\varepsilon_{t-1} + \xi v_t.$$

From these expressions we find that

$$\mathbb{CR}(\Delta\varepsilon_t,\, \Delta\varepsilon_{t-i}) = -\frac{\phi^i(1-\phi)\xi^2\sigma_v^2}{1+\phi} < 0,$$

for $i > 0$, and

$$\mathbb{CR}(\Delta s_t,\, \Delta\varepsilon_t) = \left(\frac{\nu-1}{\nu}\right)\left(\frac{1-b}{1-b\phi}\right)\sqrt{\frac{\mathbb{V}(\Delta\varepsilon_t)}{\mathbb{V}(\Delta s_t)}},$$

where

$$\frac{\mathbb{V}(\Delta\varepsilon_t)}{\mathbb{V}(\Delta s_t)} = \frac{\xi^2\sigma_v^2}{\sigma_e^2\,(1+\phi) + \left(\frac{\nu-1}{\nu}\right)^2\left(\frac{1-b}{1-b\phi}\right)^2\xi^2\sigma_v^2}.$$

We can therefore account for features (1), (2), and (3) if σ_e^2/σ_v^2 is small, ϕ is close to one, and ν is very large.

These calculations show that it is not difficult to replicate the statistical behavior of real and nominal exchange rates in the complete markets model *if* we are free to choose the exogenous processes. What they do not show is that the model can simultaneously account for the behavior of exchange rates and other variables. Most importantly, the assumed endowment process implies that the variance ratio $\mathbb{V}(\Delta\varepsilon_t)/\mathbb{V}(\Delta c_t)$ is far smaller than we observe in the data for any values of the risk-aversion parameter γ. Thus, if we choose the variance of the relative endowment shocks, σ_v^2, to replicate the volatility of the real depreciation rate, the implied volatility of consumption growth under complete risk-sharing is much larger than what we observe in the data.

7. Since the money supplies and endowments follow homoskedastic processes, the second moments involving all exogenous and endogenous variables are constant. Equation (1.62) showed that the risk premium must also be constant under these circumstances, so it has no effect on the nominal depreciation rate, Δs_t.

1.3.4 Exchange Rates in a Production Economy

We now extend the model to include the production decisions of the firms. This extension does not change the nominal-exchange-rate equation presented in (1.61) because the links among the spot rate, money supplies, risk premium, and the real exchange rate are not affected by firms' decisionmaking. However, these decisions do affect the behavior of the terms of trade, so this is where we concentrate our attention.

In the endowment economy with complete markets, the U.S. terms of trade are proportional to the relative supplies of U.S. to E.U. traded goods. As a result, variations in the terms of trade are directly tied to the relative supplies of traded goods, and so display no intrinsic dynamics (i.e., dynamics generated within the economy). In a production economy, by contrast, the supplies of traded goods are determined optimally by firms taking their current capital stock and the state of current and expected future productivity into account. In an international setting, firms also have to account for both the current and expected future terms of trade. This means that current production decisions not only influence the current terms of trade via their implications for the relative supplies of goods, but are themselves affected by firms' expectations about the future path of the terms of trade. This feedback effect from the terms of trade to production is ruled out in an endowment economy, but is a source of intrinsic terms-of-trade dynamics in a production economy.

Firms

Assume that there is a single industry in each country and that both industries are populated by a continuum of identical firms distributed on the interval [0,1]. A representative U.S. firm owns all of its capital stock, K_t, and produces output, Y_t, according to $Y_t = A_t K_t^{\eta}$ with $1 > \eta \geq 0$, where A_t denotes the exogenous state of productivity. The output of a representative E.U. firm, $\hat{Y}_t$, is given by an identical production function using its own capital, $\hat{K}_t$, and productivity $\hat{A}_t$. At the beginning of each period, firms observe productivity and decide how to allocate their output between investment and consumption goods. Output allocated to consumption is supplied competitively to U.S. and E.U. households and the proceeds are used to finance dividend payments to the owners of the firms' equity. Output allocated to investment adds to the stock of physical capital available for production in the next period. We also assume that both traded goods can be costlessly transported between the United States and Europe, so the LOOP applies to both goods.

Let us first focus on the problem facing a representative U.S. firm. The firm's objective is to choose investment, I_t, so as to maximize its total value to its shareholders, that is, households. Let Q_t denote the ex-dividend dollar price in period t of U.S. equity providing a claim to a dollar dividend payment of D_{t+1} at the start of period $t+1$. The firm's problem can now be written as

$$\max_{I_t}(D_t + Q_t) \quad \text{s.t.}$$

$$K_{t+1} = (1-\eth)K_t + I_t \quad \text{and} \quad D_t = P_t^{\text{US}}(A_t K_t^{\eta} - I_t), \tag{1.71}$$

where $\eth > 0$ is the depreciation rate of physical capital. Note that the dividends paid by the firm are equal to the dollar value of goods sold to households, $A_t K_t^{\eta} - I_t$.

