Chapter 23

International Macroeconomics

Giancarlo Corsetti, Luca Dedola and Simon Lloyd

23.1 Introduction

Previous chapters have examined macroeconomic dynamics within closed economies, abstracting from interactions with the rest of the world via international trade and financial flows. Yet, in reality, the global economy is deeply interconnected; no country operates in isolation. National business cycles may be impacted by spillovers from foreign shocks, while many events—like the IT revolution, the 2007-2009 Global Financial Crisis (GFC), the economic fallout of the Covid-19 pandemic, and recent geopolitical developments—are global in nature. At the same time, openness may help a country to smooth domestic shocks. All this underscores the importance of modeling open economies. In this chapter, we broaden our analysis, moving from a closed economy to a financially integrated one.

We start by describing how national accounting incorporates trade in goods and financial assets between countries. We then summarize the main stylized facts of international business cycles. Grouping countries by GDP per capita, these facts lay foundations for the next two chapters: this one dedicated to the study of interactions between large advanced countries with strong institutions (i.e., those committed to repaying their debt); and the next focused on small-open economies that may default. Throughout, we show how business-cycle evidence challenges theory, highlighting numerous "puzzles" in international macroeconomics.

In this chapter, we introduce a workhorse two-country open-economy model designed to represent the dynamics of large, developed economies. We use the model to examine how TFP shocks in one country spill over to others, affect international relative prices and alter external balances. We also note that two stylized empirical facts challenge the predictions of standard models. First, net exports and the current account (i.e., external savings) are counter-cyclical; that is, external deficits become larger during booms.[1] Second, international

[1] In the conventional Mundell-Fleming model presented in undergraduate textbooks, a counter-cyclical external balance is generated with the ad hoc assumption that the 'marginal propensity to import' from income is positive, resulting in strong complementarities between foreign and domestic goods and positive cross-border transmission.

relative prices (i.e., the real exchange rate and the terms of trade) tend to appreciate when output is relatively high.

We explain how the degree of international risk sharing by households—governed by the structure of international financial markets—and the degree of substitutability between domestic and imported goods are key for reconciling these facts within the model. Following an increase in domestic productivity, a fall in the international price of home output redirects domestic and foreign demand towards goods produced in the home economy. At the same time, however, a lower price also reduces the value of any given quantity of home output—all else equal, reducing the relative income and wealth of households in the home country. We explain that model specifications in which these income fluctuations are imperfectly insured can better capture international business-cycle facts.

The core of the chapter focuses on settings with fully flexible prices. Nevertheless, the setup can be augmented to a two-country New Keynesian framework, with nominal rigidities, as we discuss briefly in our conclusion alongside other model extensions.

23.1.1 National accounting in the open economy

Recall the definition of Gross Domestic Product (GDP) from Chapter 1:

$$Y = C + I + G + NX,$$

where Y represents GDP, C is private consumption, I is investment, G is government spending and NX is net exports. NX represent the difference between a country's exports (X) and imports (M) of goods and services. They can be positive (*trade surplus*) or negative (*trade deficit*) depending on economic interactions (in goods and asset markets) with the rest of the world.

A comprehensive view of how a country interacts economically with other nations through trade, investment and financial transfers is provided by the Balance of Payments (BOP). The BOP records all international transactions. In the U.S., BOP data can be found at the Bureau of Economic Analysis, in the 'International Transactions' section.[2] BOP information for other countries is compiled by the IMF on their 'BOP information' website.[3]

The BOP has three components: the Current Account (CA), the Financial Account (FA), and the Capital Account (KA), satisfying $CA + FA + KA = 0$. The CA records transactions related to goods, services, income, and current transfers. It includes the trade balance (or net exports NX), as well income from abroad.[4] In a country like the U.S., where net factor income and net transfers from abroad are small, $CA \simeq NX$. The FA records changes in ownership of financial assets and liabilities between a country and the rest of the world.[5] The KA tracks capital transfers and the acquisition or disposal of non-produced,

[2]See https://www.bea.gov/data/intl-trade-investment/international-transactions, Table 1.1, Balance of Current Account (annual data).

[3]See https://data.imf.org/?sk=7A51304B-6426-40C0-83DD-CA473CA1FD52.

[4]The $CA = NX + NFI + NUT$, where NFI is Net Factor Income (income earned from foreign investments minus income paid to foreign investors) and NUT is Net Unilateral Transfers (gifts from foreign countries minus gifts to them).

[5]It includes Direct Investment (long-term investments in foreign businesses or assets), Portfolio Investment (investments in foreign stocks, bonds, or other financial instruments) and Other Investments (other financial transactions).

non-financial assets (e.g., patents and copyrights). It is typically smaller in magnitude compared to the other two accounts, implying that $CA \simeq -FA$.

The evolution of the BOP in the U.S. (left) and detail on the three accounts (right) are depicted in Figure 23.1.[6] The sharp decline in the external balance of the U.S. from the mid-1990s, captured in Figure 23.1 has given rise to a large literature on "global imbalances" (see, e.g., Caballero, Farhi, and Gourinchas, 2008) and a modern reconsideration of the U.S. "exorbitant privilege" from issuing the global reserve currency (see, e.g., Farhi and Maggiori, 2017).

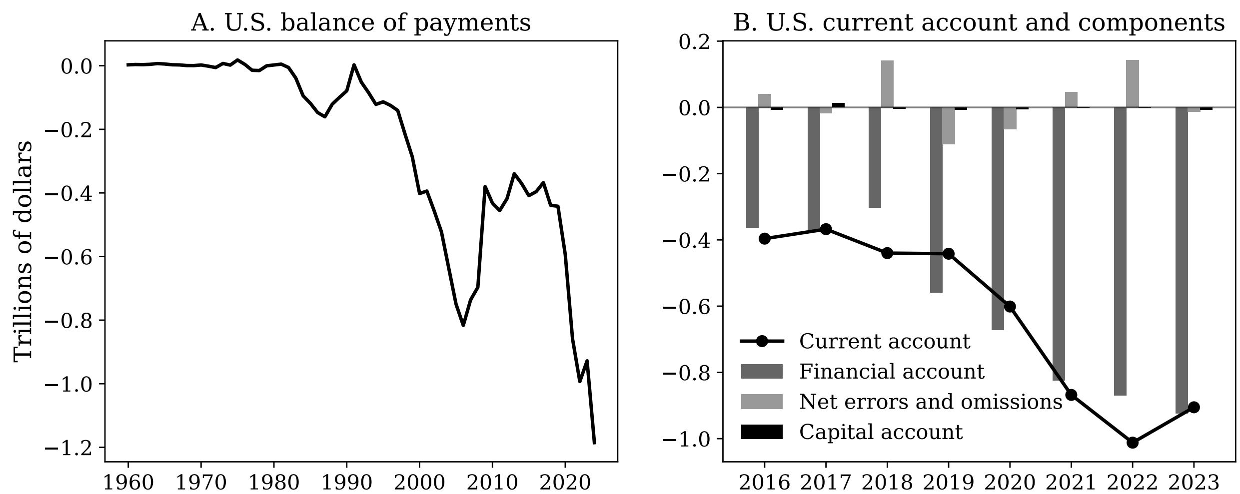

Figure 23.1: Balance of payments (left) and current, capital, and financial accounts (right) in the U.S.

Source: BEA International Transactions and IMF Data, Balance of Payments.

From national-accounting identities, we know that national saving satisfies $S = Y - G - C = I + NX$. Because $CA \simeq NX$, a current-account deficit is equivalent to a trade deficit. And since $CA \simeq -FA$, this implies that trade deficits are financed by capital inflows, as a country borrows from abroad. So, the negative current-account balance seen in the right-panel of Figure 23.1 indicates both that the U.S. has been importing more goods than it has been exporting and that it has been saving less than it has been investing domestically.

23.1.2 Cross-country differences in GDP per capita and size

Relationships between trade balances, savings and investment also shed light on broader patterns across countries, as nations with persistent trade surpluses or deficits often experience distinct trajectories of economic growth, resource allocation and external borrowing or lending. These patterns are influenced by trade specialization, where countries focus on producing goods and services in which they have a comparative advantage. This specialization shapes international trade flows, affects income levels, and contributes to the heterogeneity

[6] The plot in the left panel was downloaded from https://www.bea.gov/data/intl-trade-investment/international-transactions and the one in the right panel from https://data.imf.org/?sk=7a51304b-6426-40c0-83dd-ca473ca1fd52&sid=1484252556980, accessed in December 2024.

observed in GDP per capita across nations. Based on these differences, the IMF categorizes countries into three main groups: 'Advanced Economies' (AE), 'Emerging Markets' (EM), and 'Low-income and Developing Countries' (LI). In 2023, the average GDP per capita (in real 2015 U.S. dollars) among AEs was almost \$90,000, whereas in EMs it stood at slightly above \$10,000. LI countries had an average income per capita below \$2,000. Wide disparities between these groups have been persistent over time, as Chapter 3 discussed.

Another critical dimension of heterogeneity is country size. In 2024, the global population was approximately 8.14 billion. The two most populated nations were India (1.45 billion) and China (1.41 billion), followed by the U.S. (340 million). Europe collectively had a population of around 745 million, while Africa's population stood at approximately 1.5 billion. Countries with a relatively low GDP per capita can nonetheless have a large weight in the global economy due to their population size.

Country size has notable influence on economic interactions between regions, whether in trade of goods and services or financial-asset transactions. Larger economies can have market power, hence their macroeconomic dynamics can influence global interest rates, terms of trade and exchange rates. In contrast, smaller economies are typically price takers in international markets, responding to, rather than shaping, global economic conditions. These differences are reflected in the open-economy academic literature, which is divided into two primary strands. One branch addresses small-open economies, emphasizing their unique dynamics and constraints. These models will be explored in the next chapter. The other branch focuses on large economies and often employs models involving two countries to capture their significant mutual interactions. This approach will be the focus of this chapter, where we analyze characteristics of international business cycles (i.e., fluctuations around trends) for large, advanced countries.[7]

23.2 International business-cycle facts

Open-economy business-cycle theory seeks to address fundamental questions about the behavior of aggregate economic variables within and across nations. How volatile is output? Are consumption, investment, and exports pro-cyclical or counter-cyclical? Are economic booms and recessions typically associated with trade balance deficits or surpluses? Are emerging markets different from advanced economies?

23.2.1 Macro variables

To explore business-cycle patterns globally, we extend the methodology introduced in Chapter 14, where similar statistics were computed for the U.S. We analyze a broad set of countries using country-specific data from the World Bank Development Indicators. Our panel consists of 96 countries with annual data spanning the period 1980-2023. Variable definitions,

[7] As shown by Galí and Monacelli (2005), one can derive a small-open economy version of the model in this chapter by taking limits with respect to country size, specifying one country as infinitesimally small in the global economy for which foreign quantities and prices are taken as given. Because of product specialization, however, as explained below, the country retains monopoly power on its terms of trade, i.e., it is not infinitesimal in the market for its own output.

countries included, and replication codes are provided in the book's GitHub page.[8]

To characterize the 'average' world business cycle, we compute business-cycle statistics for individual countries and aggregate them using population-weighted averages. The main statistics are presented in the second column of Table 23.1. The first column includes the values for the U.S., for comparison.[9] The table also includes the average trade-balance-to-output ratio, $tb/y = (x - m)/y$, and the average openness ratio, $(x + m)/y$, where x denote exports, m imports and y GDP (all expressed in per capita terms).

As Table 23.1 shows, national business cycles are similar to each other in numerous dimensions. In particular:

Fact 1 **Aggregate demand and its components are pro-cyclical**: consumption, investment, exports and imports are all positively correlated with output.

Fact 2 **Persistence**: output and its components—consumption, investment, exports and imports—exhibit strong persistence.

Fact 3 **Ranking of volatilities**: imports, exports and private investment are the most volatile components of GDP, followed by government spending and consumption.

Fact 4 **Trade balances and current accounts are counter-cyclical**: the trade balance, trade-balance-to-output ratio, current account and current-account-to-output ratio are negatively correlated with output. This implies that countries tend to import more than they export during expansions and run trade surpluses during recessions.

The following two facts point to significant differences between AE and EM-LI countries—setting the stage for analysis in the next chapter.

Fact 5 **Business cycles in EM and LI economies are more volatile**: the average standard deviation of output is around 2% in AEs, but the corresponding volatilities for EM and LI countries are, respectively, about 50 and 25% higher.

Fact 6 **Consumption smoothing is lower in LI economies**: private consumption, including durables, is about as volatile as output, or slightly more volatile, in AEs and EMs. This may appear counterintuitive and at odds with consumption smoothing. However, durable goods expenditures, viewed as investment in household capital, account for much of this volatility. Consumption volatility is, however, more than twice the volatility of output in LIs.

Finally, relative to the closed-economy literature, open-macro research is confronted with the issue of interdependence. Not only do countries' business cycles share similar features (i.e., for most, consumption is usually less volatile than investment), they are typically positively correlated across borders, and the correlation is especially high among advanced countries—as summarized by the following fact.

[8] It can be accessed in https://github.com/PhD-Macroeconomics/Codes-and-Data.

[9] These differ somewhat from those presented in Chapter 14 because we here use annual, rather than quarterly data, and the sample period is different in order to support cross-country comparison.

