

COPYRIGHT NOTICE:

Piet Sercu: International Finance: Theory into Practice

is published by Princeton University Press and copyrighted, © 2009, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

Follow links for Class Use and other Permissions. For more information send email to: permissions@press.princeton.edu

3

Spot Markets for Foreign Currency

In this chapter, we study the mechanics of the spot exchange market. Section 3.1 explains the various ways in which exchange rates can be quoted, and section 3.2 how the exchange markets themselves operate. Section 3.3 then considers exchange transactions in greater detail, focusing on bid and ask rates (that is, the rates at which a bank buys and sells). This also gives us an opportunity to learn about arbitrage. Specifically, in the third section, we shall already apply arbitrage arguments to the simplest possible problem, the relation between rates quoted by different banks for the same currency. Understanding this simple application now will make it easier to digest more complicated versions of similar arguments later. One such application already occurs in section 3.4, where we use arbitrage arguments to explain how exchange rates quoted by, for example, German banks (against EUR) relate to rates offered by New Zealand banks (against the NZD).

The chapter ends with the concepts of, and empirical evidence on, “purchasing power parity” (PPP) rates and real exchange rates. The conclusion of that part will be that exchange rates can make or break an exporting company, not just because of capital losses on foreign-currency-denominated receivables but possibly also because of a loss of competitiveness. Exchange risk even interferes with capital market equilibrium and the CAPM. These findings motivate the attention given to exchange rates in this book.

3.1 Exchange Rates

As we begin exploring exchange rates, we first provide a definition. We then describe the convention used to quote exchange rates throughout this book, as well as the conventions used in the exchange market. Finally, we explain how exchange rates are quoted in the presence of bid-ask spreads.

3.1.1 Definition of Exchange Rates

An exchange rate is the amount of a currency that one needs in order to buy one unit of another currency, or it is the amount of a currency that one receives

when selling one unit of another currency. An example of an exchange rate quote is 0.8 USD per CAD (which we will usually denote as “USD/CAD 0.8”): you can, for instance, buy a CAD by paying USD 0.80.

In the above, we have combined currency names following the conventions in physics: EUR/USD means euros per dollar just as “km/h” means kilometers per hour. This is the most logical convention. For instance, if you exchange 3m dollars into euros at a rate of 0.8 euros per dollar, the result is 2.4m euros—a number of euros. This fits with our notation:

$$\text{USD } 3\text{m} \times \text{EUR/USD } 0.8 = \text{EUR } 2.4\text{m}. \quad (3.1)$$

This may seem self-evident. The reason why we bring this up is that the pros do it differently. In the convention typically adopted by traders, bankers, and journalists, EUR/USD is not the dimension of the quote but the name of the exchange rate: it is the *value* of the euro, *expressed* in dollars, not its dimension. That is, the pros write “EUR/USD = 1.2345,” whereas we write “ $S_t = \text{USD/EUR } 1.2345$.” The dimension the trader asks for is USD/EUR, the inverse of what they write—but they do not mean a dimension, they mean a name.¹ In all our examples we use dimensions. The “name” notation pops up occasionally in press clippings or in pictures of trading screens, etc., and should not be a problem. To harden yourself, stare at the following entries for a full minute:

Currency name	Value
EUR/USD	USD/EUR 0.75
EUR/GBP	GBP/EUR 0.60
USD/CHF	CHF/USD 1.05

The telltale difference is that the dimension is immediately followed (or, occasionally, preceded) by the number. If there is no number, or if there is an “=” or “is” or “equals,” etc., between the ratio and the number, it must be the name of a rate. Sometimes practitioners drop the slash in the name and write EURUSD or EUR:USD instead of EUR/USD, which makes more sense.

It is even more crucial that you understand how exchange rates are quoted. While the notation is occasionally confusing—are we using dimensions or names?—there could be even more confusion as to which currency should be used as the numéraire. While you are familiar with the idea of buying goods and services, you may be less used to buying money with money. With exchange transactions, you need to agree which money is being bought or sold. There would be no ambiguity if one of the currencies were your home currency. A purchase then means that you obtained foreign currency and paid in home currency, the way you would do it with your other purchases too; and a sale

¹ It is sometimes whispered that the trader notation comes from a kind of pseudo-math like “EUR 1 = USD 1.2345,” where one then “divides both sides by USD.” The mind boggles. This is like denoting a speed as “1 h = 100 km” instead of $v = 100 \text{ km/h}$.

means that you delivered foreign currency and received home currency. If neither currency is your home currency, then you need to establish which of the two acts as the home currency.