To complete the description of the firm's problem, we have to identify how investment decisions affect the price of U.S. equity, Q_t. This is easily accomplished under complete markets because the payout from holding U.S. equity in period $t+1$, $D_{t+1} + Q_{t+1}$, can be replicated by a portfolio of AD securities. In particular, let $Q_t(z)$ denote the price of a portfolio of AD securities with a payoff of $D_{t+1}(z) + Q_{t+1}(z)$ dollars in period $t+1$ when the state is z and zero otherwise. Equation (1.37) tells us that the equilibrium price of an AD security for state z is equal to $\beta\pi_t(z)\left(C_{t+1}(z)/C_t\right)^{-\gamma}(P_t/P_{t+1}(z))$. Hence, the price of the portfolio is

$$Q_t(z) = \beta\pi_t(z)\left(\frac{C_{t+1}(z)}{C_t}\right)^{-\gamma}\left(\frac{P_t\{D_{t+1}(z) + Q_{t+1}(z)\}}{P_{t+1}(z)}\right). \qquad (1.72)$$

The price of U.S. equity, Q_t, is nothing other than the price of a claim to a payout of $D_{t+1}(z) + Q_{t+1}(z)$ in all possible states (i.e., for all $z \in \mathcal{Z}$). Thus, Q_t must equal $\sum_{\mathcal{Z}} Q_t(z)$. Hence, under complete markets, the price of U.S. equity satisfies

$$Q_t = \beta\mathbb{E}_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\left(\frac{P_t\{D_{t+1} + Q_{t+1}\}}{P_{t+1}}\right)\right]. \qquad (1.73)$$

The optimal investment decision for the representative U.S. firm is the solution to the problem in (1.71) with the restriction that the equity price, Q_t, satisfies (1.73). Solving this problem produces the following first-order condition:

$$1 = \beta\mathbb{E}_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\exp\left(\Delta p^{\text{US}}_{t+1} - \Delta p_{t+1}\right)R^k_{t+1}\right], \qquad (1.74)$$

where $R^k_{t+1} = 1 - \mathfrak{d} + \eta A_{t+1}K^{\eta-1}_{t+1}$ denotes the marginal product of U.S. capital.

The problem facing the representative E.U. firm is analogous. Investment, $\hat{I}_t$, is chosen to maximize the value of the firm, $\hat{D}_t + \hat{Q}_t$, where $\hat{Q}_t$ is the ex-dividend dollar price of E.U. equity and $\hat{D}_t$ are dollar E.U. dividends, subject to

$$\hat{K}_{t+1} = (1-\mathfrak{d})\hat{K}_t + \hat{I}_t, \qquad \hat{D}_t = P^{\text{EU}}_t(\hat{A}_t\hat{K}^{\eta}_t - \hat{I}_t), \qquad (1.75)$$

and the following restriction on E.U. equity:

$$\hat{Q}_t = \beta\mathbb{E}_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\left(\frac{P_t\{\hat{D}_{t+1} + \hat{Q}_{t+1}\}}{P_{t+1}}\right)\right]. \qquad (1.76)$$

The associated first-order condition is

$$1 = \beta\mathbb{E}_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\exp\left(\Delta p^{\text{EU}}_{t+1} - \Delta p_{t+1}\right)\hat{R}^k_{t+1}\right], \qquad (1.77)$$

where $\hat{R}^k_{t+1} = 1 - \mathfrak{d} + \eta\hat{A}_{t+1}\hat{K}^{\eta-1}_{t+1}$ denotes the marginal product of E.U. capital.

Two aspects of these characterizations deserve comment. First, it may appear from equations (1.73) and (1.76) that the price of equity depends only on how U.S. households value future payoffs, because both equations include the MRS for U.S.

households and not E.U. households. This impression is incorrect. In Section 1.3.2, equation (1.43) showed that the nominal MRS for U.S. and E.U. households measured in terms of a common currency are equalized internationally when markets are complete. This means that the dollar value of a claim to a future dollar equity payout is the same for both U.S. and E.U. households. Consequently, it makes no difference whether we express the equity restriction in (1.73) and (1.76) using the MRS for U.S. households or E.U. households. The optimal investment decisions of firms will be exactly the same in both cases. By the same token, it also makes no difference whether the equity issued by any firm is held by U.S. or E.U. households. When markets are complete, the problem facing each firm does not depend on the identity of the shareholders.

The second aspect concerns the presence of relative prices in the first-order conditions, (1.74) and (1.77). These equations show that firms choose investment so that the marginal utility of current consumption equals the discounted expected marginal utility of additional capital. Importantly, the latter depends on the marginal product of capital and the change in relative prices. For example, if the relative price of the U.S. good rises (i.e., P^{US}/P increases), the marginal utility of additional U.S. investment increases even when the marginal product of U.S. capital remains unchanged. Consequently, firms' investment decisions respond to variations in both productivity (which affects the marginal product of capital) and relative prices. As we shall see, variations in relative prices generate the feedback effects from the terms of trade to production.