Table 23.1: HP-filtered Business Cycles Around the World

Statistic	U.S.	All Economies	Advanced	Emerging Markets	Low-Income Countries
Standard Deviations					
σ_y	1.96	2.89	2.13	3.10	2.41
σ_c/σ_y	0.94	1.24	1.03	1.16	2.35
σ_g/σ_y	1.18	2.02	1.17	1.91	4.50
σ_i/σ_y	4.55	3.04	3.46	2.76	4.77
σ_x/σ_y	4.11	3.91	3.43	3.55	7.87
σ_m/σ_y	3.47	4.08	3.40	3.87	7.19
$\sigma_{tb/y}$	0.61	1.75	1.07	1.82	2.41
$\sigma_{ca/y}$	0.72	1.79	1.20	1.82	2.51
Correlations with y					
c	0.93	0.67	0.84	0.66	0.41
g	0.02	0.33	0.04	0.38	0.35
i	0.65	0.70	0.76	0.71	0.47
x	0.17	0.17	0.32	0.13	0.30
m	0.67	0.40	0.62	0.37	0.35
tb/y	-0.59	-0.27	-0.41	-0.27	-0.09
tb	-0.62	-0.28	-0.41	-0.27	-0.15
ca/y	-0.53	-0.31	-0.36	-0.32	-0.09
ca	-0.59	-0.28	-0.38	-0.28	-0.10
Auto correlations					
y	0.51	0.54	0.47	0.55	0.49
c	0.53	0.47	0.45	0.49	0.37
g/y	-0.63	-0.21	-0.61	-0.17	0.09
i	0.21	0.51	0.47	0.53	0.43
x	0.55	0.44	0.43	0.46	0.36
m	0.23	0.40	0.29	0.43	0.39
tb/y	0.62	0.41	0.50	0.41	0.20
CA/y	0.67	0.40	0.51	0.40	0.15
Means					
tb/y	-2.3	-1.0	-0.7	-0.1	-9.2
$(x+m)/y$	22.0	28.1	32.5	26.5	35.3

Notes: The variables y, c, g, i, x, m, $tb \equiv (x-m)$, and ca denote GDP, total private consumption, government spending, investment, exports, imports, the trade balance, and the current account, respectively. Variables are expressed in real per capita terms. The variables y, c, g, i, x, m are Hodrick-Prescott-filtered ($\lambda = 100$) in logs and expressed as percentage deviations from trend. The variables tb/y, g/y, and ca/y are HP-filtered in levels. The variables tb and ca are detrended first using the long-term component of y and then HP filtered. There are 96 countries in the sample covering, approximately, the interval 1980-2023 at annual frequency. Moments are averaged across countries using long-run population weights. Countries included in Advanced Economies, Emerging Markets, and Developing Countries follow the IMF definition of these groups. The lists of countries, individual statistics, and replication material are provided the book's GitHub page: GitHub page

.

Fact 7 **Output co-moves positively across borders**: business cycles tend to be synchronized across AE countries.

Figure 23.2 illustrates this by plotting the deviations from trend in per-capita GDP for the

U.S. (solid line), Canada (dashed line) and the United Kingdom (dotted line) between 1970 and 2023. The detrended GDP series for these countries move closely in tandem, reflecting the interconnected nature of their economies and the global transmission of economic shocks. Table 23.2 shows the average correlation of output, consumption, investment and prices between the U.S. and a subset of developed economies between 1971 and 2018, further illustrating synchronicity in international business cycles.

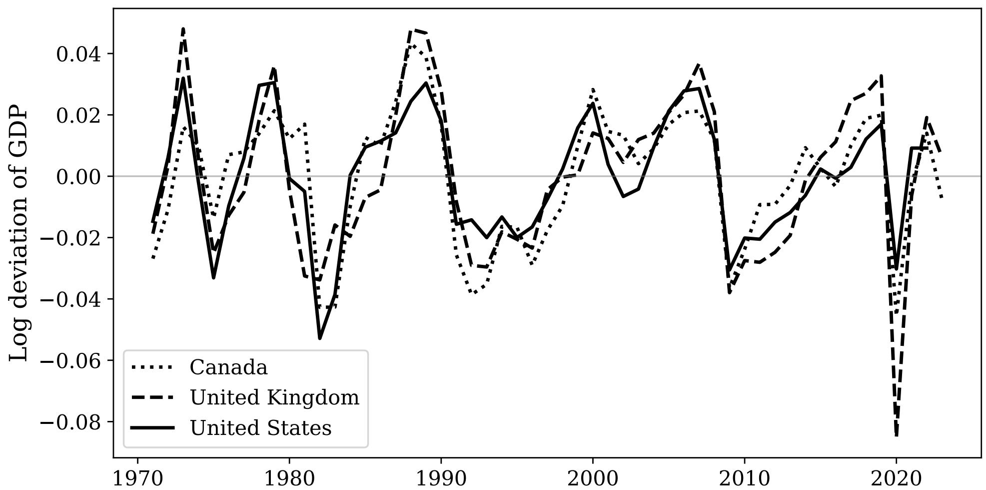

Figure 23.2: Log-deviations in GDP per capita (HP-filtered)

Notes: Hodrick-Prescott filtered log GDP, using $\lambda = 100$ for this annual data.

Table 23.2: Synchronized business cycles

Statistic	1971:Q3-2018:Q2	1971:Q3-2007:Q4
Correlations (U.S. vs ROW)		
GDP	0.58	0.56
C	0.46	0.44
I	0.56	0.48
CPI	0.62	0.57
Backus-Smith Correlation (U.S. vs ROW)		
Rel C vs RER	-0.49	-0.53

Notes: ROW includes Aus, Can, Fra, Ger, Ire, Ita, Jap, Swe, UK. All data detrended over 1971:Q3-2018:Q2 period using Hodrick-Prescott filter with $\lambda = 1600$. Source: update of Corsetti, Dedola, and Leduc (2008b).

Cross-border business-cycle correlations have long challenged the open-macro literature. At country level, TFP, one of the key drivers of business cycles, exhibits a low degree of synchronization.[10] In order to match business-cycle co-movements, quantitative models

[10] As shown by Backus, Kehoe, and Kydland (1992) and subsequent literature, joint estimation of TFP among advanced countries typically yields a positive but small covariance of contemporaneous shocks, and

therefore need to embed a sufficiently strong transmission mechanism for country-specific productivity innovations. We discuss the evidence on this mechanism at the end of this section, looking at the effects of identified TFP shocks in U.S. tradables production on macro variables, including output, in other G10 economies.[11]

International financial and economic integration puzzles. Before proceeding, we highlight four "puzzles" concerning the role of financial and goods-market integration in shaping the international business cycle.

A long-standing debate in open macro concerns the high correlation between national saving and investment. From the perspective of a world with perfect capital mobility, this appears counterfactual: investment should be higher in countries where returns are high—with capital flows naturally supplementing national saving. This is dubbed the "Feldstein-Horioka puzzle." While a high correlation may be produced by shocks that move saving and investment in tandem, the puzzle is commonly used to highlight the incidence of frictions in international capital mobility and/or goods trade—both ultimately impeding intertemporal trade and risk sharing.

A related debate revolves around the "home bias in equity portfolios puzzle." This is based on evidence that investors exhibit a strong preference for domestic assets, despite opportunities for international diversification. The question is whether home bias is a feature of optimal portfolio allocation, as opposed to being driven by trading and information frictions.

By the same token, one may note that consumption patterns across countries are less correlated than predicted by models that assume frictionless international risk-sharing. This discrepancy, known as the "consumption correlation puzzle," also suggests the need to investigate the nature and incidence of frictions in risk sharing, as opposed to fundamental determinants of portfolio diversification (e.g., different preferences or risk exposures).

Lastly, trade within countries significantly exceeds international trade, even among highly integrated economic regions. This phenomenon, known as the "home bias in trade puzzle," requires models of international trade to incorporate and calibrate parameters that replicate the home bias in demand. Obstfeld and Rogoff (2001) provide a thorough discussion of each of these puzzles, and more.

23.2.2 Exchange rates and relative prices

We now introduce international relative prices. While the analysis in this chapter is developed in real terms, it is useful to start with the definition of the nominal exchange rate (NER), denoted by $\mathcal{E}_t$. The NER represents the relative price of currencies, defined as the number of home currency units required for one unit of foreign currency. For example, when the home country is the U.S. and the foreign country is the EU, our notation means that one euro can be bought with $\mathcal{E}_t$ U.S. dollars. An *increase* in $\mathcal{E}_t$ represents a nominal *depreciation*

a positive but moderate spillover coefficient, by which positive innovations in a country drive TFP abroad with a delay.

[11] As further discussed below, the literature distinguishes between tradable goods, which can be exported and imported by a country, and non-tradable goods. The latter are generally defined as goods that are not traded internationally, due to prohibitively high transportation costs relative to their value added, (see, e.g., Obstfeld and Rogoff, 1996, p.199).

of the dollar (home currency), because more dollars are necessary to buy one euro, whereas a *decrease* in $\mathcal{E}_t$ indicates a dollar *appreciation*. In contrast with *bilateral* exchange rates, *trade-weighted* exchange rates (also known as *effective* exchange rate) offer a broader measure reflecting the value of a currency relative to a basket of foreign currencies (e.g., taking into account the yen, the renminbi, etc.).

The *Law of One Price* (LOOP) states that a good should cost the same abroad as at home when expressed in a common currency. Formally, LOOP holds for a specific good i if

$$P_{i,t} = \mathcal{E}_t P_{i,t}^*,$$

where $P_{i,t}$ is the price of good i in dollars at time t and $P_{i,t}^*$ the price of the same good in euros. Hence, LOOP implies that the cost of the good in euros multiplied by the NER must equal the cost in dollars. For example, if a computer can sell at $P_{i,t}$ dollars in the U.S. and $P_{i,t}^*$ euros in Europe, a U.S. producer can sell it in Europe, exchange the euros for dollars and receive $\mathcal{E}_t P_{i,t}^*$ dollars. While LOOP generally holds well for tradable goods, such as commodities and luxury consumer items, it often breaks down for non-tradable goods, like consumer services, housing, transportation and utilities (Schmitt-Grohé, Uribe, and Woodford, 2022).

Purchasing Power Parity (PPP) is a similar concept, said to hold when the cost of a representative basket of goods in each country is identical. That is, PPP extends LOOP from individual goods to a broad basket of goods representative of households' actual consumption, reflecting the relative cost of living between countries. Mathematically, it is expressed as

$$P_t = \mathcal{E}_t P_t^*.$$

Note that if consumption baskets are the same in both countries, LOOP implies PPP. In general though, there are large deviations from absolute PPP in the data. Deviations from PPP arise from factors such as home bias (the difference in baskets), the existence of non-tradable goods and LOOP violations for traded goods. In the model that we present later, we focus on home bias as the main source of discrepancy between LOOP and PPP.

The *real exchange rate* (RER), denoted by $\mathcal{Q}_t$, reflects the relative cost of goods between a foreign and the home country, capturing how the price of a representative basket of goods abroad evolves over time relative to a basket of goods at home. It is defined as the nominal exchange rate adjusted for relative price levels at home P_t and abroad P_t^*,

$$\mathcal{Q}_t = \frac{\mathcal{E}_t P_t^*}{P_t}.$$

The representative basket at home costs P_t dollars and the representative basket abroad (which may have different components) P_t^* euros (or $\mathcal{E}_t P_t^*$ dollars). An increase in $\mathcal{Q}_t$ means that it becomes cheaper (in dollars) to buy the home basket compared to the foreign basket. In other words, the home currency can buy relatively less of the representative basket abroad. An increase in $\mathcal{Q}_t$ therefore reflects a *real depreciation* for the home country. As the equations above show, *absolute PPP* holds when $\mathcal{Q}_t = 1$.[12]

Most studies of PPP emphasize changes in the RER rather than its absolute level. This focus is primarily because changes in the RER can be calculated using consumer price indices

[12]Note that whereas we can use the nominal exchange rate $\mathcal{E}_t$ with a unit ("$\mathcal{E}_t$ dollars"), the real exchange rate does not have a unit (it is a relative cost of two baskets, both measured in a common currency).

(CPIs), which are widely available for many countries at relatively high frequencies, typically on a monthly basis. While the CPI provides information about the price of a basket of goods only up to a scalar, it effectively measures how the price of the basket changes over time. Mathematically, the RER can be decomposed into (log) changes in the NER and inflation differentials between countries:

$$\Delta q_t = \Delta e_t + \pi_t^* - \pi_t,$$

where $\Delta q_t = \ln \mathcal{Q}_t - \ln \mathcal{Q}_{t-1}$, $\Delta e_t = \ln \mathcal{E}_t - \ln \mathcal{E}_{t-1}$, and $\pi_t^{(*)} = \log P_t^{(*)} - \log P_{t-1}^{(*)}$ is home (foreign) CPI inflation. *Relative PPP* holds if $\Delta q_t = 0$ or $\Delta e_t = \pi_t - \pi_t^*$; that is if the depreciation rate of a country's currency against the domestic currency, Δe_t, matches the inflation differential between the foreign country and the domestic one.