Example 3.1. In a Paris bank, a tourist hands over USD 1,000 to the bank clerk and receives CAD 1,250 in return. This event would be described differently depending on whether the person is a U.S. tourist, a Canadian, or a Frenchman:

- The U.S. tourist would view this as a purchase of CAD 1,250 at a total cost of USD 1,000, implying a unit price of $[\text{USD } 1,000]/[\text{CAD } 1,250] = \text{USD}/\text{CAD } 0.8$.
- The Canadian would think of this transaction as a sale of USD 1,000 for CAD 1,250, implying a unit price of $[\text{CAD } 1,250]/[\text{USD } 1,000] = \text{CAD}/\text{USD } 1.25$.
- The Frenchman would regard this as an exchange of two foreign currencies, and would be at a loss if he were asked which of these is being sold and which bought.

Among pros, the currency in which the price is expressed is called the *quoting currency*, and the currency whose price is being quoted is called the *base currency* or *reference currency*. We avoid the terms, except in the next two lines. We have just noted that pros denote a rate as base/quoting (or, better, base:quoting) while its dimensions are quoting/base. A different issue is whether the quoting currency is the home or the foreign one.

3.1.2 Our Convention: Home Currency per Unit of Foreign Currency

Once we agree which country is, or acts as, the home country, we can agree to quote exchange rates as the price in units of home currency (HC) per unit of foreign currency (FC). That is, we quote the rate as HC/FC throughout this text, meaning that one unit of FC is worth N HC units (dimension HC/FC). As we shall see, some people do it differently and state that, with one unit of home currency, they can buy $M = 1/N$ units of foreign currency (FC/HC). We adopt the HC/FC convention because it is the most natural one. It is the convention we use when buying goods. For example, we say “the price is five dollars per umbrella” (HC/umbrella), not “with one dollar you can buy one-fifth of an umbrella” (umbrellas per unit of home currency).

Example 3.2.

1. A quote like USD/EUR 1.25 is an American’s natural quote for the EUR; it is the USD price an American gets or pays per EUR. For Germans or other Eurolanders, a quote as EUR/USD (euros per dollar) is the more natural one.
2. A quote like USD/CAD 0.75 is an American’s natural quote for the CAD, since the CAD is the currency in the denominator: a price in USD per CAD.

Expressing prices in HC is the convention for not just umbrellas but also for financial assets. Thus, standard finance results hold: the current market value is the expected future value (including interest earned), discounted at a

rate that takes into account the risk. Under the alternative quotation, confusingly, the current value would be determined by the inverse of the expected inverse of future value, multiplied by unity plus the required return. (If you just felt you had to read this sentence twice, you may want to consider reading technical note 3.1 at the end of the chapter instead.)

The direct (HC/FC) quoting convention used to be standard in continental Europe, and is called the “direct” quote, or the “right” quote. In the United States, a price with dimension USD/FC is called “American terms.” The alternative is called the “indirect” or “left” quote or, in the United States, “European terms.” Let’s see who uses which and why.

3.1.3 The Indirect Quoting Convention

One group of people using mostly indirect quotes are professional traders in the United States. Between 1944 and the mid 1980s, each and every exchange deal went through the USD; even when a German needed to buy CHF, the DEM would first be converted into USD and these dollars were then exchanged for CHF. Naturally, when New York traders talk to, say, their German counterparts, both must talk the same language, quotewise; otherwise too much time would be wasted inverting each other’s rates all the time. Both Germans and Americans actually preferred to quote in terms of DEM/USD rather than USD/DEM, for the simple reason that the official parities, set by the German government, were expressed in DEM/USD.² More generally, U.S. professionals use the exchange-rate convention as quoted in the other country. Thus, for countries that quote directly themselves, like Japan, New York traders would talk JPY/USD. But in the case of countries that quote indirectly themselves, like the United Kingdom, the pros would also use USD/GBP. Thus, U.S. pros use indirect quotes for countries that themselves quote directly, and direct quotes for countries that themselves quote indirectly.

As already hinted at, in the United Kingdom one uses the reverse quote, the number of foreign units that can be bought with one pound, or FC/HC. Some former British or Commonwealth countries (e.g., Australia and New Zealand) and, until 1979, Ireland do likewise.³ One reason is that, before World War I, the pound was the world’s reserve currency and played the role taken over by the dollar after World War II. In addition, until 1967 the GBP was still severely nondecimal—one pound consisted of twenty shilling, each worth twelve pence⁴—while currencies not based on the pound had gone decimal long before. It is much easier to multiply or divide by a decimal number, say

²Recall from the previous chapter that, until 1972, countries declared an official parity in relation to the USD, say DEM/USD 4. Intervention kept the actual rates between an upper and lower bound expressed, likewise, in DEM/USD.