Solving the Model

Although the structure of the model is quite straightforward, it is still too complex to examine analytically (even with the use of approximations). We must therefore study numerical solutions to the model—an approach that requires taking a stand on the time series process for productivity.

We adopt the conventional assumption that log productivity in each industry (i.e., $a_t = \ln A_t$ and $\hat{a}_t = \ln \hat{A}_t$) follows

$$\begin{bmatrix} a_t \\ \hat{a}_t \end{bmatrix} = \begin{bmatrix} \phi & 0 \\ 0 & \phi \end{bmatrix} \begin{bmatrix} a_{t-1} \\ \hat{a}_{t-1} \end{bmatrix} + \begin{bmatrix} u_t \\ \hat{u}_t \end{bmatrix}, \tag{1.78}$$

with $|\phi| < 1$, where u_t and $\hat{u}_t$ are mean-zero i.i.d. productivity shocks. To clarify how firms' decisions lead to the propagation of productivity shocks between countries, we assume that $\mathbb{CV}(u_t, \hat{u}_t) = 0$, so productivity in each industry follows an independent stationary AR(1) process.

A numerical solution to the model requires values for the parameters describing households' preferences, firms' technology, and the productivity processes. Here we face two options: The first is to choose parameter values via an estimation procedure that compares moments in the observed data with moments implied by the numerical solution of the model. Although details differ across procedures, implementing this option is computationally demanding because the model has to be solved for a great many different sets of parameter values. The second option is to calibrate the model, that is, to choose parameter values so that the distribution of observed data matches the distribution of data simulated from the numerical solution of the model. Again, the calibration procedures found in the literature differ in their details, but usually involve the following steps:

1. Choose some measurements from observed data that the model is to explain.
2. Choose functional forms and parameter values to ensure that the model's long-run (steady state) properties are consistent with the observed data, make sense, or are consistent with the estimates in other studies.
3. Find the numerical solution to the model and use it to generate time series for the variables the model is to explain.
4. Compare the simulated time series with the observed data.

Implementing these steps is much less computationally demanding than the estimation procedures because the model is only solved for the particular parameterization chosen in step 2.

Although the process just outlined appears relatively straightforward, several problems may undermine the economic relevance of a calibration exercise. For example, it may not be possible to pin down values for all the parameters by matching the steady state of the model against the long-run properties of available data. In these cases the preferred approach is to set the values for the unidentified parameters equal to appropriate estimates from microeconomic data. Unfortunately, this is often easier said than done because microeconomic studies rarely estimate models that are directly applicable (see, e.g., Browning, Hansen, and Heckman 1999). Parameter uncertainty creates further problems. Even when the parameters can be identified from long-run properties of the observed data and applicable microeconomic studies, precise parameter values are always obscured by some uncertainty. The best way to address this issue is to examine how the model simulations vary across different parameter choices. However, such sensitivity analyses are hard to conduct and report in a systematic and transparent manner for all but the simplest models. Finally, there is the issue of *how* to compare the simulated time series with the observed data. In other words, what metric should be used when comparing moments of the simulated and observed data? Should the metric take account of how parameter uncertainty affects moments of the simulated data, or how sampling uncertainty affects moments in the observed data, or both?

These issues have received much less attention than they deserve in the many papers that use calibration to study international macro models. Indeed, it is common practice to solve models with parameter values that are simply taken from earlier studies. One pedagogical advantage of this approach is that it allows the reader to easily compare the dynamic properties of models that share a common set of core features. Since our aim is to illustrate the properties of a standard model, it is the approach we follow in this chapter. More generally, the downside of this approach is that the calibrated parameter values in the original studies have been chosen to address a specific question (see step 1 above) that may only be loosely related to the focus of the current study. In short, full-blown calibration requires more than simply choosing parameter values used by other researchers.

The parameter values used to solve the model are reported in Table 1.3. Household preferences and firms' technologies are assumed symmetric across the two countries, and a single period in the model is interpreted as 1 month. Consistent with many models, households' discount factor and risk-aversion parameters are set to 0.997

TABLE 1.3
Production Economy Parameters

Parameter	Symbol	Value
Discount factor	β	0.997
Risk aversion	γ	2.000
Consumption share	λ	0.850
Consumption elasticity	θ	0.740
Depreciation rate	$\mathfrak{d}$	0.010
Capital share	η	0.360
Productivity AR(1)	ϕ	0.980

and 2, respectively. The share parameter, λ, and the elasticity of substitution, θ, govern preferences over the consumption of U.S. and E.U. goods. We set $\lambda = 0.74$ and $\theta = 0.85$, so the model displays home bias in consumption and imperfect substitutability between traded goods. These choices are within the wide range of values found in the literature (see Hnatkovska 2010), but are clear candidates for sensitivity analyses in a full-blown calibration exercise. Parameters on the production side of the model are chosen to be consistent with Backus, Kehoe, and Kydland (1994). The capital share in production, η, is set to 0.36 and the depreciation rate, $\mathfrak{d}$, is set to 0.01. Finally, we set the autoregressive coefficient in the productivity processes equal to 0.98.