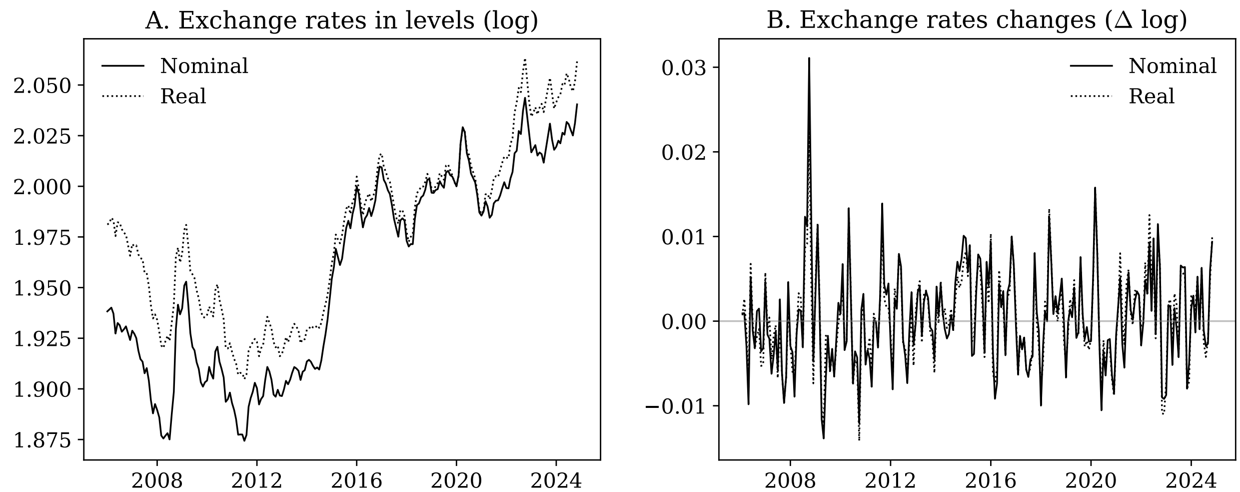

Figure 23.3: Nominal and Real U.S. Exchange Rates, in levels (left) and changes (right)

Source: FRED. The NER corresponds to the series DTWEXBGS in FRED, "Nominal Broad U.S. Dollar Index, Index Jan 2020=100, Monthly, Not Seasonally Adjusted." The RER corresponds to the series RBUSBIS, "Real Broad Effective Exchange Rate for United States, Index Jan 2020=100, Monthly, Not Seasonally Adjusted."

In the data, particularly for major currencies, the RER closely tracks the NER, as Figure 23.3 shows. In general, exchange rates in nominal and real terms are more volatile than macroeconomic aggregates, but less volatile than other financial market variables. Moreover, exchange-rate fluctuations exhibit no robust contemporaneous correlation with macroeconomic fundamentals (such as the money supply, interest rates and output levels). This fact, which has been labeled the "exchange-rate disconnect", has long challenged both macroeconomic and financial models. In the same vein, Meese and Rogoff (1983) famously established that empirical exchange-rate models cannot outperform the random walk in out-of-sample forecasting of exchange rates—the "Meese and Rogoff puzzle."

The exchange rate also displays a high degree of persistence, with weak mean reversion, as the left panel of Figure 23.3 demonstrates. This slow adjustment is commonly referred to as the "PPP puzzle." The right panel of the figure highlights the substantial volatility of the RER, indicating that PPP does not hold in the short run. However, the tendency of RER changes to center around zero suggests that PPP holds in the long run.[13] Studies have shown that the RER is almost an order of magnitude more volatile than macroeconomic fundamentals such as inflation, consumption, and output. Moreover, it tends to be negatively correlated (if only weakly) with the ratio of domestic to foreign consumption. A negative correlation between the RER and the home-to-foreign relative consumption is puzzling: when the RER is appreciated, home goods are relatively expensive; if this reflects scarcity of the home good, one would expect home consumption to be lower, not higher, than foreign—the "Backus-Smith puzzle." For the U.S., the puzzle is illustrated by Figure 23.4, with a correlation between RER and relative consumption of around -0.5 (see also bottom row of Table 23.2).[14]

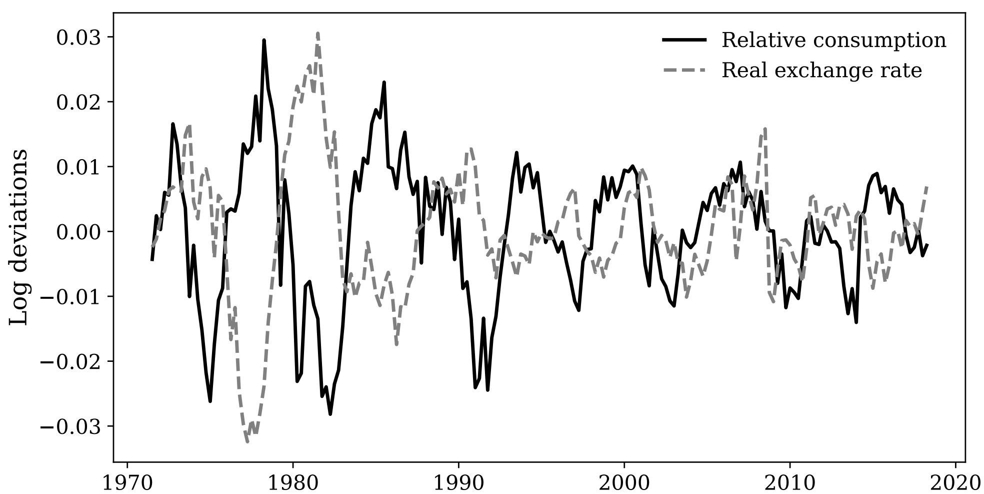

Figure 23.4: The U.S. real exchange rate and relative consumption (HP-filtered)

Source: update of Corsetti et al. (2008b). U.S. real exchange rate and relative consumption constructed relative to remaining PPP-weighted G10 economies. Hodrick-Prescott filtered series, using $\lambda = 1600$ for this quarterly data.

The "terms of trade" (TOT), denoted by $\mathcal{T}_t$, is the final important international relative price. It measures the relative price of the import basket ($P_{F,t}$, price of imported foreign goods at home) to the export basket ($P^*_{H,t}$, price of exported home goods abroad):

$$\mathcal{T}_t = \frac{P_{F,t}}{P^*_{H,t}\mathcal{E}_t},$$

[13]This is also true if the RER are computed using PPI- or wage-based price indexes. Empirical studies suggest that the RER is stationary, suggesting that despite large fluctuations at business-cycle frequencies, there is some version of PPP anchoring its long-run level (Taylor, 2002; Taylor and Taylor, 2004).

[14]The "Mussa Puzzle," first documented by Mussa (1986), highlights a dramatic increase in the volatility of nominal and real exchange rates following the transition of industrial countries from fixed to floating exchange rate regimes after the collapse of the Bretton Woods System in 1973. The rise in volatility cannot fully be explained by changes in domestic price levels, suggesting a disconnect between exchange rate movements and inflation. Since we focus on a real economy in this chapter, we will not have much to say about this.

where $P_{F,t}$ ($P^*_{H,t}$) is the home (foreign) import price index in local currency (that is, $P_{F,t}$ is measured in dollars and $P^*_{H,t}$ is measured in euros). The TOT and RER differ in that the import basket is different from the representative consumption basket at home and the export basket is different from the representative consumption basket abroad. In other words, while the RER reflects deviations in purchasing power, the TOT captures actual relative prices of traded goods. Empirically, the RER and TOT are only weakly correlated, with the TOT being less volatile.

23.2.3 Cross-country transmission of productivity shocks

We conclude this section with conditional evidence, which sets the stage for the rest of the chapter. Specifically, we complement the business-cycle statistics discussed so far with evidence on the cross-country transmission of technology improvements, highlighting how advancements in one country affect relative quantities, prices, and external balances. Figure 23.5 presents estimated impulse responses from Corsetti, Dedola, and Leduc (2014), which empirically examine the cross-country transmission of U.S. supply-side shocks relative to the rest of the world. This analysis employs a sign-restricted Bayesian vector auto-regression methodology. The figure illustrates the responses of U.S. relative quantities (manufacturing output and consumption), relative prices (real exchange rate and terms of trade), and external balances (net exports and net foreign assets) to a temporary increase in U.S. labor productivity in the tradable sector, *vis-à-vis* other G10 economies.

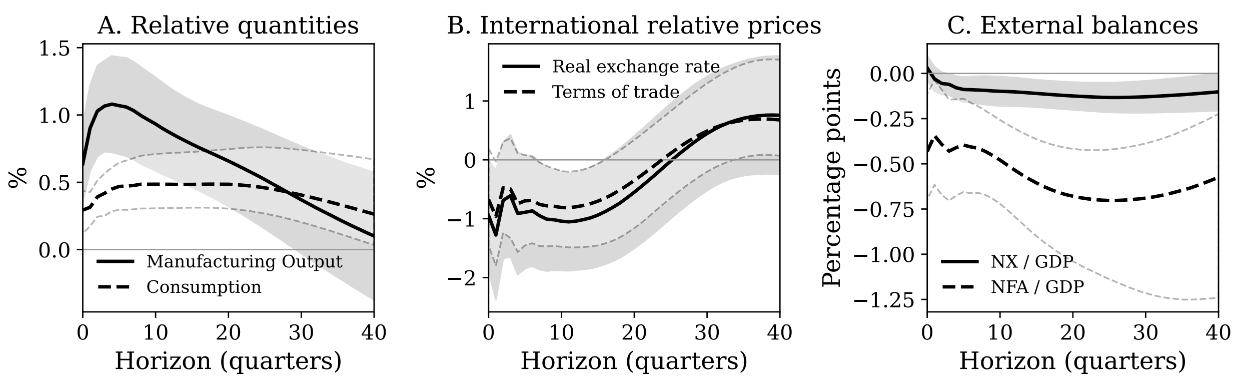

Figure 23.5: Empirical estimates of the cross-country transmission of productivity shocks: U.S. to rest of world

Source: Corsetti, Dedola, and Leduc (2014) (Figures 1 and 2). *Notes*: estimated impulse response to a positive U.S. labor productivity shock. Charts report median response of U.S. variables relative to G10 rest of the world, as well as the 16th and 84th percentiles of the posterior distribution, satisfying the sign restrictions detailed in Section 2.1 of Corsetti et al. (2014). Rest of the world is PPP-weighted and comprises the following advanced economies: Japan, Germany, UK, Italy, France, Canada, Australia, Sweden, and Ireland. RER = real exchange rate; TOT = terms of trade; NX/GDP = net exports to GDP; NFA/GDP = net foreign assets to GDP. Estimation sample: 1973Q1-2004Q4.

Although the U.S. and other G10 business cycles are to a large extent synchronized, the estimated responses reveal that a relative productivity improvement in the U.S. leads to an increase in U.S. tradable output *relative* to the rest of the world; this increase in relative output goes hand in hand with an increase in relative consumption and appreciation of the domestic real exchange rate (the ratio of rest-of-the-world to U.S. consumer prices declines) and TOT, such that imported goods are relatively cheaper. External balances—net exports and net-foreign-asset positions—are counter-cyclical; conditional on the shock, the domestic economy imports and consumes more today, funded by international borrowing.

These patterns have been further explored by subsequent literature, relying on different methodologies and/or corroborating analyses of identified contemporaneous, *but persistent*, productivity shocks with analysis identifying 'news shocks' to productivity (see, e.g., Nam and Wang, 2015; Chahrour, Cormun, De Leo, Guerrón-Quintana, and Valchev, 2024). In the closed-economy literature, leading studies developed after the Covid-19 pandemic have recently elaborated on the notion of "Keynesian Supply shocks" (Guerrieri, Lorenzoni, Straub, and Werning, 2022), similar to the mechanism explained in this chapter (i.e., the idea that productivity shocks move both supply and demand, and under certain conditions, the demand effects can be dominant).

The pattern identified by the empirical estimates highlight the two key stylized facts that challenge open-macro theory: (1) net exports and the current account (external saving) tend to be counter-cyclical; and (2) international relative prices (i.e., the RER and TOT) tend to appreciate when output and consumption are relatively high. These empirical regularities form the central puzzles motivating recent open-economy macroeconomics. In what follows, we show how to bring modern international business-cycle theory closer in line with this evidence.

23.3 The workhorse open-economy model

We do so by presenting a two-country macro model, widely used to study the cross-border transmission of country-specific shocks. Our focus will be on productivity improvements. Throughout, we will confront the model with the stylized facts presented in the previous section—namely the pro-cyclicality of external deficits, and the co-movements of output, demand and exchange rates. In this section, we simplify our analysis by modeling endowment economies; in the next, we generalize our results to economies with endogenous production; in the third, we synthesize the main conclusions with three numerical exercises.

23.3.1 Model setup: preferences and technology

The model world economy consists of two countries: home H and foreign F, each populated by a continuum of identical households with unit mass. Throughout, we label foreign variables with an asterisk. We assume, for simplicity, that countries are of equal size. A representative home household receives an exogenous endowment of (tradable) home goods, $Y_{H,t}$, while the foreign agent has an exogenous endowment of the foreign good, $Y^*_{F,t}$. For consumers, the home and foreign goods are imperfect substitutes, so there is trade in goods internationally. In the global equilibrium, the separate markets for home and foreign goods, along with the markets for internationally traded assets, must simultaneously clear.

Household utility. In each country, the representative household derives utility from consuming a combination of domestic and imported goods. In the home country, aggregate consumption is denoted by C_t. The expected lifetime utility of the representative home household U_t is:

$$U_t = \mathbb{E}_t \left[\sum_{j=0}^{\infty} \beta^j u(C_{t+j}) \right] \tag{23.1}$$

where $\beta \in (0,1)$ is the discount factor and $u(C) = (C^{1-\sigma} - 1)/(1-\sigma)$ is instantaneous utility from consumption (with $\sigma > 0$ representing the coefficient of relative risk aversion). Preferences of the representative foreign consumer are defined analogously, with consumption C_t^*, expected lifetime utility U_t^*, as well as $\beta^* = \beta$ and $\sigma^* = \sigma$.

The final consumption basket C_t is a CES aggregator of home, $c_{H,t}$, and foreign goods, $c_{F,t}$, known as the 'Armington aggregator' (Armington, 1969), of the form:

$$C_t = \left[a_H^{\frac{1}{\phi}} c_{H,t}^{\frac{\phi-1}{\phi}} + a_F^{\frac{1}{\phi}} c_{F,t}^{\frac{\phi-1}{\phi}} \right]^{\frac{\phi}{\phi-1}}. \tag{23.2}$$

This allows for imperfect substitutability between domestic and foreign goods—encompassing both the case of complementarity and substitutability—as well as 'home bias' in consumption, in line with the stylized fact that a greater proportion of domestic expenditure is devoted to domestic goods. With $a_H + a_F = 1$, the parameter a_H (a_F) represents the steady-state share of the home (foreign) goods in home expenditure such that when $a_H \in (0.5, 1]$ we have home bias in consumption. The parameter $\phi > 0$ denotes the constant elasticity of substitution between home and foreign goods, with higher (lower) values representing greater substitutability (complementarity) between home and foreign goods. As a result, this parameter is often referred to as the 'trade elasticity'.[15]

Given separable and time-additive preferences in equation (23.1), we can solve the household optimization problem by separating its intertemporal and intratemporal dimensions. As shown below, the representative household's intertemporal preferences will pin down aggregate consumption and savings, once the structure of international financial markets is defined. Taking aggregate consumption C_t as given, one can determine the demand for home and foreign goods, and derive a welfare-based price index of consumption, solving an intratemporal optimization in each period.