³Canada and South Africa went off the pound ages ago, which is why they quote differently.

⁴Recall there also was a dollar (10s.), a crown (5s.), and a guinea, worth 21s. in the end; and in Elizabethan times many wages were expressed in marks (13s.4d., i.e., 160d.). But by modern times most prices were in pounds, shillings, and pence.

Table 3.1. Key exchange rates: pros' notation, dimensions, and nicknames.

Symbol	Currency pair	Dimension	Trading terminology
USDJPY	U.S. dollar, in Japanese yen	JPY/USD	Dollar yen
USDCHE	U.S. dollar, in Swiss francs	CHF/USD	Dollar Swiss or Swissy
USDCAD	U.S. dollar, in Canadian dollars	CAD/USD	Dollar Canada
USDZAR	U.S. dollar, in South African rand	ZAR/USD	Dollar ZAR or South African rand
GBPUSD	British pound, in U.S. dollars	USD/GBP	Cable
GBPCHF*	British pound, in Swiss francs	CHF/GBP	Sterling Swiss
GBPJPY*	British pound, in Japanese yen	JPY/GBP	Sterling yen
AUDUSD	Australian dollar, in U.S. dollars	USD/AUD	Australian dollar
NZDUSD	New Zealand dollar, in U.S. dollars	USD/NZD	New Zealand dollar or Kiwi
EURUSD	Euro, in U.S. dollars	USD/EUR	Euro
EURGBP*	Euro, in British pounds	GBP/EUR	Euro sterling
EURJPY*	Euro, in Japanese yen	JPY/EUR	Euro yen
EURCHF*	Euro, in Swiss francs	CHF/EUR	Euro Swiss
CHFJPY*	Swiss franc, in Japanese yen	JPY/CHF	Swiss yen
GLDUSD	Gold, in U.S. dollars per troy ounce	USD/ozXAU	Gold
SLVUSD	Silver, in U.S. dollars per troy ounce	USD/ozXAG	Silver

*Cross rate, from the U.S. perspective. Most names should be obvious, except perhaps CHF (*Confederatio Helvetica*, Latin for Switzerland—the way a four-language country solves a political conundrum). The ZAR, the South African rand, is not to be confused with SAR, the Saudi riyal. GLD and SLV are unorthodox: the official codes as used by, for example, Swift are XAU and XAG, with X signalling a nonstandard currency (also like the CFA franc and the ecu of old), and the Latin *aurum* and *argentum*. “Cable” for USDGBP refers to the fact that it is about bank-account money, with payment instructions wired by telegram cable rather than sent by surface mail. There was a time when telegraphy was cutting-edge technology.

FC/GBP 0.792 08, than with a number like £1/s5/d3 (one pound, five shillings, three pence). So everyone preferred to talk FC units per pound.

A third (and more recent) class of people using the indirect quote are the Eurolanders, who always quote rates as USD/EUR or JPY/EUR even though they traditionally quoted directly (as DEM/USD). Cynics conjecture that the Europeans may have coveted the reserve-currency status associated with an indirect quote. Another possible reason is that, initially, the euro was foreign to all existing currencies. For example, to Germans the euro was introduced as worth 2 DEM, so they would quite naturally introduce it to Americans and Japanese as being worth 1.20 USD or 110 JPY. When, eventually, the euro had become the home currency, the habit simply stuck.

Example 3.3. Have a look at table 3.1, showing the most important rates in the way they are always quoted by pros. The primary rates are in non-U.S. currencies except for the GBP, NZD, and AUD, or for the EUR; you know why. Cross rates for the EUR are in non-EUR currencies, and likewise for the GBP.

Table 3.2. Sample spot exchange rate quotes. (From the *Wall Street Journal Europe*, which sensibly shows both the natural and indirect quotes.)