Terms-of-Trade Dynamics

We now examine the implications of productivity shocks for the behavior of the U.S. terms of trade when firms make optimal investment decisions. Specifically, our task is to examine the behavior of the equilibrium terms of trade when productivity follows the process in (1.78). As a benchmark for what follows, let us first consider the implication of (1.78) in the special case where production requires no capital (i.e., when $\eta = 0$). In this case, market clearing in the markets for traded goods requires $X_t = A_t$ and $\hat{X}_t = \hat{A}_t$, where X_t and $\hat{X}_t$ are the aggregate demands for U.S. and E.U. goods derived in (1.34). We are thus back to the endowment version of the model, with the endowments following the productivity process in (1.78). The behavior of the terms of trade is determined by combining the goods market-clearing conditions with (1.78) and the relation between aggregate demand and the terms of trade implied by (1.52) and (1.65). In the absence of LOOP deviations, this gives

$$\tau_t = \phi\tau_{t-1} + \frac{\gamma}{4\gamma\lambda\theta(1-\lambda) + (2\lambda-1)^2}(u_t - \hat{u}_t).$$

Hence, the log terms of trade follow an AR(1) process with the same degree of persistence as the differential in log productivity across industries, $a_t - \hat{a}_t$. Clearly, there are no intrinsic terms-of-trade dynamics in this special case.

Let us now return to the general case where $1 > \eta > 0$. Three sets of equations hold the key to understanding how the terms of trade behave in response to productivity shocks when firms choose investment optimally. The first set comprises the goods' market-clearing conditions,

$$A_t K_t^{\eta} = Y_t = X_t + I_t \quad \text{and} \quad \hat{A}_t \hat{K}_t^{\eta} = \hat{Y}_t = \hat{X}_t + \hat{I}_t. \tag{1.79}$$

These expressions imply that each firm's choice of investment is equivalent to a choice of how much current output it will supply to satisfy the aggregate demand for its product. The second set of equations describes the implication of market clearing for the dynamics of the capital stock:

$$K_{t+1} = (1-\mathfrak{d})K_t + A_t K_t^{\eta} - X_t \quad \text{and} \quad \hat{K}_{t+1} = (1-\mathfrak{d})\hat{K}_t + \hat{A}_t \hat{K}_t^{\eta} - \hat{X}_t. \tag{1.80}$$

These equations express the intertemporal trade-off between current consumption of each good, (i.e., X_t and $\hat{X}_t$) and the accumulation of capital for future production.

The third set of equations links the expected marginal product of U.S. and E.U. capital with expected changes in the terms of trade. To derive this link, we first approximate the first-order conditions in (1.74) and (1.77) as

$$\gamma \mathbb{E}_t \Delta c_{t+1} = \mathbb{E}_t r_{t+1}^k + \mathbb{E}_t[\Delta p_{t+1}^{\text{US}} - \Delta p_{t+1}] + \varkappa_t \tag{1.81a}$$

and

$$\gamma \mathbb{E}_t \Delta c_{t+1} = \mathbb{E}_t \hat{r}_{t+1}^k + \mathbb{E}_t[\Delta p_{t+1}^{\text{EU}} - \Delta p_{t+1}] + \hat{\varkappa}_t, \tag{1.81b}$$

where

$$\varkappa_t = \ln \beta + \tfrac{1}{2}\mathbb{V}_t \left(\Delta p^{\text{US}} - \Delta p + r_{t+1}^k - \gamma \Delta c_{t+1} \right)$$

and

$$\hat{\varkappa}_t = \ln \beta + \tfrac{1}{2}\mathbb{V}_t \left(\Delta p^{\text{EU}} - \Delta p + \hat{r}_{t+1}^k - \gamma \Delta c_{t+1} \right).$$

Appendix 1.A.2 shows that these approximations hold exactly when log consumption, prices, and the marginal product of capital are jointly normally distributed. Given the symmetry of the model (in terms of productivity, production functions, and preferences), the variance terms in $\varkappa_t$ and $\hat{\varkappa}_t$ will be equal. Under these circumstances, (1.81) implies that

$$\mathbb{E}_t r_{t+1}^k - \mathbb{E}_t \hat{r}_{t+1}^k = \mathbb{E}_t[\Delta p_{t+1}^{\text{EU}} - \Delta p_{t+1}^{\text{US}}] = \mathbb{E}_t \Delta \tau_{t+1}. \tag{1.82}$$

Thus, when firms follow optimal investment policies, the difference between the expected log marginal product for U.S. and E.U. capital must equal the expected deterioration in the log U.S. terms of trade.