Let $p_{H,t}$ and $p_{F,t}$ represent the price of home and foreign goods, respectively, in the home economy. Consider a level of expenditures $\mathbf{E}_t = P_t C_t$ on home and foreign goods

$$\mathbf{E}_t = p_{H,t} c_{H,t} + p_{F,t} c_{F,t}.$$

The intratemporal decision of households involves minimizing their expenditure with respect to $c_{H,t}$ and $c_{F,t}$ subject to the Armington aggregator.

Solving this, home demand functions for home and foreign goods are, respectively:

$$c_{H,t} = a_H \left(\frac{p_{H,t}}{P_t} \right)^{-\phi} C_t \quad \text{and} \quad c_{F,t} = a_F \left(\frac{p_{F,t}}{P_t} \right)^{-\phi} C_t, \tag{23.3}$$

[15] For given aggregate consumption C_t, one can show that home and foreign goods are substitutes (complements) when $\phi > 1$ ($\phi < 1$) by calculating the second-order cross-partial derivatives of equation (23.2), $(\partial^2 C_t)/(\partial c_{H,t} \partial c_{F,t})$.

where the home CPI is defined as:

$$P_t = \left(a_H p_{H,t}^{1-\phi} + a_F p_{F,t}^{1-\phi}\right)^{\frac{1}{1-\phi}}. \tag{23.4}$$

The foreign CPI P_t^* is similarly defined, with $p_{H,t}^*$ and $p_{F,t}^*$ denoting the prices of home and foreign goods in the foreign economy.[16]

Market clearing. Given that the world is a closed endowment economy, equilibrium involves market clearing for each good: $Y_{H,t} = c_{H,t} + c_{H,t}^*$ and $Y_{F,t} = c_{F,t} + c_{F,t}^*$ for all t.

Exchange rates and relative prices. Since we abstract from nominal rigidities, we normalize $\mathcal{E}_t$ to 1 and focus, instead, on relative prices: the TOT and RER. We also assume that the LOOP holds: $p_{i,t} = \mathcal{E}_t p_{i,t}^* = p_{i,t}^*$ for $i = \{H, F\}$. The TOT from the perspective of the home country is then:

$$\mathcal{T}_t = \frac{p_{F,t}}{p_{H,t}}. \tag{23.5}$$

Hence, $\mathcal{T}_t$ goes up when foreign imports become relatively more expensive, corresponding to a worsening in the home TOT. From the perspective of consumers, taking the TOT as given, the intratemporal consumption demand for domestic and foreign goods $\{c_{H,t}, c_{F,t}, c_{H,t}^*, c_{F,t}^*\}$ arising from the ratio of the equations in (23.3) must satisfy:

$$\frac{1}{\mathcal{T}_t} = \left(\frac{a_H}{a_F}\right)^{\frac{1}{\phi}} \left(\frac{c_{H,t}}{c_{F,t}}\right)^{-\frac{1}{\phi}} = \left(\frac{a_H^*}{a_F^*}\right)^{\frac{1}{\phi}} \left(\frac{c_{H,t}^*}{c_{F,t}^*}\right)^{-\frac{1}{\phi}}. \tag{23.6}$$

Given the normalization of $\mathcal{E}_t = 1$, the RER $\mathcal{Q}_t$ satisfies

$$\mathcal{Q}_t = \frac{\mathcal{E}_t P_t^*}{P_t} = \frac{P_t^*}{P_t}. \tag{23.7}$$

An increase in $\mathcal{Q}_t$ is a home real depreciation, with foreign consumption becoming more expensive relative to home consumption. Importantly, although the LOOP is assumed to hold for each tradable good, the LOOP is not sufficient for the price of consumption bundles to be equal across countries (i.e., for PPP to hold). PPP requires that $P_t = \mathcal{E}_t P_t^* = P_t^*$, which would imply $\mathcal{Q}_t = 1$, a condition regularly rejected in the data, as explained in the previous section. An attractive feature of Armington aggregation is that absolute PPP will not generally hold. It will only hold if two conditions are satisfied: (i) LOOP holding in both goods markets, and (ii) consumption baskets being identical across countries.[17]

[16] To be clear, we use H to denote a specific good. This means it is "home-produced" for the domestic economy but actually "foreign-produced" from the perspective of the other country. Thus, a_H^* is the CES weight on the H good for the foreign economy. If countries are symmetric and have a stronger taste for their own good, then $a_H = a_F^* > a_H^* = a_F$.

[17] In addition to accounting for home bias in consumption, as we do here, one can also model deviations from PPP by: (i) modeling deviations from the LOOP, for example by allowing 'pricing-to-market' by firms, or (ii) allowing for non-tradable goods within the model, for which the LOOP will not necessarily hold.

By using the RER definition (23.7), home CPI (23.4) and the analogous expression for foreign CPI, we can express a relationship between the RER and TOT as:

$$\mathcal{Q}_t^{1-\phi} = \frac{a_H^* + a_F^* \mathcal{T}_t^{1-\phi}}{a_H + a_F \mathcal{T}_t^{1-\phi}}. \tag{23.8}$$

According to the model, the co-movement between the RER and TOT will depend on the degree of home bias. If $a_H = a_F^* > 0.5$, then the co-movement will be positive; while if $a_H = a_F^* = 0.5$ then PPP will hold and fluctuations in TOT will not translate into fluctuations in RER. Counterfactually, away from PPP, the model predicts a perfect correlation between $\mathcal{Q}$ and $\mathcal{T}$. Model specifications that allow for (plausible) differences between consumption baskets—e.g., via the presence of non-traded goods and distribution costs (see Benigno and Thoenissen, 2008; Corsetti, Dedola, and Leduc, 2008a)—can resolve this.

So far, we have not characterized the consumption-saving decision by households. Specifically, in open economies, since domestic aggregate consumption can deviate from domestic output via trade in financial assets (i.e., in the model, $Y_{H,t}$ need not equal C_t), the structure of financial markets is crucial for determining the global equilibrium, including the exchange-rate determination and capital flows across borders. To close the model, a specification for international financial markets is needed.

23.3.2 International financial markets and intertemporal choices

To understand the role of international financial markets, and their degree of 'openness', in shaping the global equilibrium, we contrast three benchmark cases (akin to those in Chapter 5), that differ in the extent to which country specific risks are insured (or 'traded'):

1. **Complete markets**: where all risk is traded via a complete set of Arrow-Debreu securities, and the RER is pinned down by the ratio of marginal utilities across borders.

2. **Incomplete markets**: where insurance is partial, so that equilibrium reflects non-traded risks:

   (i) **Financial autarky**: with no markets for international assets, the RER must adjust to balance international trade in goods.

   (ii) **Riskless bonds only**: with trade in a single non-contingent bond, through which households can self insure, the RER is pinned down by the 'uncovered interest parity' (UIP) condition.

We describe each in turn in the remainder of this sub-section. Throughout the analysis we assume that agents commit to repay bonds (regardless of whether they are contingent or not). The commitment assumption will be relaxed in Chapter 24.

Complete markets

Let s_t denote the time-t state of the world and suppose $s_t \in S$, where S is the exogenous part of the state space (i.e., values of the two endowments). If markets are complete—for instance, if there is a full set of Arrow-Debreu securities for each potential state of the world—all risk is traded. The representative home household's intertemporal budget constraint is then:

$$P_t C_t + \int_{s_{t+1}} q_t(s_{t+1}) \mathbb{B}_{H,t}(s_{t+1})\, \mathrm{d}s_{t+1} \leq \underbrace{\mathbb{B}_{H,t-1} + p_{H,t} Y_{H,t}}_{\text{current income}},$$

where $\mathbb{B}_{H,t}(s_{t+1})$ denotes the quantity of Arrow-Debreu securities paying one unit of the home consumption basket upon realization of the state s_{t+1} at time $t+1$, traded at the price $q_t(s_{t+1})$ at time t. For convenience, this notation suppresses the dependence on history before t (the whole history of shocks). For example, $q_t(s_{t+1})$ is the price at time t, given a history up until and including t, of a one-unit payoff of H goods at $t+1$, and for different histories this $q_t(s_{t+1})$ will in general take on different values.[18]

Maximizing the representative home household's expected lifetime utility subject to this budget constraint yields the following Euler equation for each state $s_{t+1} \in S$:

$$u'(C_t) \frac{q_t(s_{t+1})}{P_t} = \Pr(s_{t+1}|s_t) \beta u'(C_{t+1}(s_{t+1})) \frac{1}{P_{t+1}(s_{t+1})}, \tag{23.9}$$

where $\Pr(s_{t+1}|s_t)$ denotes the probability of transitioning from state s_t at time t to s_{t+1} at $t+1$.

Consider guaranteeing one unit of consumption good for a home investor at time $t+1$. For this, at period t, the investor has to buy $P_{t+1}(s_{t+1})$ units of the Arrow-Debreu security for each $s_{t+1} \in S$. This purchase would cost $q_t(s_{t+1})P_{t+1}(s_{t+1})$ in terms of the home basket and in total $\int q_t(s_{t+1})P_{t+1}(s_{t+1})\mathrm{d}s_{t+1}$ units as a bundle. In terms of time-t consumption goods, this bundle costs $\int q_t(s_{t+1})P_{t+1}(s_{t+1})\mathrm{d}s_{t+1}/P_t$ units of the home consumption bundle today. Let us call the gross real return $R_t \geq 1$ the inverse of the cost of this bundle. That is, because $1/R_t$ units of the home consumption bundle today can guarantee one unit of the home consumption bundle tomorrow, it follows that one unit today would guarantee R_t units tomorrow. Using (23.9), that is

$$\frac{1}{R_t} = \frac{1}{P_t} \int q_t(s_{t+1}) P_{t+1}(s_{t+1}) \mathrm{d}s_{t+1} = \int \Pr(s_{t+1}|s_t) \beta \frac{u'(C_{t+1}(s_{t+1}))}{u'(C_t)} \mathrm{d}s_{t+1}.$$

In the foreign economy, there is equivalently a full set of Arrow-Debreu securities with prices q^*, and hence we can define R_t^* to be the number of units of the foreign consumption bundle obtained by a riskless investment of one unit of such a bundle today.[19] Thus the gross real returns R_t and R_t^* for the representative home and foreign households satisfy

$$1 = \beta R_t \mathbb{E}_t \left[\frac{u'(C_{t+1})}{u'(C_t)} \right] = \beta R_t^* \mathbb{E}_t \left[\frac{u'(C_{t+1}^*)}{u'(C_t^*)} \right]. \tag{23.10}$$

Now, consider the following: rather than investing in home Arrow-Debreu securities, first exchange one unit of the home consumption bundle into $1/\mathcal{Q}_t$ units of the foreign bundle.

[18] The notation based on Chapter 7 would therefore have been $q_t(S^{t+1})$, with $S^{t+1} = \{S^t, s_{t+1}\}$.

[19] The existence of Arrow-Debreu markets in both countries means an excess of assets—twice as many as we need to complete markets. We use this assumption for symmetry and will of course invoke the absence of arbitrage, as in Chapter 14, when needed.

Then invest it in the foreign Arrow-Debreu securities and receive $R_t^*/\mathcal{Q}_t$ units of foreign bundle at period $t+1$, which can be converted into $R_t^*\mathcal{Q}_{t+1}/\mathcal{Q}_t$ units of the home bundle. No-arbitrage implies

$$1 = \beta R_t \mathbb{E}_t\left[\frac{u'(C_{t+1})}{u'(C_t)}\right] = \beta R_t^* \mathbb{E}_t\left[\frac{u'(C_{t+1})}{u'(C_t)}\frac{\mathcal{Q}_{t+1}}{\mathcal{Q}_t}\right]. \tag{23.11}$$

The above is a core condition in international finance: the (real) Uncovered Interest Parity (UIP) condition. The expected rate of RER depreciation is a function of (real) interest differentials and expected growth of marginal utilities.

When markets are complete, no-arbitrage implies a stronger condition, holding state by state. Combining home and foreign state-by-state Euler equations yields

$$\frac{u'(C_t)}{u'(C_t^*)}\frac{P_t^*}{P_t} = \frac{u'(C_{t+1}(s_{t+1}))}{u'(C_{t+1}^*(s_{t+1}))}\frac{P_{t+1}^*(s_{t+1})}{P_{t+1}(s_{t+1})} \quad \forall\, t. \tag{23.12}$$

The above equation leads to the key implication of *full risk sharing* under complete markets. Define $t = 0$ as the initial period when the economy starts with the complete set of financial contracts already in place. At any point in time after that, it must be the case that:

$$\frac{u'(C_0)}{u'(C_0^*)}\mathcal{Q}_0 \equiv \kappa_0 = \frac{u'(C_t(s_t))}{u'(C_t^*(s_t))}\mathcal{Q}_t(s_t) \quad \forall\, t.$$

With perfect insurance, the marginal-utility ratio across two countries is pinned down by the wealth distribution—reflected by the difference in consumption—at time 0, and remains constant forever. The initial distribution at time 0 is of course endogenous, as it depends on the equilibrium asset and goods prices, which are in turn a function of relative endowments.[20] If countries are assumed to be initially perfectly symmetric (as they typically are in the workhorse model we consider), then $\kappa_0 = 1$; otherwise it endogenously takes values above or below one that cannot be ignored in the computation of the equilibrium. However, the key is that under complete markets κ (whatever value it takes) is independent of time.