AMERICAS	Per		In		EUROPE	Per		In	
	Per euro	In euros	U.S. dollar	U.S. dollars		Per euro	In euros	U.S. dollar	U.S. dollars
Argentino peso-a	3.9628	0.2523	3.0838	0.3243	Euro zone euro	1	1	0.7782	1.2851
Brazil real	2.9588	0.3380	2.3025	0.4343	Czech Rep koruna-b	28.260	0.0354	21.992	0.0455
Canada dollar	1.438	0.7073	1.1002	0.9089	Denmark krone	7.4576	0.1341	5.8034	0.1723
Chile peso	683.07	0.001464	531.55	0.001881	Hungary forint	262.82	0.003805	204.52	0.004890
Columbia peso	3186.28	0.0003138	2479.50	0.0004033	Malta lira	0.4294	2.3288	0.3342	2.9926
Ecuador US dollar-f	1.2850	0.7782	1	1	Norway krone	7.800	0.1282	6.0698	0.1648
Mexico peso-a	14.5307	0.0688	11.3075	0.0884	Poland zloty	3.9369	0.2540	3.0637	0.3264
Peru sol	4.2368	0.2360	3.2970	0.3033	Russia ruble-d	34.669	0.02884	26.979	0.03707
Uruguay peso-e	30.841	0.0324	24.000	0.0417	Slovak Rep koruna	37.7856	0.02647	29.4040	0.03401
U.S. Dollar	1.2850	0.7782	1	1	Sweden krona	9.2662	0.1079	7.2108	0.1387
Venezuala bolivar	2759.39	0.000362	2147.30	0.000466	Switzerland franc	1.5604	0.6109	1.2103	0.8262
...

a—floating rate b—commercial rate c—government rate d—Russian Central Bank rate f—Special Drawing Rights from the International Monetary Fund; based on exchange rates for U.S., British and Japanese currencies.

Note: Based on trading among banks in amounts of \$1million and more, as quoted by Reuters

Example 3.4. Look at the *Wall Street Journal Europe* excerpt in table 3.2, conveniently showing both quotes; the value in USD or EUR of one unit of the third (“foreign”) currency, and the value of 1 USD or EUR in units of that third (“foreign”) currency. The natural quote for Americans or Europeans would be the first one, but U.S. traders and Eurolanders may use the other quote. Take a minute to look at table 3.2 and see if you understand the exchange rates as quoted.

Q1. What is the dollar equivalent of one euro, according to the quotes in the *Wall Street Journal*?

A1. If your answer is USD 1.285, you are correct.

Q2. Determine the amount of Peruvian soles per EUR.

A2. If you answered 4.2368 soles per EUR, you are right.

3.1.4 Bid and Ask Rates

When you deal with foreign currency, you will discover that you pay a higher price at the time of purchase than when you sell one currency for another. For example, for dollar-rouble deals the currency booth in your hotel will quote two numbers, say RUB/USD 35–36. This means that if you sell USD for RUB, you receive RUB 35, while if you wish to buy USD you will have to pay RUB 36. The rate at which the bank will buy a currency from you is called the *bid* rate: they bid (i.e., they announce that they are willing to pay) 35 per dollar; and the rate at which the bank will sell a currency to you is the *ask* rate (they ask 36 per dollar). It is, initially, safer not to think about the meaning of bidding and asking because the words refer to the bank’s view, not yours. Just remember that *you* buy at the bank’s ask rate, and *you* sell at the bank’s bid rate. The bid is the lower quote, and ask is the higher one. The ask comes higher in the alphabet—use any trick that works, until you get used to it.

Indeed, if exchange rates are being quoted with the currency of interest—the currency you are buying or selling—in the denominator, then the ask rate will be higher than the bid rate. Obviously, it could not be the other way around: with a bid rate above the ask rate you would be able to make huge risk-free profits by buying at the ask and immediately reselling at the assumedly higher bid. No bank will allow you to buy low and then immediately resell at a profit without taking any risk, because your sure gains would obviously mean sure losses for the bank. In theory, there could still be room for a situation “bid rate = ask rate” (which offers no such arbitrage opportunities). Yet the real-world situation is invariably “bid rate < ask rate”: banks want to make some money from foreign-currency transactions.

Another way to think of this difference between the ask and the bid rates is that the difference contains the bank’s commission for exchanging currencies. The difference between the buying and selling rates is called the *spread*, and you can think of the bank’s implicit commission as being equal to half the spread. The following example explains why the commission is half of the spread rather than the spread itself.

Example 3.5. Suppose that you can buy CAD at RUB/CAD 38.6 and sell at RUB/CAD 38.0. With these rates, you can think of a purchase as occurring at the midpoint rate (RUB/CAD 38.3), grossed up with a commission of 0.30. Likewise, a sale can be thought of as a sale at the midpoint, 38.3, from which the bank withholds a commission of 0.30. Thus, the equivalent commission per one-way transaction is the difference between the bid (or ask) and the midpoint rate, that is, half the spread. (The spread itself would be the cost of a round-trip deal—buy and then sell.)