To describe the equilibrium dynamics of the terms of trade, we first have to identify how the aggregate consumption of each good varies in response to productivity shocks. The easiest way to do this is to conjecture that the log consumption of each good follows a particular process and then use the equilibrium conditions of the model to verify the conjecture. In this case, we conjecture that

$$\begin{bmatrix} x_t \\ \hat{x}_t \end{bmatrix} = \begin{bmatrix} \kappa_a & \hat{\kappa}_a \\ \hat{\kappa}_a & \kappa_a \end{bmatrix} \begin{bmatrix} a_t \\ \hat{a}_t \end{bmatrix} + \begin{bmatrix} \kappa_k & \hat{\kappa}_k \\ \hat{\kappa}_k & \kappa_k \end{bmatrix} \begin{bmatrix} k_t - k \\ \hat{k}_t - k \end{bmatrix}, \tag{1.83}$$

where k denotes the log capital stock in the steady state. At this point we do not know the values of the coefficients, κ_a, $\hat{\kappa}_a$, κ_k, and $\hat{\kappa}_k$. However, if we combine the log-linearized version of (1.80) with (1.81) and the productivity process in (1.78), we can verify that log consumption follows (1.83) and solve for the values of the coefficients. With these values in hand, we can combine (1.83) with (1.78) and the linearized version of (1.80) to describe the equilibrium dynamics of productivity, capital, and any other variable, including the terms of trade. A full description of this solution procedure is provided in Appendix 1.A.4.

We are now ready to examine how productivity shocks affect the terms of trade. Since the dynamics in this model are quite complex, we proceed in three steps: First, we consider how a U.S. productivity shock affects the behavior of U.S. output and investment in a closed-economy setting. Second, we study how the shock is transmitted between firms when U.S. and E.U. goods are perfect substitutes. This special case rules out variations in the terms of trade. Finally, we incorporate variations in the terms of trade by examining how the shock affects both firms when their goods are imperfect substitutes.

A positive productivity shock has two impacts on U.S. firms: it increases current output, and raises the expected marginal product of capital. The key choice facing each U.S. firm is how much of the additional output to allocate to investment. Here a firm faces an intertemporal trade-off. On one hand, the less output it devotes to investment, the more it can pay out in terms of current dividends to its shareholders. On the other, accumulating more capital via higher current investment will allow it to produce more in the future, and hence pay higher future dividends. Since their shareholders value both current and future dividends, firms will use some of the additional output to boost the current dividend and some to boost future dividends via increased investment. As a result, the immediate effects of a positive U.S. productivity shock in a closed economy are an increase in U.S. consumption, dividends, and investment. Thereafter, the marginal product of U.S. capital declines as capital accumulates and the effects of the productivity shock die out. As this happens, further investment becomes less attractive, so the accumulation of capital stops and output begins to fall. In the absence of further shocks, this process continues until the capital stock and productivity return to their original levels.

To understand how this transmission process differs in an open-economy setting, let us assume (for the moment) that U.S. and E.U. goods are perfect substitutes for one another. In this case, goods arbitrage implies that $P_t^{\mathrm{US}} = P_t^{\mathrm{EU}}$, so the equilibrium terms of trade always equal one. As a result, the expected marginal product of capital will be equalized across industries [e.g., $\mathbb{E}_t r_{t+1}^k = \mathbb{E}_t \hat{r}_{t+1}^k$ from equation (1.82)]. This implication of firms' optimal investment plans leads to the international transmission of productivity shocks. In particular, when a U.S. productivity shock raises the expected marginal product of U.S. capital, E.U. firms must adjust their investment plans to raise the expected marginal product of E.U. capital by the same amount. To do this, E.U. firms *reduce* investment because the marginal product of capital is a decreasing function of the capital stock. Consequently, the consumption of E.U.-produced goods immediately increases in response to a positive U.S. productivity shock. Thereafter, the consumption of E.U. production continues to rise until the loss

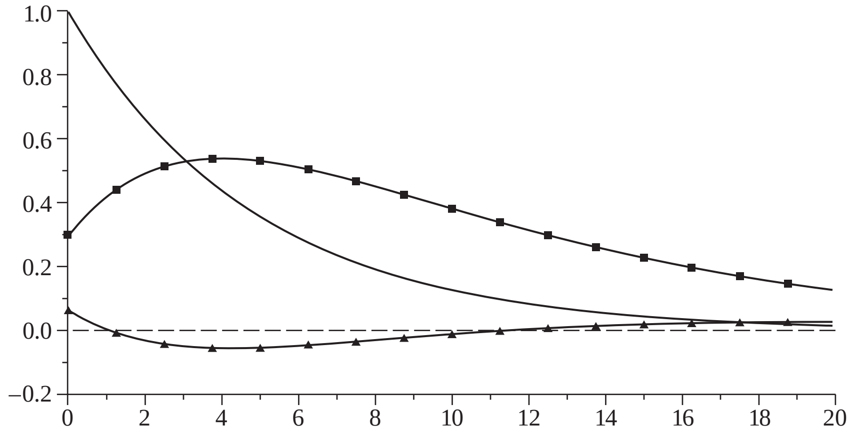

FIGURE 1.3 Impulse responses of U.S. productivity (a_t, solid) and traded goods consumption (x_t U.S., solid squares; $\hat{x}_t$ E.U., triangles).

of E.U. capital and the falling marginal product of U.S. capital makes E.U. investment attractive. From this point on, E.U. investment rises and the consumption of E.U. goods falls back toward its original level.