Using our specification of preferences, the *full risk-sharing* condition simplifies to:

$$\mathcal{Q}_t = \kappa_0\left(\frac{C_t}{C_t^*}\right)^{\sigma} \quad \forall t. \tag{23.13}$$

Under complete markets, there is a tight equilibrium relation between relative marginal utilities of consumption and the RER, the implications of which we study in the next section.[21] The full risk-sharing condition states that home consumption can increase relative to foreign consumption only if the home country experiences a real depreciation (i.e., domestic consumption becomes relatively cheaper).[22] This outcome is ensured by *ex post* contingent

[20]See Corsetti, Lipinska, and Lombardo (2025) for an analysis of risk sharing that distinguishes between the effects on consumption smoothing and relative wealth, driven by the repricing of assets.

[21]The correlation between relative consumptions and RER is 1 according to the model, but is negative in the data (see Figure 23.4 and Table 23.2), illustrating the Backus-Smith puzzle.

[22]In this chapter we abstract from (time) preference shocks, under which the condition includes an additional term as a function of this shock.

transfers across borders. From a welfare perspective, these transfers are efficient: a low price indicates that the home consumption good is relatively more abundant. As long as there are no other distortions (e.g., nominal rigidities), the complete-market specification of our real economy is first-best from the perspective of a global planner.

Closed-form solutions for prices and quantities can be computed as follows. Within a given period, normalize prices to $P_t = 1$, take C_t and C_t^* as given, and then combine the demand functions (23.3) with market clearing conditions for each good. Together with (23.13), this leaves the same number of equations as unknowns, with κ_0 determined by relative wealth/income.[23]

Financial autarky

Define the trade balance as the value of exports minus value of imports. In units of home consumption this is:

$$tb_t \equiv \frac{p_{F,t}}{P_t} c_{F,t} - \frac{p_{H,t}}{P_t} c^*_{H,t}. \tag{23.14}$$

Under perfect risk sharing, any non-zero (real) trade balance would be entirely financed *ex post* by state-contingent payments from Arrow-Debreu securities. At the other extreme, with no trade in international financial markets, *financial autarky*, external trade must balance in each and every period (i.e., $tb_t = 0$):

$$\mathcal{T}_t c_{F,t} - c^*_{H,t} = 0 \iff C_t = \frac{p_{H,t}}{P_t} Y_{H,t}. \tag{23.15}$$

The TOT adjusts in equilibrium to ensure balanced trade. As a result, the expected real depreciation no longer obeys the UIP condition (23.11). Since no risk can be insured through financial contracts, in general any shock will drive a wedge in marginal utilities between home and foreign countries. Without the ability to transfer wealth intertemporally, in effect, the model is the two-country counterpart to closed-economy models with hand-to-mouth consumers. However, in open economies, consumption and output comprise different goods, so that real income crucially depends on the TOT, since from (23.4) we have

$$\left(\frac{p_{H,t}}{P_t}\right)^{1-\phi} = \frac{1}{a_H + (1-a_H)\mathcal{T}_t^{1-\phi}}.$$

Under financial autarky, home households face a static maximization each period:

$$\max_{C_t} \; u(C_t) \quad \text{subject to} \quad P_t C_t = p_{H,t} Y_{H,t}.$$

Therefore, this case can be solved period by period, since there are no borrowing and saving mechanisms.[24]

When markets are incomplete, it turns out to be useful to define a notion of relative wealth levels of the two countries, $\mathcal{W}_t$, as many of the key equilibrium objects depend on how this

[23] In the complete-markets model, the relevant initial wealth levels are total present-value wealth amounts; under financial autarky, discussed next, they are given by $p_{H,t}Y_{H,t}$ and $p_{F,t}Y^*_{F,t}$, respectively.

[24] Note, of course, that domestic and foreign interest rates can be computed as the values that clear the respective saving markets at every period.

notion evolves over time. $\mathcal{W}_t$ is constant over time in the complete-markets model but will evolve stochastically over time in the autarky model and with riskless bonds only. To begin, let us denote by $\chi_t \equiv u'(C_t)/P_t$ the Lagrange multiplier on the budget constraint above. This variable can be interpreted as the marginal utility of increased consumer resources (measured in units of the home consumption basket) and, as the box below argues, will be strictly decreasing in the consumer's current wealth. Given that, we define

$$\mathcal{W}_t \equiv \frac{\chi_t^*}{\chi_t} = \frac{u'(C_t^*)/P_t^*}{u'(C_t)/P_t} = \left(\frac{C_t}{C_t^*}\right)^{\sigma} \frac{1}{\mathcal{Q}_t}. \tag{23.16}$$

Recall our definition of the constant κ under complete markets. We see from the above that $\mathcal{W}_t$ is its inverse (i.e., $1/\kappa_t = \mathcal{W}$), but away from complete markets nothing insures that the ratio of Lagrange multipliers remains constant over time, i.e., non-traded risk causes the analog of κ_0 under complete markets to become state contingent and time varying ($\kappa_t(s)$). As such, it can be interpreted as capturing, at each point in time, the *wealth gap* created by uninsured risk, relative to complete markets.

Lagrange multiplier as the marginal utility of income (wealth)

Consider a static model where a consumer is choosing a vector of consumption goods to maximize $u(C)$, where C is a CES aggregate of goods, subject to a budget $PC = Y$. Here, Y is total income, in units of the CES aggregate, and P is the price defined as the smallest cost of obtaining one unit of the CES aggregate. That is, we have a static version of the model described in equations (23.1)–(23.4) above. For this problem, we can write down the Lagrangian

$$\mathcal{L} \equiv u(C) + \chi\,(Y - PC)\,.$$

The first-order condition is $u'(C) = \chi P$, as in the text. Moreover, quite trivially, the maximum utility equals $u(Y/P)$, which means that the derivative of maximized utility with respect to income, Y, equals $u'(Y/P)/P = \chi$. Given a strictly concave utility function u, therefore, χ is strictly decreasing in income.

The interpretation that the Lagrangian multiplier represents a marginal utility of wealth is also a general result for any utility maximization problem with a standard budget constraint. In a dynamic model, as in the main text, Y/P needs to be replaced by period-t resources minus saving. Thus, given prices, χ_t would represent (the inverse of) this wealth measure. Under financial autarky, that wealth measure is simply the current world market value of the home endowment; in the bond economy, it is the amount the consumer chooses to optimally consume, which in turn is increasing in consumer wealth due to consumption being a normal good.

Hence $\mathcal{W}_t$ represents how rich the home country is compared to the foreign country in equilibrium. When $\mathcal{W}_t > 1$, the home country is wealthier than the foreign country at period t, which is also reflected in a change in the relative demands for the two goods.

Riskless bonds only

In our final case, we model a particular form of market incompleteness, allowing for cross-border trade in discount bonds—a common approach for capturing financial-market incompleteness in the literature. To keep the analysis simple, and relying on our normalization $\mathcal{E}_t = 1$, we posit trade in a single nominal one-period bond B_t, with nominal gross return $(1+i_t)$ and price $Q_t = 1/(1+i_t)$. The budget constraint of the home household is now

$$C_t = \frac{p_{H,t}Y_{H,t} - (Q_t B_t - B_{t-1})}{P_t}.$$

Current consumption can now be financed not only out of the domestic endowment, but also by borrowing (for $B_t < 0$). Therefore, in equilibrium, the trade balance will be $(Q_t B_t - B_{t-1})/P_t$. The current account, defined as the change in the net foreign asset position of the country (in our model, $ca_t = (B_t - B_{t-1})/P_t$), results from the sum of the trade balance and the earnings on net foreign assets. A positive current-account balance means that the home country is accumulating net claims *vis-à-vis* the foreign one.

The household maximizes expected discounted lifetime utility (23.1) subject to

$$p_{H,t+j}Y_{H,t+j} + B_{t+j-1} - P_{t+j}C_{t+j} - Q_{t+j}B_{t+j},$$

a constraint we associate with the Lagrange multiplier χ_{t+j}, as above.[25] This problem yields the following Euler equation for the home household:

$$Q_t = \beta \mathbb{E}_t \left[\frac{u'(C_{t+1})}{u'(C_t)} \frac{P_t}{P_{t+1}} \right] \tag{23.17}$$

and analogously for the representative foreign household (where again we have that $\chi_t = u'(C_t)/P_t$). Under our assumption that there is trade in one asset only, we can always define the home real interest rate as the expected marginal utility growth

$$R_t = \left(\beta \mathbb{E}_t \left[\frac{u'(C_{t+1})}{u'(C_t)} \right] \right)^{-1},$$

Because of our assumption that the NER is identically equal to one, for no arbitrage in the bond market, $Q_t = Q_t^*$ has to hold. Thus,

$$\mathbb{E}_t \left[\frac{u'(C_{t+1})}{u'(C_t)} \frac{P_t}{P_{t+1}} \right] = \mathbb{E}_t \left[\frac{u'(C^*_{t+1})}{u'(C^*_t)} \frac{P^*_t}{P^*_{t+1}} \right], \tag{23.18}$$

which is a version of the real UIP condition. Using the definition of real interest rates, it can be rewritten as follows:

$$(R_t)^{-1}\mathbb{E}_t \left[\frac{P_t}{P_{t+1}} \right] + \mathrm{cov}_t \left(\frac{u'(C_{t+1})}{u'(C_t)}, \frac{P_t}{P_{t+1}} \right) = (R^*_t)^{-1}\mathbb{E}_t \left[\frac{P^*_t}{P^*_{t+1}} \right] + \mathrm{cov}_t \left(\frac{u'(C^*_{t+1})}{u'(C^*_t)}, \frac{P^*_t}{P^*_{t+1}} \right). \tag{23.19}$$

[25] Implicit here is also a borrowing constraint that will not bind in the approximate solution studied below.

This relates real interest rate differentials to expected inflation differentials. Under our normalization of the NER (equal to one), these are tightly linked to expected real depreciation.

A key difference between this 'bond economy', relative to complete markets, is that the state-by-state condition (23.12) no longer holds. With trade in a single bond, shocks still create an *ex-post* marginal-utility wedge, reflecting risk that cannot be traded by intertemporal smoothing, namely:

$$\epsilon_{t+1} \equiv \frac{u'(C_{t+1})}{u'(C_t)}\frac{P_t}{P_{t+1}} - \frac{u'(C^*_{t+1})}{u'(C^*_t)}\frac{P^*_t}{P^*_{t+1}} = \frac{\chi_{t+1}}{\chi_t} - \frac{\chi^*_{t+1}}{\chi^*_t}. \tag{23.20}$$

From the real UIP condition, we see that this wedge must equal zero in expectation: the relative evolution of wealth is random with zero drift in expectation. Unlike under complete markets, these wedges do not cancel state-by-state; only their expectation is zero.

23.3.3 Solving the model

Allowing shocks to emanate from country endowments, we define equilibrium as prices and quantities satisfying the key conditions above: consumer maximization for the two countries and market clearing.

Log-linearization. We log-linearize the model around its symmetric non-stochastic steady state with $\overline{\mathcal{Q}} = \overline{TOT} = 1$, imposing home bias in preferences ($a_H = a^*_F > 1/2$). From hereon, steady-state variables will be denoted with a bar (i.e., $\overline{x}$) and percent deviations from the steady state will be labelled with a hat (i.e., $\hat{x}_t = (x_t - \overline{x})/\overline{x}$).

Two key log-linearized expressions will be used recurrently in the derivation of the equilibrium below. The first links the RER to the TOT (23.8):

$$\widehat{\mathcal{Q}}_t = (2a_H - 1)\widehat{\mathcal{T}}_t. \tag{23.21}$$

Absent home bias ($a_H = 1/2$), PPP holds and $\widehat{\mathcal{Q}}_t = 0$ independently of equilibrium fluctuations in the TOT.

The second expression characterizes how the wealth gap $\mathcal{W}_t$ (23.16) can be written as a function of the consumption gap and the real exchange rate:

$$\hat{\mathcal{W}}_t = \sigma(\hat{C}_t - \hat{C}^*_t) - \widehat{\mathcal{Q}}_t, \tag{23.22}$$

an expression we have seen must be zero under full risk sharing. Using the above and taking expectations of the first-order approximation of (23.20), we see that

$$\mathbb{E}_t[\hat{\mathcal{W}}_{t+1} - \hat{\mathcal{W}}_t] = \mathbb{E}_t[\hat{\epsilon}_{t+1}] = 0.$$

In a bond economy, shocks result in the accumulation (or reduction) of net foreign assets. Hence, relative wealth will change in response to shocks at each point in time. This condition also implies that, in the linearized version of the bond economy, relative wealth is non-stationary: the economy does not converge back to the symmetric steady state around which the equilibrium is approximated. The non-stationarity of the wealth distribution is due to the fact that the model is linearized around a non-stochastic steady state in which the real interest rate is equal to the inverse of the discount factor ($\beta R = 1$). The true, non-linear solution of the model does not, however, exhibit non-stationarity: in response to shocks, the economy is drifting slowly around an average relative wealth position of zero but never settles at zero since shocks are reoccurring and are large enough to influence all prices.[26] This is the classic 'unit-root' problem in small open economy models with linearization; adding debt-elastic discounting restores stationarity. In Chapter 11, we studied closed-economy heterogeneous-agent economies and one can imagine extending the present model to include a continuum of countries (and a continuum of goods), rather than just two (of each). A law of large numbers for the productivity shocks, which would then be idiosyncratic and hence 'small', would then imply the existence of a steady state where $\beta R < 1$.