To get an idea of whether your house bank charges a low commission, you can ask for a two-way quote to see if the spread is small. If this is the case, you probably do not have to check with other banks. However, for large transactions, you should also compare the spot quotes given by different banks. (This will be examined further in section 3.3, especially section 3.3.3.) We discuss the determinants of spreads later, after we have described the market microstructure.

3.1.5 Primary Rates versus Cross Rates

As of 1945 and until well into the 1980s, all exchange rates in the wholesale segment were against the USD. They were and are called *primary rates*, while any rate not involving the USD would be called a *cross rate* and would traditionally be regarded as just implied by the primary rates. You will find an example for midpoint rates in table 3.3. The primary rates are in the first column (FC/USD) or the bottom line (USD/FC).⁵ The rest of the table is obtained

⁵ Many newspapers give currency j the j th row and the j th column instead of the $(N - j)$ th row and the j th column, but the layout is not crucial. The orientation of the empty diagonal (or the unit diagonal, as other tables might show it) is the sign to watch.

Table 3.3. Cross rates as in the *Wall Street Journal Europe*.

Cross rates U.S. dollar and euro foreign-exchange rates in global trading												
	USD	GBP	CHF	SEK	RUB	NOK	JPY	ILS	EUR	DKK	CDN	AUD
Australia	1.3253	2.4818	1.0915	0.1838	0.0491	0.2183	0.0118	0.2934	1.7031	0.2284	1.2046	—
Canada	1.1002	2.0603	0.9061	0.1526	0.0408	0.1813	0.0098	0.2436	1.4038	0.1896	—	0.8302
Denmark	5.8034	10.867	4.7794	0.8048	0.2151	0.9561	0.0518	1.2847	7.4576	—	5.2748	4.3789
<i>Euro</i>	0.7782	1.4573	0.6409	0.1079	0.0288	0.1282	0.0069	0.1723	—	0.1341	0.7073	0.5872
Israel	4.5173	8.4592	3.7202	0.6265	0.1674	0.7442	0.0403	—	5.8049	0.7784	4.1058	3.4085
Japan	112.11	209.93	92.325	15.547	4.1553	18.469	—	24.817	144.06	19.317	101.90	84.589
Norway	6.0698	11.367	4.9988	0.8418	0.2250	—	0.0541	1.3437	7.8000	1.0459	5.5170	4.5800
Russia	26.978	50.521	22.218	3.7414	—	4.4448	0.2407	5.9724	34.669	4.6488	24.521	20.357
Sweden	7.2108	13.503	5.385	—	0.2673	1.1880	0.0643	1.5963	9.2662	1.2425	6.5541	5.4409
Switzerland	1.2145	2.2739	—	0.1684	0.0405	0.2000	0.108	0.2688	1.5604	0.2092	1.1037	0.9162
U.K.	0.5340	—	0.4398	0.0741	0.0198	0.0880	0.0048	0.1182	0.6862	0.0920	0.4854	0.4029
U.S.	—	1.8726	82.236	0.1387	0.0371	0.1648	0.0089	0.2214	1.2850	0.1723	0.9089	0.7546

The numbers in the “EUR” column, for instance, show the values of the euro in other currencies (that is, the EUR acts as FC), while those in the “EUR” row show the values of the other currencies in euros (that is, the EUR acts as HC).

by division or multiplication: $GBP/EUR = GBP/USD \times USD/EUR$, for example. Each of the resulting new rows or columns is a set of quotes in HC/FC (row) or FC/HC (column). With 12 currencies you have 144 entries, of which 12 are on the information-free diagonal, and half of the remaining 132 are just the inverses of the others.

We have a whole section on the relation between primary and cross rates in the presence of spreads, so at this stage we just consider why, among pros, there were until the 1980s just primary rates. There were several reasons:

- Official parities were against the USD; there was no official parity (in the sense of being defended by any central bank) for rates against other currencies.
- The USD market had the lowest spreads, so all real-world transactions would effectively be done via the dollar anyway. That is, pounds were converted into marks by buying dollars first and then exchanging these for marks, for example, because that was the cheapest way to do so (see below). The cross-rate would just be the rate implied by the two primary rates used in the transaction.
- In pre-electronic days it would be quite laborious to keep track of, say, a 30×30 matrix of cross rates with 435 distinct meaningful entries, making sure all cross rates are consistent with the primary ones all the time. So rather than quoting cross rates all the time, banks just showed primary quotes and then computed cross rates if and when needed.