Now we examine the international transmission of productivity when traded goods are imperfect substitutes. For this purpose Figure 1.3 plots the impulse responses of log U.S. productivity, a_t, and the log aggregate consumption of U.S. and E.U. goods, x_t and $\hat{x}_t$, for 20 years following a one-unit productivity shock. These impulse responses are computed from the equilibrium dynamics of the model using the parameter values reported in Table 1.3.

Figure 1.3 shows that U.S. consumption rises for approximately 4 years following the productivity shock. Then, after the accumulation of U.S. capital comes to an end and the effects of the productivity shock diminish, the consumption of U.S. goods falls back toward its long-run value. The effects on E.U. consumption are quite the opposite. Consumption falls for the first 4 years before gradually returning to its long-run level. This pattern is completely different from the response of E.U. consumption when goods are perfect substitutes. To understand why, we have to consider the role played by the terms of trade in the international transmission of the U.S. productivity shock.

Figure 1.4 shows the impulse responses for the log marginal products of capital and the U.S. terms of trade. The square and triangle plots indicate the response of the log marginal product for U.S. and E.U. capital, respectively, and the solid plot depicts the log U.S. terms of trade. For comparison purposes, the figure also shows the response of the log U.S. terms of trade in the endowment economy (i.e., when $\eta = 0$) as a diamond plot. As we can see in the figure, the marginal products of U.S. and E.U. capital differ significantly in response to the productivity shock. As one would expect, a positive U.S. productivity shock immediately increases the marginal product of U.S. capital, r^k. What is more surprising is that r^k falls *below* its long-run level 4 years after the shock, and only begins to rise again after 10 years. By contrast, the productivity shock has very little impact on the marginal product of E.U. capital.

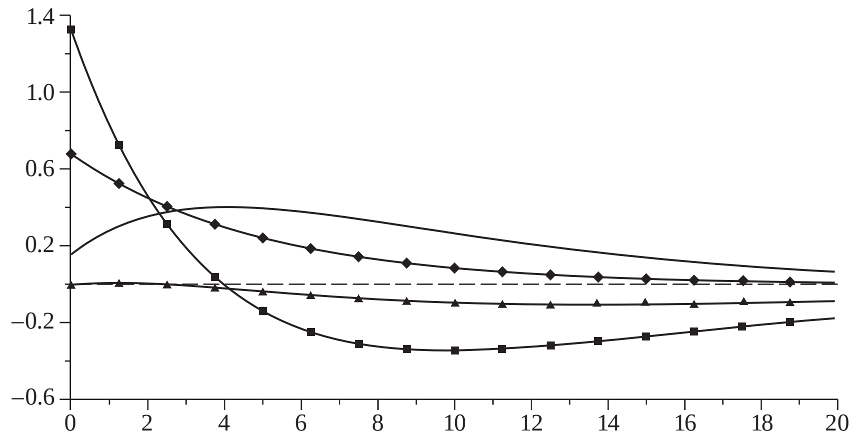

FIGURE 1.4 Impulse responses of the terms of trade (τ, solid), marginal product of capital (r^k U.S., squares; $\hat{r}^k$ E.U., triangles), and terms of trade in an endowment economy (τ, diamonds).

The terms of trade provide the key to understanding these dynamics. Figure 1.3 showed that the relative consumption of U.S. versus E.U. goods falls as the effects of the productivity shock die out. From our analysis of the endowment economy, we know that this can only happen under complete risk-sharing when there is an improvement on the U.S. terms of trade (i.e., a fall in τ). Consequently, firms must expect the U.S. terms of trade to continue to improve from some point after the productivity shock occurs. Figure 1.4 shows that this point occurs approximately 4 years after the shock. Thereafter, the expected marginal product of U.S. capital is lower than that of E.U. capital, and the slope of the terms-of-trade plot is negative. U.S. firms run down their capital stocks more rapidly than E.U. firms during this period, so the output of U.S. goods returns more quickly to its original level (see Figure 1.3).