Since the log-linearized setting is a very convenient tool, we will instead slightly change the model so as to yield convergence back to steady state. This is accomplished with the aid of 'Uzawa-style' preferences, where the discount factor is assumed to rise with the accumulation of net debt.[27]

23.3.4 Global equilibrium: a relative demand-relative supply framework

In the rest of this section, we study the properties of the model, discussing how it can be brought in line with the empirical responses shown in Figure 23.5. At an aggregate level, global output—which we treat as exogenously given for now—must be equal to global consumption, independently of the structure of financial markets:

$$\hat{Y}_{H,t} + \hat{Y}^*_{F,t} = \hat{C}_t + \hat{C}^*_t.$$

The equilibrium allocation of consumption across borders and international relative prices will nonetheless vary with the degree of risk sharing. For this reason, it is instructive to solve and study the model in relative terms, showing first how the demand for the goods produced in each economy vary as a function of the TOT (the relative price of output), and the wealth gap. Using this schedule in conjunction with relative supply, we then characterize the equilibrium allocation as a function of fundamental (supply) shocks, and analyze their international propagation.

Relative demand. We now introduce a notion of *relative demand* for home vs. foreign goods as a function of wealth gap and international relative prices. Define the total demands for each good by $D_{H,t} \equiv c_{H,t} + c^*_{H,t}$ and $D^*_{F,t} \equiv c_{F,t} + c^*_{F,t}$. Combining this definition with (23.6), showing how the ratio of the home to foreign demand for Home (Foreign) goods depend on the TOT, and (23.8), showing how the TOT are related to the RER, we obtain the following expression for the (log-linearized) relative total demand

$$\hat{D}_{H,t} - \hat{D}^*_{F,t} = (2a_H - 1)\left(\hat{C}_t - \hat{C}^*_t\right) + 4a_H(1 - a_H)\phi\widehat{\mathcal{T}}_t. \tag{23.23}$$

[26] The borrowing constraints will prevent relative wealth from exploding.

[27] Alternatively, one can assume infinitesimal costs of holding foreign assets, sufficient to motivate agents to run their external assets down in the long run. Key differences across solutions are analyzed by Schmitt-Grohé and Uribe (2003) and Bodenstein (2011).

Equation (23.23) decomposes demand into wealth-driven and substitution-driven components. Differences in demand for home goods and foreign goods come from two sources: (i) differences in the level of consumption, driving the demand for each goods depending on home bias, and (ii) international relative prices. We know, from (23.22), that the consumption difference comes from (i) difference in wealth and (ii) relative price level of each country. Using (23.22), we can obtain:

$$\hat{D}_{H,t} - \hat{D}^*_{F,t} = \sigma^{-1}(2a_H - 1)\hat{\mathcal{W}}_t + \sigma^{-1}\left[4a_H(1-a_H)\left(\sigma\phi - 1\right) + 1\right]\widehat{\mathcal{T}}_t, \tag{23.24}$$

This expression highlights the two key channels affecting relative demand for home and foreign goods. For given international relative prices $\widehat{\mathcal{T}}$, the wealth gap drives relative demand in proportion to the degree of home bias a_H—as captured by the coefficient $(2a_H - 1)$. For a given wealth gap $\hat{\mathcal{W}}_t$, a fall in the relative price of home goods (higher $\widehat{\mathcal{T}}$) unambiguously increases the relative demand for home goods.[28] This *substitution effect* of relative price movements is always negative, since relative demand of home goods falls with an increase in its relative price (the inverse of the terms of trade).

This highlights a crucial result: substitution is the only active channel either when markets are complete ($\hat{\mathcal{W}}_t = 0$), in which case the adjustment of prices supports the first-best allocation, or if there is no home bias ($a_H = 1/2$) irrespective of the degree of risk sharing. With no home bias, the RER is constant, so relative demand is independent of transfers of purchasing power across borders.[29] With home bias, under incomplete markets, wealth effects associated with non-traded risk can work against substitution effects (in contrast with the first best and thus inefficiently) as a worsening in the TOT (i.e., a drop in the price of a country goods supply) pushes down on (imperfectly insured) households' purchasing power.

Relative supply. The relative supply of the home and foreign goods is independent of relative prices—simply $\hat{Y}_{H,t} - \hat{Y}^*_{F,t}$ here as we, for now, work with an endowment economy.

Equilibrium. The equilibrium is obtained combining the two schedules. To gain insights on its properties, the case of financial autarky is particularly instructive, since here the relative wealth gap is tightly linked to current real incomes and thus relative prices. An expression for the wealth gap under financial autarky can be derived using equation (23.22) and log-linearizing (23.15), which invokes equilibrium in that it uses the autarky budget constraint, along with its foreign analog. Evaluated at a given supply of home and foreign goods, the wealth gap can be written as follows:

$$\hat{\mathcal{W}}^{FA}_t = \sigma\left[\hat{Y}^S_{H,t} - \hat{Y}^{*S}_{F,t} - 2(1-a_H)\widehat{\mathcal{T}}_t\right] - (2a_H - 1)\widehat{\mathcal{T}}_t. \tag{23.25}$$

Substituting this into (23.24), and equating demand and supply, the relative demand schedule can be written as:

$$\hat{D}_{H,t} - \hat{D}^*_{F,t} = \left[1 - 2a_H\left(1-\phi\right)\right]\widehat{\mathcal{T}}_t. \tag{23.26}$$

[28] Note that the coefficient on $\widehat{\mathcal{T}}_t$ in brackets is unambiguously positive, and increasing in the trade elasticity ϕ.

[29] See the controversy between Keynes and Ohlin on the effects of 'transfers' in the 1920s.

Under financial autarky, the coefficient on $\widehat{\mathcal{T}_t}$ will unambiguously be positive if PPP holds (i.e., if $a_H = \frac{1}{2}$), as is also the case under complete markets, such that relative home demand is unambiguously downward sloping in the international relative price of home output. But, under financial autarky, relative demand can be negatively or positively associated with relative price when there is home bias, depending on the elasticity ϕ. Here, relative demand is ultimately shaped by the endogenous response of the wealth gap to shocks. For example, the equilibrium relative demand for the home goods will be increasing in its relative price (i.e., $\frac{\partial(\hat{D}_{H,t}-\hat{D}^*_{F,t})}{\partial \mathcal{T}_t} < 0$) with an elasticity $\phi < 1 - \frac{1}{2a_H}$. Intuitively, with home bias in consumption and strong complementarity between home and foreign goods, home demand needs to rise enough to buy the higher supply of home goods. This is only possible if home goods prices rise in equilibrium, driving up home incomes.

Figure 23.6 illustrates the equilibrium by plotting relative demand (RD) and relative supply (RS) in the space $\{Y_H/Y_F,\, 1/\mathcal{T}\}$—note that y-axis reports the inverse of the TOT. RD denotes the relative demand schedule for $\phi > 1-\frac{1}{2a_H}$, upward sloping in this space, while RD' represents relative demand with sufficient complementarity which slopes downward. Relative supply in the endowment economy is a vertical line. The figure illustrates an example in which a country-specific productivity innovation brings Home output above Foreign output, shifting the relative supply schedule from RS to RS'.

Figure 23.6: Relative demand and relative supply

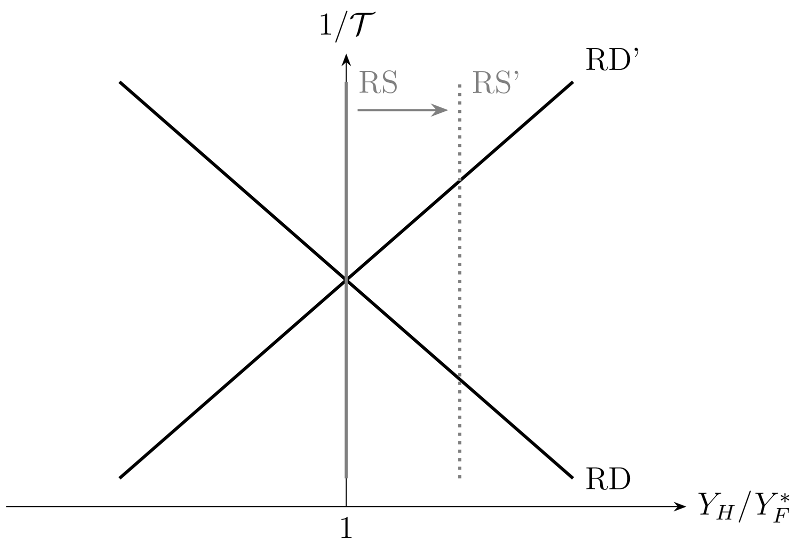

The upward schedule in the figure (RD') illustrates the case in which, for a small enough ϕ, the wealth effects from TOT movements (captured by equation (23.25)) are the primary drivers of relative demand, dominating substitution effects. In response to a productivity innovation at home, relative home consumption and the demand for home goods both grow higher relative to their foreign counterpart, the home TOT strengthens, and the home RER appreciates,[30] resulting in a negative correlation of the exchange rate with both relative consumption and output. This suggests that models with incomplete markets (here in an

[30]Using equilibrium conditions, and further considering the limit of full home bias, $a_H \to 1$, we can show that $\mathcal{Q}_t = \frac{\hat{Y}_{H,t}-\hat{Y}^*_{F,t}}{2\phi-1}$.

extreme form), specified to reflect strong wealth effects of supply shocks, may provide an explanation of the Backus-Smith puzzle.

The same applies in the bond economy. To see this, define real assets as $b_t \equiv B_t/P_t$, and the *ex post* (gross) real interest rate as $r_t \equiv (1/Q_t)(P_{t-1}/P_t)$, which in steady state is equal to β^{-1}. Then write the bond-economy wealth gap, up to first order around a symmetric steady state with zero net foreign assets (evaluated again at $\hat{Y}_{H,t} = \hat{D}_{H,t}, \hat{Y}^*_{F,t} = \hat{D}^*_{F,t}$), by substituting relative consumption in (23.23) with the relative budget constraint and imposing equilibrium:

$$\hat{\mathcal{W}}_t^{BE} = \sigma \left[\hat{D}_{H,t} - \hat{D}^*_{F,t} - 2(1-a_H)\widehat{\mathcal{T}}_t\right] - (2a_H - 1)\widehat{\mathcal{T}}_t + 2\sigma\beta^{-1}[\tilde{b}_{t-1} - \beta\tilde{b}_t]. \tag{23.27}$$

where $\tilde{b}_t = b_t - \bar{b}$ and $\bar{b} = 0$. This expression differs from the expression for $\hat{\mathcal{W}}_t^{FA}$ (23.25) (evaluated again at $\hat{Y}_{H,t} = \hat{D}_{H,t}, \hat{Y}^*_{F,t} = \hat{D}^*_{F,t}$) only by the last term in equation (23.27), which generally depends on both *current* and *future* relative prices. As agents smooth consumption in response to shocks, borrowing and lending may attenuate the wealth effects of inefficient relative price adjustment on current relative demand, compared to the financial autarky case. However, in the next section we show that intertemporal trade also magnifies the wealth effects of persistent or anticipated future productivity innovations on current demand. Remarkably, with these types of shocks, the condition on the trade elasticity for strong wealth and demand effects to materialize are different.

23.3.5 International transmission of productivity shocks via relative prices, wealth and demand

We now elaborate on the international transmission of supply shocks, focusing on their effects on relative prices and financial and capital flows. This will facilitate the comparison of the model with the empirical responses shown in Figure 23.5.

Terms of trade and real exchange rate. With complete markets, combining the relative demand equation, (23.24), with relative supply, and solving for the TOT yields:

$$\widehat{\mathcal{T}}_t = \frac{\sigma}{4a_H(1-a_H)(\sigma\phi - 1) + 1}(\hat{Y}_{H,t} - \hat{Y}^*_{H,t}). \tag{23.28}$$

In this expression, the denominator is unambiguously positive when $a_H \in (0.5, 1)$. So, in equilibrium, the home TOT unambiguously worsen and the RER depreciates with an increase in relative domestic productivity. Note that, other things equal, the coefficient on the productivity differential vanishes when goods are highly substitutable and is largest as ϕ approaches zero.

These results do not extend to cases with incomplete markets. Under financial autarky, using (23.25), the equilibrium TOT become:

$$\widehat{\mathcal{T}}_t = \frac{1}{1 - 2a_H(1-\phi)}(\hat{Y}_{H,t} - \hat{Y}^*_{H,t}). \tag{23.29}$$

In this expression, the denominator is not necessarily positive. The home TOT worsen following a domestic productivity improvement if and only if

$$\phi > \tilde{\phi}^{FA}_{TOT} \equiv 1 - \frac{1}{2a_H}.$$

For values of the elasticity below this threshold, a home productivity gain results in *stronger* TOT, as in the RD' schedule in Figure 23.6. Thus, the sign of the TOT response depends on ϕ under autarky, but not under complete markets.

Comparing the expressions for the equilibrium TOT under complete markets, (23.28), and financial autarky, (23.29), yields a remarkable result. Observe that the equilibrium association between the TOT and relative productivity is identical across complete markets and financial autarky for $\sigma = \phi = 1$. This resonates with the contribution by Cole and Obstfeld (1991), where TOT movements move one-to-one with relative output, ensuring efficient sharing of productivity risk independently of trade in assets.[31] In this case, a relative rise in home output unambigously worsens the TOT under any financial market structure.