By the 1980s desktop computers were ubiquitous and, for many pairs of “big” currencies the volume of cross transactions had become large enough to make direct cross exchanges competitive compared with exchanges via the USD. Official exchange rates were gone in many cases, or in the ERM case had become multilateral. So we now see explicit quotes for some of the cross rates. Look at

figure 3.4 to see what rates have active multilateral electronic markets—a good indication of there being a reasonable volume. Note also that for some new EU members the market against the EUR works well while the market against the USD lacks liquidity; that is, for these countries the rate against the euro is economically the key one, even though Americans would regard it as just a cross rate.

3.1.6 Inverting Exchange Rates in the Presence of Spreads

The next issue is how a pair of quotes for one currency can be translated into a pair of quotes for a different currency. The rule is that the inverse of a bid quote is an ask quote, and vice versa. To conceptualize this, consider the following illustration.

Example 3.6. An Indian investor wants to convert her CAD into USD and contacts her house bank, Standard Chartered. Being neither American nor Canadian, the bank has no natural preference for either currency and might quote the exchange rate as either USD/CAD or CAD/USD. The Indian bank would make sure that its potential quotes are perfectly compatible. If it quotes from a Canadian viewpoint, the bank gives a CAD/USD quote (which says how many CAD the investor must pay for 1 USD, for instance, CAD/USD 1.5). If it uses the U.S. perspective, the bank gives a USD/CAD quote, which says how many USD the U.S. investor gets for 1 CAD, 0.666 67.

The bank's alternative ways of quoting will be fully compatible if

$$S_{\text{bid},t}^{\text{CAD/USD}} = \frac{1}{S_{\text{ask},t}^{\text{USD/CAD}}}, \quad (3.2)$$

$$S_{\text{ask},t}^{\text{CAD/USD}} = \frac{1}{S_{\text{bid},t}^{\text{USD/CAD}}}. \quad (3.3)$$

To fully understand this, recall that what looks like buying (at the ask) to a U.S. resident looks like selling to a Canadian, at the Canadian's bid. Alternatively, recall that the ask is the higher of the two quotes. But if you invert two numbers, the inverse of the larger number will, of course, be smaller than the inverse of the smaller number. Because the inverse of a larger number is a smaller number, the inverse ask must become the bid, and vice versa.

Example 3.7. Suppose that you read the following quote on the Reuters screen: USD/CAD 1.000–1.005.

Q1. What is the bank's buying and selling rate for CAD?

A1. The bank's buying rate for CAD is USD 1.000 and its selling rate is USD 1.005; that is, *you* sell CAD at USD 1.000 and buy at 1.005.

Q2. What, therefore, are the bank's buying and selling rates for USD (in CAD)?

A2. The bank's buying rate or bid for USD is $1/1.005 = \text{CAD/USD } 0.995\ 025$ (probably rounded to 0.9950) and the selling rate or ask is $1/1.000 = 1.000$; that is, wearing your Canadian hat, *you* sell USD at CAD 0.9950 and buy at 1.000.

One corollary is that in countries like the United Kingdom, where the reverse or indirect quote is used, the rate relevant when you buy is the lower of the two, while the higher quote is the relevant rate when you sell. Thus, it is important to be aware of what the foreign currency is, and what convention is being used for quoting the exchange rate. Again, it is always easier and more convenient to have the foreign currency in the denominator. That way the usual logic will work: banks buy low and sell high.

3.2 Major Markets for Foreign Exchange

In this section, we describe the size and structure of the exchange market and the types of transactions one can make in this market.

3.2.1 How Exchange Markets Work

The foreign exchange market is not an organized market. Stock markets or futures markets are: they have fixed opening hours, a more or less centralized mechanism to match supply and demand, standardized contracts, an official publication channel for data on volumes and prices, and a specific location or one designated group of computers running everything. In contrast, the exchange market consists of a wholesale tier, which is an informal network of about 500 banks and currency brokerages that deal with each other and with large corporations, and a retail tier, where you and I buy and sell foreign exchange. At any point in time, wholesale exchange markets on at least one continent are active, so that the worldwide exchange market is open twenty-four hours a day (see figure 3.1). Until the mid 1990s, most interbank dealing was done over the telephone; most conversations were tape-recorded, and later confirmed by mail, telex, or fax. Reuters, which was already omnipresent with its information screens, and EBS⁶ have now built computer networks which allow direct trading and which now largely replace the phone market. The way the computer systems are used depends on the role the bank wants to play. We make a distinction between deals via (i) market makers, (ii) auction platforms, or (iii) brokers.

⁶Electronic Broking Services (EBS) was created by a partnership of the world's largest foreign exchange market-making banks. Over approximately USD 150 billion in spot foreign exchange transaction and hundreds of tonnes of gold and silver are traded every day over the EBS Spot Dealing System. It was created in 1990 to challenge Reuters's threatened monopoly in interbank spot foreign exchange and provide effective competition. ICAP Plc, the world's largest broker of transactions between banks, agreed in 2006 to buy EBS.