Variations in the terms of trade also play an important role when the productivity shock hits. Figure 1.4 shows that there is an immediate deterioration in U.S. terms of trade. This has two effects. First, it changes prices so that household demand will absorb the greater supply of U.S. goods not going into investment. Second, it increases the real value of E.U. dividends. Recall that $(\hat{A}_t\hat{K}_t^{\eta} - \hat{I}_t)(P_t^{\mathrm{EU}}/P_t)$ identifies the value of E.U. dividends when measured in terms of U.S. household consumption. Since a deterioration in the U.S. terms of trade raises P_t^{EU}/P_t, it also increases the real value of E.U. dividends for a given level of E.U. production and investment. In other words, a positive U.S. productivity shock benefits the holders of E.U. equity via its impact on the terms of trade. This valuation effect reduces the incentive for E.U. firms to initially cut back on investment so the consumption of E.U. goods rises. Indeed, as U.S. firms accumulate capital and the consumption of U.S. goods rises, the U.S. terms of trade continue to deteriorate so that the valuation effect on E.U. firms becomes greater. This allows E.U. firms to increase investment while their shareholders enjoy a rise in the real value of dividends. As a result, the deterioration in the U.S. terms of trade following the productivity shock is accompanied by a fall in both the marginal product of E.U. capital and the consumption of E.U. goods (see Figure 1.3).

To summarize, productivity shocks have different dynamic implications for the behavior of the terms of trade in exchange and production economies. In the former, the terms of trade change in response to the exogenously varying supplies of traded goods simply to clear markets. In the latter, variations in the terms of trade also affect the investment decisions of firms via a valuation channel and hence the supplies of traded goods available for consumption. It is this valuation channel that provides the feedback from the terms of trade to consumption that is absent in endowment economies.

Exchange-Rate Implications

The exchange-rate implications of productivity shocks in the production economy are straightforward. The equilibrium nominal exchange continues to follow (1.70):

$$s_t = f_t + \mathbb{E}_t \sum_{i=1}^{\infty} b^i \Delta f_{t+i}, \tag{1.84}$$

with fundamentals, $f_t = m_t - \hat{m}_t - \sigma\delta_t + \left(\frac{\nu-1}{\nu}\right)\xi(x_t - \hat{x}_t)$, but now x_t and $\hat{x}_t$ vary endogenously with the capital stocks and productivity according to (1.83). All the exchange-rate implications of productivity shocks are therefore captured via their effects on x_t and $\hat{x}_t$. Moreover, since $\xi(x_t - \hat{x}_t) = \varepsilon_t = (2\lambda - 1)\tau_t$, we can rewrite fundamentals using the terms of trade as

$$f_t = m_t - \hat{m}_t - \sigma\delta_t + \frac{(\nu - 1)(2\lambda - 1)}{\nu}\tau_t,$$

and hence trace the effects of productivity shocks on the spot rate via their implications for the terms of trade and fundamentals.

Productivity shocks generally induce more persistence and greater spot rate volatility in the production economy than in the endowment economy. The reason for these differences are easily understood with the aid of Figure 1.4. There we saw that U.S. productivity shocks lead to a smaller initial deterioration in the U.S. terms of trade than in the endowment economy, but that the variations are nonmonotonic and longer lasting. This means that a productivity shock has a smaller impact on current fundamentals in the production economy, but a larger impact on the present value of expected future changes in fundamentals. The balance of these effects determines how the spot rate reacts to a productivity shock. When the discount parameter, b, is close to one, the spot rate more closely reflects variations in the expected future growth of fundamentals; when b is close to zero, the spot rate more closely reflects variations in current fundamentals. In this model, the discount parameter, b, is equal to $\sigma/(1+\sigma)$, where σ is the semi-interest elasticity of money demand. Estimates in the literature give values for σ between 20 and 60, which implies that b is very close to one. Consequently, variations in the present value term will be an important source of spot rate variation when productivity shocks have persistent effects on fundamentals via their impact on the terms of trade.

To illustrate these effects, Figure 1.5 compares the impulse response of fundamentals and the spot rate to a U.S. productivity shock in both the endowment and

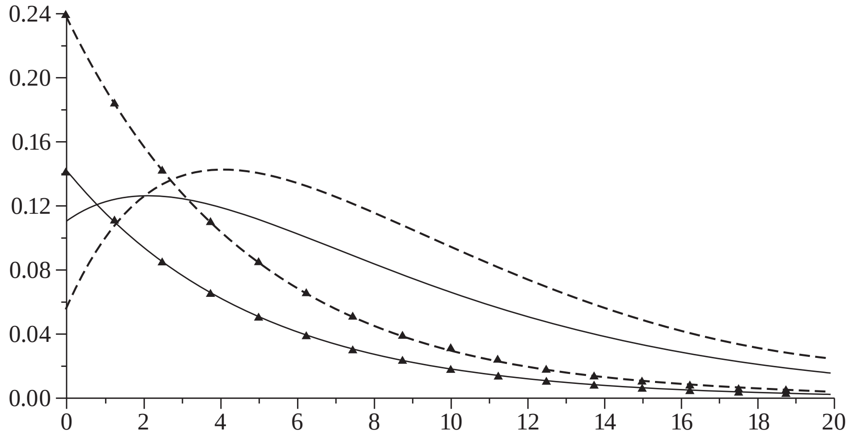

FIGURE 1.5 Impulse responses of the spot exchange rate, s_t (solid), and fundamentals, f_t (dashed), to a U.S. productivity shock in the production and endowment (with triangles) economies.