Cross-border real and financial flows. Key insights on the response of capital flows can be derived through further analysis of the complete-markets case. There are many equivalent ways to decentralize efficient state-contingent payments that support full risk sharing under complete markets. For our purposes, the most instructive approach consists of modeling them in analogy to movements in net foreign assets in the bond economy around the symmetric steady state with zero net foreign wealth. Following this approach, we introduce a 'notional real foreign asset,' denoted by $\widehat{\mathcal{B}}_t = \mathcal{B}_t - \overline{\mathcal{B}}$ with $\overline{\mathcal{B}} = 0$, defined as the cumulative real net exports scaled by steady-state output. Recall that real net exports are always uniquely defined in the complete-market allocation. Using the resource constraints together with the risk-sharing condition and equilibrium output, we then derive how net exports respond to productivity shocks, driving our notional flows and their accumulation over time, consistent with complete markets:

$$\widehat{\mathcal{B}}_t - \beta^{-1}\widehat{\mathcal{B}}_{t-1} = (1 - a_{\mathrm{H}})\,\sigma^{-1}\left[(2a_{\mathrm{H}}\,(\sigma\phi - 1) + 1 - \sigma)\,\widehat{\mathcal{T}}_t\right]. \tag{23.30}$$

The key result here is that net exports and their counterpart, financial flows, are proportional to the TOT (and hence the RER), but the relation can have either sign. To appreciate why, suppose that, at the initial steady state (i.e., $\widehat{\mathcal{B}}_{t-1} = 0$), the home TOT unexpectedly worsen because of a domestic output boom (i.e., $\widehat{\mathcal{T}}_t > 0$). Net exports turn positive and capital flows from the more productive country to the less productive country (hence from home into foreign) provided

$$\phi > \frac{1}{\sigma} + \frac{1}{2a_H}\left(1 - \frac{1}{\sigma}\right). \tag{23.31}$$

[31] With $\phi = \sigma = 1$ and symmetric preferences in consumption, a country experiencing an increase in its (relative) output, will also experience a proportional fall in the international price of its goods. Hence, the value of national output remains constant, but consumption increases in both countries—residents abroad have higher income in real terms as they can buy the goods produced by the more productive country more cheaply.

Note that this condition on the trade elasticities can be interpreted as the general-equilibrium counterpart of the classic Marshall-Lerner condition (on exports and imports elasticities) for an improvement in the trade balance in response to a worsening of the TOT.[32] Notably, with log utility (i.e., $\sigma = 1$), the condition will boil down to gross substitutability between home and foreign goods (i.e., $\phi > 1$).

If the above condition is violated, financial resources will flow *into* the more productive country, which will thus run a trade deficit. It is worth spelling out the economics of this result, focusing again on the case $\sigma = 1$. In an efficient allocation, when home experiences an output boom, home households will increase their consumption of domestically-produced output. If $\phi < 1$ (home and foreign goods are gross complements), the home marginal utility from consuming home imports rises. It is then efficient for the foreign country to produce and export more. In a decentralized equilibrium of our economy with complete markets, this allocation is supported by a sharp deterioration in the home TOT, increasing the value of foreign output and triggering a flow of financial payments from abroad to home. In equilibrium, the home country runs an external deficit even though its output is high relative to foreign output.

In a bond economy, the general-equilibrium analog of the Marshall-Lerner conditions is more complex. For an elasticity sufficiently low, the model can generate a counter-cyclical trade deficit in response to productivity gains also when these gains lead to a TOT appreciation (hence irrespective of whether (23.31) holds).

23.4 The production economy

In the previous sections, we studied the transmission of productivity shocks in an endowment economy. Here, we extend our analysis to an environment in which labor supply is endogenous, so that we can discuss the ability of the model to account for the correlation of national business cycles seen in the data.

23.4.1 Model setup

The model is identical to the one in the previous section, with the following differences: (i) labor supply is endogenous and (ii) output is produced by competitive firms. In particular, the expected lifetime utility of the representative household is now:

$$U_t = \mathbb{E}_t \left[\sum_{j=0}^{\infty} \beta^j \left[u(C_{t+j}) - \zeta \nu(L_{t+j}) \right] \right], \tag{23.32}$$

where $\nu(L) = L^{1+\eta}/(1+\eta)$ is the instantaneous disutility from household labor supply (with $\eta > 0$ representing the inverse Frisch elasticity of labor supply), and $\zeta > 0$ is a constant parameter. With this specification, and regardless of the international financial market structure, the representative home household's labor supply in a competitive labor market

[32]For a classic discussion of Marshall-Lerner conditions in modern general equilibrium models see Backus, Kehoe, and Kydland (1994a).

will be determined by:

$$C_t^{-\sigma} w_t = \zeta L_t^{\eta}, \tag{23.33}$$

where w_t denotes the real wage. This expression equalizes the marginal benefit, in consumption terms, of an extra hour of work with its marginal disutility cost to the household. The preferences of the representative foreign consumer are defined analogously, with consumption C_t^*, labor L_t^*, expected lifetime utility U_t^*, as well as $\beta^* = \beta$, $\sigma^* = \sigma$, $\zeta^* = \zeta$ and $\eta^* = \eta$.

On the supply side, firms in each country produce specialized goods and are perfectly competitive. The representative home firm has the following production function:

$$Y_{H,t} = Z_t L_t^{\alpha}, \tag{23.34}$$

where Z_t represents an exogenous productivity process and $\alpha \in (0, 1]$ is the labor share. The representative foreign firm has an analogous production function, with foreign productivity Z_t^*. The stochastic processes for home and foreign productivity satisfies:

$$Z_t = \rho_1 Z_{t-1} + (1 - \rho_1)\overline{Z} + \varepsilon_t \quad \text{and} \quad Z_t^* = \rho_1 Z_{t-1}^* + (1 - \rho_1)\overline{Z} + \varepsilon_t^*, \tag{23.35}$$

where $\rho_1 \in (0, 1)$ is shock persistence and $\overline{Z} = 1$ denotes the steady-state value of productivity. In computational exercises, we impose that $\varepsilon_t, \varepsilon_t^* \sim \mathcal{N}(0, \sigma^2)$ where $\varepsilon_t \perp \varepsilon_t^*$ for all t.

The problem of the competitive firm is standard: hire labor to maximize profits, $p_{H,t}Y_{H,t} - W_t L_t$, with nominal wage W_t, such that real wages w_t equalize the marginal product of labor. Firms make zero profits in equilibrium (i.e., $p_{H,t}Y_{H,t} = W_t L_t$), so labor demand is determined by:

$$w_t = \frac{p_{H,t}}{P_t} Z_t L_t^{\alpha - 1}, \tag{23.36}$$

where $w_t \equiv W_t / P_t$. This expression highlights an important feature of the open-economy setting. The wage relevant for labor demand is not the real consumption wage w_t, but the nominal wage deflated by the product price—i.e., the product real wage $W_t / p_{H,t}$. Therefore, the relative price of domestic production in terms of consumption goods $p_{H,t}/P_t$ can act as a shifter of labor demand and thus introduces a channel of spillovers from shocks abroad.

23.4.2 Relative supply, relative demand, and global output

The relative demand is identical to the one in the endowment economy, given by equation (23.24). The supply of the home and foreign goods as a function of relative wealth and international relative prices can be derived by combining the optimizing conditions for labor demand (23.33) and labor supply (23.36), with the production function (23.34), for the two economies, using the wealth-gap definition (23.22) to substitute out relative consumption. Up to a first order, the home *relative to* foreign supply is:

$$\hat{Y}_{H,t}^{S} - \hat{Y}_{F,t}^{*S} = \frac{1 + \eta}{1 - \alpha + \eta} \left(\hat{Z}_t - \hat{Z}_t^* \right) - \frac{\alpha}{1 - \alpha + \eta} \left(\hat{\mathcal{W}}_t + \widehat{\mathcal{T}}_t \right), \tag{23.37}$$

where the superscript 'S' on output denotes supply and (recalling our definitions) η is the inverse Frisch elasticity, α captures the returns to scale in production, and Z_t^* is a home (foreign) productivity shock. The TOT now enter the relative supply expression (23.37) because the equilibrium demand for labor (23.36) depends on the relative price of the marginal product of labor in terms of domestic consumption, and hence varies with the relative price of imports.

Of course, the wealth gap also depends on relative prices, but in the case of complete market it is identically equal to zero, $\hat{\mathcal{W}}_t^{CM} = 0$. Since the expression multiplying $\widehat{\mathcal{T}}_t$ is always negative, (23.37) establishes that under complete markets the relative supply of home in terms of foreign goods is invariably increasing in the relative price of the home goods as a function of parameters related to technology and labor supply. Under full risk sharing, a real appreciation is associated with a fall in relative consumption: via equation (23.33), a real appreciation is also associated with a rise in the relative supply of labor. Remarkably, under complete markets the slope of the relative supply curve does not depend on openness.

Under financial autarky, as already seen, the relative wealth gap is tightly linked to current real incomes and thus relative prices (see equation (23.25)). Substituting this into (23.37), the relative supply schedule becomes:

$$\hat{Y}_{H,t}^S - \hat{Y}_{F,t}^{*S} = \frac{1+\eta}{1-\alpha+\eta+\alpha\sigma}\left(\hat{Z}_t - \hat{Z}_t^*\right) + \frac{2\alpha(1-a_H)(\sigma-1)}{1-\alpha+\eta+\alpha\sigma}\widehat{\mathcal{T}}_t. \tag{23.38}$$

With no cross-border trade in assets, the relative supply of home in terms of foreign goods is either *increasing* or *decreasing* in the relative price of the home goods, depending on $\sigma \lesseqgtr 1$. When $\sigma > 1$, for any given amount of output supplied, the purchasing power of domestic residents rises with higher output prices, driving up leisure. That is, the wealth effects from a RER appreciation result in a fall in labor supply—the opposite of the complete-markets case. Remarkably, relative supply is independent of the TOT if preferences are logarithmic ($\sigma = 1$).[33]

To calculate global output, we combine the resource constraint with the sum of the national output supply, to express global output as a function of productivity shocks. Up to a first order, and dropping superscripts 'D' and 'S', we have the following log-linearized resource constraint:

$$\hat{Y}_{H,t} + \hat{Y}_{F,t}^* = \hat{C}_t + \hat{C}_t^* = \frac{1+\eta}{1-\alpha+\eta+\alpha\sigma}\left(\hat{Z}_t + \hat{Z}_t^*\right).$$

This shows that, globally, productivity gains unambiguously raise global output in equilibrium, regardless of the structure of international financial markets.

23.5 Substitution and wealth effects in the international transmission mechanism

Using the production economy presented above, we close our study with an educated selection of quantitative exercises meant to illustrate the workings of substitution and wealth effects in the international transmission of shocks. Figure 23.7 plots the effects of a positive home

[33] See Heathcote and Perri (2002) for further analysis of the international real business cycle under financial autarky.

productivity shock under complete markets, financial autarky, and with riskless bonds only, for three alternative parameterizations. The first two columns bring together the results we have presented so far, demonstrating the effects of a transitory increase in home productivity for two values of the trade elasticity—where calibrations draw on the long-standing debates around the empirical estimation of this parameter.[34] The last column introduces a novel set of results, highlighting the specific, dynamic nature of the bond economy. Here, we assume that current shocks create expectations of further increases in output in the future, impacting wealth and demand as national agents reassess the present discounted value of their current and future income and smooth consumption by borrowing and lending internationally.

Strong substitution effects. In the first two columns, we plot impulse responses to a home productivity shock, which we model using the AR(1) process in (23.35) with persistence $\rho_1 = 0.97$. In column (a) we set $\phi = 1.5$. As shown analytically above, in this case, substitution effects dominate regardless of the degree of cross-border risk sharing, such that the home TOT deteriorates in response to the shock and cross-border spillovers are negative.

Under complete markets, the correlation between home and foreign business cycles is counter-factually (and strongly) negative. Substitution effects from the home RER depreciation boost home production relative foreign—home consumption is less volatile than output. Output co-movements remain negative also under financial autarky, although moderate in comparison: home output rises by less—and foreign output falls by much less—than under full risk sharing. The depreciation of the home TOT now erodes the purchasing power of home consumers, while boosting that of foreign consumers.[35] Remarkably, impulse responses for the bond economy lie between the other two cases. The bond economy is quantitatively different from the complete-markets benchmark, as intertemporal consumption smoothing falls short of replicating full risk sharing. Yet capital flows in the same direction. In both cases (holding (23.31) and its counterpart for the bond economy), home real net exports are positive, such that resources flow from the home economy to foreign.

Strong wealth effects with a low trade elasticity. In column (b), we plot impulse responses to the same AR(1) home productivity shock, but lower the trade elasticity to $\phi = 0.3$. In this case, business cycles are positively correlated and under, imperfect risk sharing, RER tend to appreciate with the home relative output boom—replicating qualitatively the empirical responses shown in Figure 23.5.

Under complete markets, the TOT deteriorates by much more than in the previous case, such that home residents are poorer. Remarkably, the notional financial flows run from foreign to home, in spite of the fact that home experiences a rise in productivity. With incomplete markets, wealth effects dominate, driving the response of the equilibrium

[34]Estimates of the elasticity at the macro level are typically low, while micro estimates high (see, e.g., Feenstra, Luck, Obstfeld, and Russ, 2018, and the references within). There are also differences across horizons, with elasticities typically higher in the long run than in the short run (Boehm, Levchenko, and Pandalai-Nayar, 2023). Example of models with short vs. long-run elasticities include those with distributive capital (Crucini and Davis, 2016), or its analog in terms of 'customer list' capital (Drozd and Nosal, 2012), as well as models with nested CES (Cooley and Quadrini, 2003).

[35]Recall from equation (23.21) that there is a linear relationship between TOT and RER in our approximation.