Figure 3.1. Trader activity over the day. Graph courtesy of Luc Bauwens, Université Catholique de Louvain. The graph shows, per 5 min interval over 24 hours, the evolution of the average number of indicative quotes entered into the Reuters FX/FX pages. Time is GMT in summer, GMT+1 in winter; that is, European time is $t + 2$ h, London +1, New York -4 h; Sydney and Tokyo time are at $t + 10$ and $t + 9$ h, respectively. Below I describe working days as 8:00–17:00, but many a trader starts earlier and/or works later. At 0:00, when the morning shift in Sydney has been up and running for about two hours and Tokyo for one hour, Hong Kong starts up, to be followed by Singapore in one hour and Bahrain in three hours. Between 6:00 and 8:00 the Far East bows out but Western Europe takes over: first the continent (6:00 GMT), then London (7:00); activity soars. A minor dip follows around the European noon but activity recovers again in the afternoon, peaking when New York takes over (12:00) and Europeans close their positions (15:00 on the continent, 16:00 in London). New York does less and less as time passes. By 22:00 Sydney is starting up, and Tokyo is preparing breakfast.

3.2.1.1 Market Making

Many players in the wholesale market act as market makers. If a market-making credit agreement between two banks has been signed, either party undertakes to provide a *two-way* quote (bid and ask) when solicited by the other party, without even knowing whether that other party intends to buy, or rather sell. Such a quote is *binding*: market makers undertake to effectively buy or sell at the price that was indicated.




Figure 3.2. A Reuters conversation and an EBS broking window. In the Reuters conversation window, GENP is an abbreviated name (Jenpi, Jean-Pierre); he asks for a quote for EUR in USD for quantity 5m (dollars); pse is GENP’s code for “please.” The counterparty answers by keying in the small numbers, and Jenpi replies he buys 5 (million) at the ask, 27, for value date September 5. The counterparty closes with “Thanks, I’ll send the dollars to your correspondent, Citibank. Bye bye.” The second picture shows part of an EBS broking screen. On top, the current date and time. Next line: the spot delivery date, September 5. For two currencies you then see in small font the “big” figure (the part of the quote that is usually omitted) and in big font the “small” quotes: bid and ask, each preceded/followed by the quantity available, in millions. Thus, somebody bids 1.2825 for 3 million dollars, another party offers 5 million dollars at 1.2827.

Example 3.8. Deutsche may ask Hong Kong and Shanghai for a quote of EUR against USD. HSBC must then provide a bid and an ask without knowing the direction of Deutsche’s possible trade; and if Deutsche replies with “I buy 10 million” then HSBC must sell that quantity at the price they quoted.

Of course there are limits to the market makers’ commitments to their quotes. First, potential customers should decide almost immediately whether to buy (“mine”), or to sell (“yours”), or not to deal; they cannot invoke a quote made, say, three minutes ago. Second, if the intended transaction exceeds a mutually agreed level, laid down in the prior credit agreement—say USD 25m—market makers can refuse. For larger transactions, the trader asking for a quote should reveal immediately what the size of the transaction will be. Third, the credit agreement also provides a limit to the total amount of open contracts that can be outstanding between the two banks at any moment;⁷ if the limit is reached, no more deals are allowed.

Transactions via binding two-way quotes are typically concluded on computers, by means of chatting windows (more grandly called “conversations”). Bank A’s trader X clicks his conversation window with trader Y at bank B—there may be up to 64 such windows open at any given point of time—and might type in, for instance, PLS EUR/USD, meaning “please provide a quote for the EUR, in USD.” Player A can also mention the quantity, in millions. The millions are omitted—that is, 5 means five million—and the quantity bears on the currency in the denominator, traditionally the USD or the GBP. B’s trader may answer, for instance, 13-16, meaning that (the last two digits of) her bid and ask are 13 and 16. (Traders never waste time by mentioning the leading

⁷Exchange transactions are settled with a delay of at least two days, so each contract remains outstanding for at least two days; many live much longer. See section 3.2.3.

numbers: everybody knows what these are. Only the “small” numbers are mentioned.) The first party can let the offer lapse; if not, he answers MINE or YOURS, mentions the quantity if not already indicated, and hits the SEND key. The deal is done, and both traders now pass on the information to their “back office,” which enters the data into the information systems. The back offices will also check with each other to see whether the inputs match; with the logs of the conversations, disputes are of course far less likely than before, when everything went by phone and when traders handed down hand-scribbled “tickets” to the accountants who then checked with each other via telexes. Voice deals still exist, but they are getting rarer.