production economies. The dashed plots show the responses to fundamentals computed from the behavior of the terms of trade with $\nu = 2$. The figure shows that the U.S. productivity shock has a greater initial impact on fundamentals in the endowment economy (depicted with triangles) than in the production economy. Thereafter, fundamentals fall monotonically in the endowment economy until the effects of the productivity shock die out. In the production economy, by contrast, fundamentals rise for 4 years before slowly falling back to their original level. The difference in these response patterns simply reflects the different behavior of the terms of trade in the endowment and production economies that was illustrated in Figure 1.4.

Figure 1.5 shows a clear distinction between the effects of the productivity shock on the spot rate in the endowment and production economies. The solid and solid with triangles plots show the impulse responses for spot rates computed with $b = 40/41$, a value for the discount parameter that is consistent with a mid-range estimate of σ equal to 40. As the figure shows, productivity shocks have a smaller impact on the spot rate than on fundamentals in the endowment economy. The reason is that everyone expects fundamentals to fall following the productivity shock, so $s_t - f_t = \mathbb{E}_t \sum_{i=1}^{\infty} b^i \Delta f_{t+i}$ is negative until the effects of the shock disappear. In other words, expectations concerning the path of future fundamentals dampen the response of the spot rate. The reverse holds true in the production economy. In this case productivity shocks have a greater impact on the spot rate than fundamentals for approximately 2 years following the shock. During this period expectations regarding the near-term rise in fundamentals dominate longer-term expectations, so $s_t - f_t = \mathbb{E}_t \sum_{i=1}^{\infty} b^i \Delta f_{t+i}$ is positive and the spot rate depreciates. Thereafter, the expected future fall in fundamentals dominates so $s_t - f_t = \mathbb{E}_t \sum_{i=1}^{\infty} b^i \Delta f_{t+i}$ is negative, and the spot rate begins to appreciate.

We can now see why it is possible for productivity shocks to induce greater volatility in the spot rate in a production economy than in an endowment economy. In the former case, expectations regarding the future path of fundamentals magnify the impact of a productivity shock on the spot rate, whereas in the latter expectations act as a damper. As a result, when b is close to one, Figure 1.5 shows that there is a larger initial depreciation of the spot rate in the production economy than in the endowment economy when a positive U.S. productivity shock arrives, even though the shock induces a smaller initial rise in fundamentals. The figure also clearly shows how greater persistence in the response of fundamentals in the production economy is reflected in the response of the spot rate. In this example, the half-life of the productivity shock increases from approximately 3.5 years to 14 years.

We saw in Section 1.3.3 that it is possible to replicate the statistical features of real and nominal depreciation rates in an endowment economy with complete risk-sharing if we are free to specify the processes for money supplies and the endowments. The problem was that these specifications implied counterfactual behavior for consumption. In this section, we have seen how productivity shocks affect real and nominal exchange rates via their impact on firms' investment decisions. Since these decisions establish the supply of goods available for consumption, productivity shocks also determine the behavior of equilibrium consumption. Indeed, this model has exactly the same implications for the *joint* behavior of consumption and the real exchange rate as the endowment model. Consequently, equation (1.69) still identifies the upper bound on the variance ratio, $\mathbb{V}(\Delta\varepsilon_t)/\mathbb{V}(\Delta c_t)$, in the absence of LOOP deviations. This upper bound has a value of 3.5 for the parameters in Table 1.2, compared to the variance ratio in the data of approximately 18. Thus, although the addition of production increases the volatility of the nominal depreciation rate, it does nothing to address the counterfactual implication of complete risk-sharing for the relative volatility of the real depreciation rate and consumption growth.

1.4 Summary

- Variations in the real exchange rate can reflect changes in the relative prices of nontraded goods across countries and/or variations in the prices of traded goods across countries. Empirically, variations in the relative price of tradables at the consumer level account for the lion's share of real-exchange-rate volatility over a wide range of horizons. These variations in the relative consumer prices are attributable to the presence of distribution costs and LOOP deviations for pure traded goods.
- Monthly changes in the log spot rate, the log real exchange rate, and the log terms of trade have similar levels of volatility and are highly correlated with one another. There is very little autocorrelation in monthly real depreciation rates. Shocks to the level of the real exchange rate appear to be very persistent, with estimated half-lives ranging from 3 to 5 years.
- The high volatility and persistence of real-exchange-rate variations has proved hard to explain theoretically—a fact often referred to as the PPP Puzzle. Although stickiness in consumer prices offers a potential explanation, the degree of price-stickiness in individual traded goods appears insufficient to account for both the high volatility and the persistence of real rate variations. The PPP Puzzle may also reflect the distortions induced by the aggregation of heterogeneous dynamics for individual traded-goods prices. The question