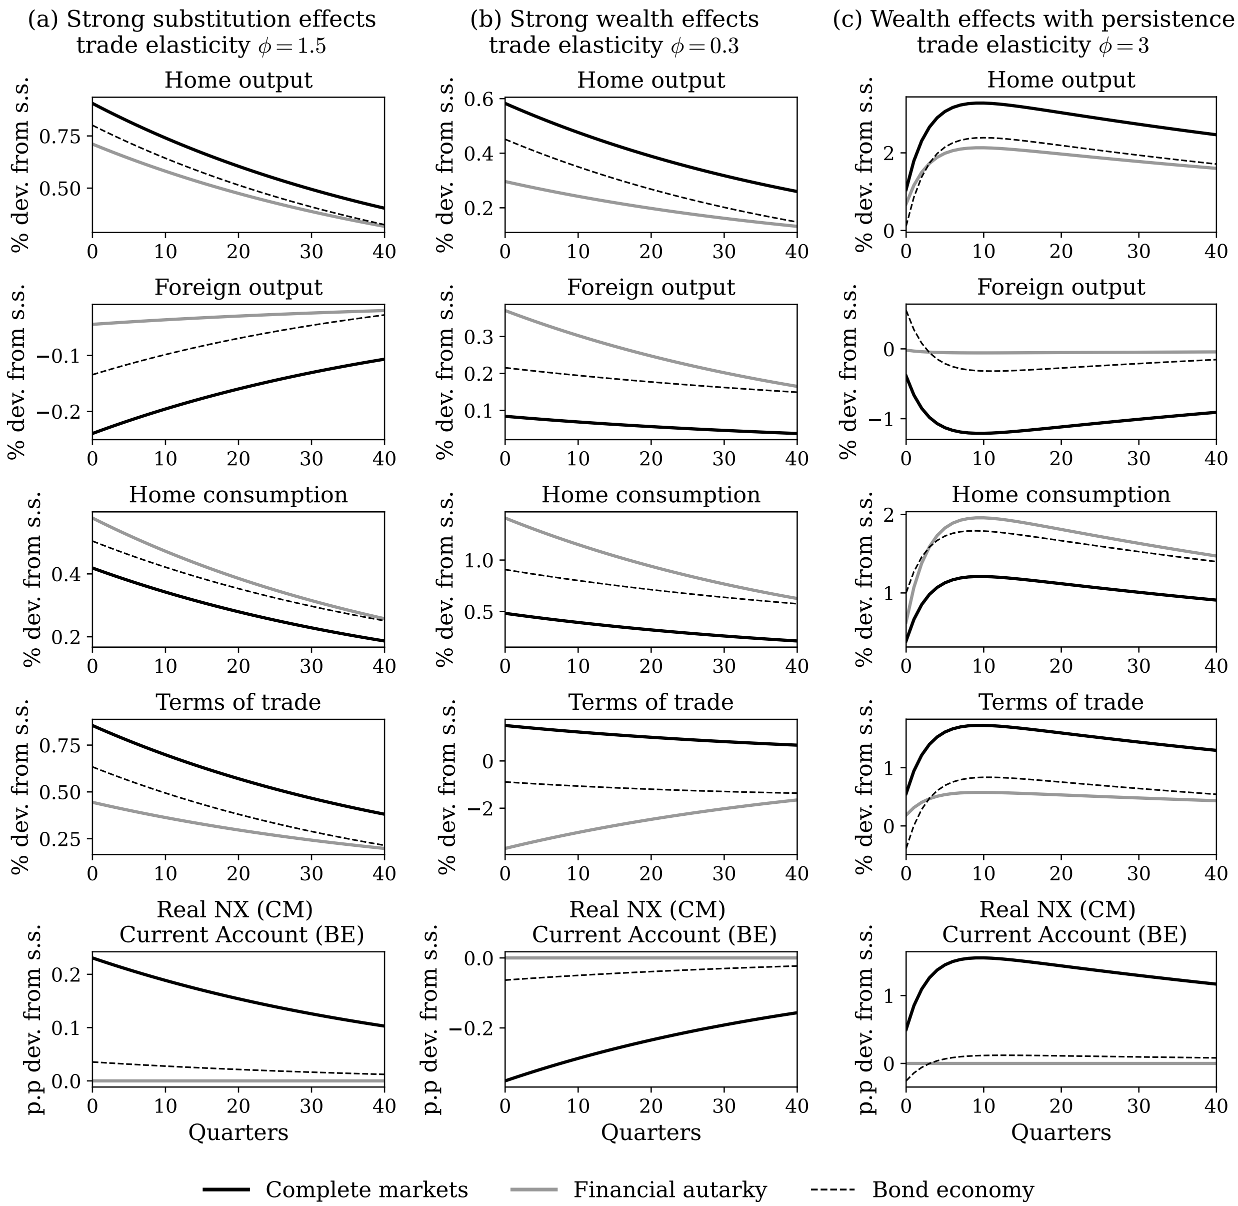

Figure 23.7: International transmission of positive home productivity shock

Note: impulse response to positive and persistent 1% increase in home productivity. In all columns: $\beta = 0.99$ (quarterly), $\sigma = 2$, $\eta = 1$, and ζ such that steady-state labor supply is 1/3, $a_H = 0.7$, $\alpha = 1$ $\overline{Z} = 1$. Bond economy made stationary using Uzawa-style preferences. *Persistence*: In columns (a) and (b) home productivity is an AR(1) process with persistence $\rho_1 = 0.97$; in column (c) home productivity follows an AR(2) process with persistence $\rho_1 = 1.7$ and $\rho_2 = -0.703$. *Trade elasticity*: in column (a) $\phi = 1.5$; (b) $\phi = 0.3$; and (c) $\phi = 3$. Parameterization is symmetric across countries.

TOT in the opposite direction. The relative price of home output improves strongly in financial autarky and the bond economy. With incomplete risk sharing, foreign output is higher relative to the complete-market case—the home appreciation boosts demand by home residents, who increase consumption of both domestic and the foreign goods. However, consumption is now counterfactually more volatile the output, contradicting one of the key stylized facts documented earlier. Note that the home country runs a persistent trade deficit both under complete markets and in the bond economy, even though the TOT move in opposite directions.

Strong wealth effects with high shock persistence and high trade elasticities. In column (c), we bring forward a different way to model cross-border wealth effects of shocks. First, we assume a hump-shaped AR(2) productivity process (with persistence parameters ρ_1 and ρ_2) so that, in effect, an increase in home productivity today also conveys news about even *higher* future productivity. It is worth noting that this is a convenient reduced-form way to capture output dynamics in the model without capital accumulation.[36] Second, we assume that home and foreign goods are highly substitutable. Under this parameterization, the expected real income from higher home output in the future is positive, i.e., higher quantities are not associated with a more than proportional fall in output prices. This means that, under incomplete markets, the shocks have a significant impact on home wealth.

The exercise highlights macroeconomic dynamics specific to the bond economy. Under complete markets and financial autarky, substitution effects dominate in equilibrium: higher home productivity is associated with a RER depreciation, but business cycles are counterfactually negatively correlated.[37]

The dynamics of the bond economy overturn these negative cross-border spillovers and ensure that consumption is less volatile than output—qualitatively mirroring the empirical responses shown in Figure 23.5. Faced with the prospect of even higher productivity in the future, home households have an incentive to run a current-account deficit and borrow to raise current consumption. Consumption smoothing thus dampens the home TOT deterioration relative to complete markets and financial autarky—it even improves them on impact, replicating the appreciation seen in Figure 23.5. Facing these changes in relative prices, foreign households produce more in the first periods after the shock. As a result, cross-border transmission is positive. Over time, the home terms of trade eventually deteriorates, and the cross-border transmission turns negative as home runs trade surpluses to repay its debt.[38]

[36] A hump-shaped output profile mimics the dynamics obtained from a model with capital accumulation, as investment responds steadily to persistent productivity gains (see, e.g., Corsetti et al., 2008b). See early work by Baxter and Crucini (1995), assessing the implications of increasing the persistence of productivity shocks.

[37] Observe that in our model without capital, anticipations of future output do not affect the current allocation under both financial autarky and complete markets. Under financial autarky, agents have no instruments to smooth consumption intertemporally. Under complete markets, future productivity variations are completely insured by financial contracts at the time they materialize. Instead, anticipation of future output dynamics becomes crucial in the bond economy.

[38] Strong wealth effects can also be obtained in model assuming that the business cycle is driven by stochastic shifts in the growth rate (Aguiar and Gopinath, 2007), 'news' shocks (Jaimovich and Rebelo, 2008) or input adjustment costs and cointegrated productivity shocks across countries (Rabanal and Rubio-Ramírez, 2015).

Taking stock. Taken together, these results suggest that, to bring the benchmark model closer to the empirical responses, an increase in output has to be associated with a 'boom' in demand. Conditional on productivity shocks, the model features that are crucial for generating these facts are incomplete financial markets and either a low trade elasticity (Figure 23.7, column (b)) or shock processes that raise the growth rate of output persistently (Figure 23.7, column (c))—see Corsetti et al. (2008b) for a generalization of this analysis to models with capital accumulation.

23.6 Richer frameworks

The framework developed and studied in this chapter provides a core building block for modern open-economy macroeconomics. Armington aggregators of domestic and foreign goods (and/or intermediate) inputs are a key feature of seminal contributions to the international real business cycle literature (e.g., Backus et al., 1992; Backus, Kehoe, and Kydland, 1994b). Building on this same structure, the New Keynesian open-economy model—–as well as earlier contributions to the New Open Economy Macroeconomics (NOEM) literature (e.g., Obstfeld and Rogoff, 1995; Corsetti and Pesenti, 2001; Benigno and Benigno, 2003; Devereux and Engel, 2003)—departs from the assumption of perfect competition, instead assuming that firms operate under monopolistic competition with nominal rigidities.[39]

Richer specifications make the model better suitable to confront the data. Investment and capital accumulation enriches the dynamic of the model, helping generating countercyclical current account balances: a positive productivity shock to domestic technology simultaneously raises output and the demand for imported capital goods—see the literature building on Backus et al. (1994b), and especially the discussion in Raffo (2008). Modeling a non-tradable good sector in each economy introduces a new relative price (nontradable to tradables) that breaks the proportionality between real exchange rate and terms of trade, potentially magnifying currency and misalignment volatility and the wealth effects from shocks (see Stockman and Tesar, 1995; Corsetti et al., 2008a, 2014). In addition to fostering data-rich micro-to-macro research, modeling firm dynamics, distributive trade and global supply chains contribute to account for possibly state-contingent and time-varying trade elasticities, driving destination-specific price and markup adjustment by exporters and possibly contributing to explain the cross-border correlation of economic activity (see, e.g., Ghironi and Melitz, 2005; Atkeson and Burstein, 2008; Johnson, 2014; Bergin and Corsetti, 2020).[40] Complementing work on destination-specific price adjustment, introducing nominal frictions can help to explain why import prices remain stable in local currency despite a high volatility of nominal exchange rates, potentially paving the way to studies addressing the Mussa Puzzle. More generally, monetary models featuring price or wage stickiness make it possible to study the international transmission of nominal demand shocks, and explore

[39] Numerous recent studies have focused on nominal rigidities in a world with dominant currency pricing (e.g., Gopinath, Boz, Casas, Díez, Gourinchas, and Plagborg-Møller, 2020; Egorov and Mukhin, 2023; McLeay and Tenreyro, 2025).

[40] Modeling trade costs along manufacturing or distribution chains introduces a wedge between the (market) trade elasticity of imports, relevant for firms' decisions, and the elasticity in production (Corsetti and Dedola, 2005). Corsetti, D'Aguanno, Dogan, Lloyd, and Sajedi (2023) embeds global value chains into the model outlined in this chapter.

normative implications for the optimal design of stabilization policy in global settings (see Corsetti, Dedola, and Leduc, 2010, 2023, for further exposition).

From a global perspective, studies into financial spillovers highlight important asymmetries in the global transmission associated with the dominant position of U.S. and the U.S. dollar in the international financial system (Gourinchas and Rey, 2007; Miranda-Agrippino and Rey, 2020). Explicitly modeling cross-border financial intermediation subject to frictions can capture these asymmetries, as shown by Gabaix and Maggiori (2015). Building on the approach by these authors, recent literature has shown that a rich specification of the model can account for international business cycle facts jointly with international finance facts. The latter includes the apparent 'disconnect' of exchange rates from fundamentals, the close correlation of the real and nominal exchange rate (the Mussa puzzle), as well as the Uncovered and Covered Interest Parity Puzzles (Itskhoki and Mukhin, 2021a,b). Financial market imperfections implying imperfect risk sharing are essential for these results. Literature has specifically stressed two imperfections. The first is moral hazard, giving rise to financial constraints that limit the risk-bearing capacity of intermediaries and makes the recourse to intermediation to borrow and lend internationally costly for the households.[41] The second is noise trading, that injects exogenous volatility in the currency and asset markets (see, e.g., Devereux and Engel, 2002; Jeanne and Rose, 2002).[42]

Relative to the rich and growing body of open-macro literature, this chapter has shown that the international transmission mechanism is best understood modeling frictions in the financial and good markets that undermine efficient risk sharing and create an economically meaningful feedback between equilibrium international prices and current and anticipated future domestic output on the one hand, and cross-border wealth and demand on the other. While following this approach may address some of the "puzzles" in the literature, others remain outstanding. The field is now moving further towards integrating trade and macro, allowing for heterogeneity across households as well as firms/sectors, and rethinking global equilibria in the context of rising geopolitical fragmentation and strategic policy games undermining cross-border policy cooperation.

[41] Because of this cost, dynamically, the model shares the same properties of the bond economy with costly cross-border bond holdings as discussed by Schmitt-Grohé and Uribe (2003). Yet, explicitly modeling balance sheets creates opportunities to explore the role of international reserve policy, swap lines and other forms of international liquidity interventions (Bahaj and Reis, 2022).

[42] Studies into the structure of international financial market, where information, incentives and/or exogenous shocks may drive portfolio reshuffling unrelated to macro fundamentals, include Bacchetta and Van Wincoop (2006) and Bippus, Lloyd, and Ostry (2023).