3.2.1.2 Implications of Market Making for the Size of the Bid-Ask Spread and the Maximum Order Size

Normally, the lower the volume in a particular market, the higher the spread. Also, during holidays, weekends, or lunch breaks, spreads widen. Spreads are also higher during periods of uncertainty, including at the open and close of the market each day. Maximum order quantities for normal quotes follow a similar pattern: a market maker is prepared to handle large lots if the market is liquid (thick) or the volatility low.

All these phenomena are explained by the risk of market making. Notably, if a customer has “hit” a market maker, the latter normally wants to get rid of that new position quickly. But in a thin or volatile market, the price may already have moved against the market maker before he or she was able to close out; thus, the market maker wants a bigger commission as compensation for the risk, and puts a lower cap on the size of the deals that can be executed at this spread. For the same reason, quotes for an unusually large position are wide too: getting rid of a very large amount takes more time, during which anything could happen. At the retail end of the market, in contrast, the spread increases for smaller transactions. This is because 100 small transactions, each for USD 100,000, cost more time and effort than one big transaction of USD 10m.

For high-volume currencies like the USD/EUR, the difference between one market-maker’s own bid and ask is often as low as three basis points (in a quote of four or five digits, like 1.2345 or 0.9876), and the difference between the best bid (across all market makers) and the best ask (also across all market makers) may be just two or one or, occasionally, zero basis points. See section 3.3.2 for more information on quoting behavior.

Table 3.3 shows the minimum and maximum amounts quoted by an internet dealer;⁸ they are smaller than interbank (and spreads are bigger than interbank), but you can still notice how the maximum amounts and the spreads relate to each other, presumably both reflecting liquidity and volatility.

⁸ “Size of 1.0 lot” (about 1m USD) shows the minimum, which is clearly targeting players out of the interbank league (where the lot size is 1m) but still above the micro-investor’s league. “Instant execution” is the maximum amount you can buy or sell at the trader’s regular quotes.

Ticker	Size of 1.0 lot	Instant execution	Spread	Limit and stop levels	March 9, 2007 rate (in pips)	Spread (‰)
EURUSD	EUR 100,000	up to 10M	2 pips	2 pips	13,115	1.5
GBPUSD	GBP 100,000	up to 10M	3 pips	3 pips	19,319	1.6
EURCHF	EUR 100,000	up to 5M	3 pips	3 pips	16,163	1.9
EURJPY	EUR 100,000	up to 10M	3 pips	3 pips	15,489	1.9
USDJPY	USD 100,000	up to 10M	3 pips	3 pips	11,810	2.5
GBPCHF	GBP 100,000	up to 5M	7 pips	7 pips	23,810	2.9
EURGBP	EUR 100,000	up to 5M	2 pips	2 pips	6,788	2.9
GBPJPY	GBP 100,000	up to 5M	7 pips	7 pips	22,817	3.1
USDCHF	USD 100,000	up to 10M	4 pips	4 pips	12,325	3.2
USDCAD	USD 100,000	up to 5M	4 pips	4 pips	11,735	3.4
AUDUSD	AUD 100,000	up to 5M	3 pips	3 pips	7,802	3.8
CHFJPY	CHF 100,000	up to 5M	4 pips	4 pips	9,583	4.2
EURCAD	EUR 100,000	up to 3M	8 pips	8 pips	15,389	5.2
NZDUSD	NZD 100,000	up to 2M	5 pips	5 pips	9,583	5.2
USDSGD	USD 100,000	up to 1M	8 pips	8 pips	15,267	5.2
EURAUD	EUR 100,000	up to 5M	10 pips	10 pips	16,810	5.9




Figure 3.3. Order limits and spreads for various rates, semi-professional. The table shows conditions for various currencies from a particular internet broker. The minimum and maximum quantities are not interbank, but still aim at semi-professionals or perhaps day traders rather than pop and mom investors, the hardcore retail. The spread and the tick size for limit and stop levels are likewise wider than interbank. Do note how the spread varies depending on liquidity and the level of the rate, and on how the maximum order size (imperfectly) relates to the spread (graph). *Source:* www.alpari.co.uk/en/cspec/ for columns 1-5; the *Wall Street Journal Europe*, March 12, 2007, for column 6; spread in basis points has been added. Data have been rearranged by increasing relative spread. For the graph the order sizes have been converted from reference currency (the FC in the quote) to USD.

3.2.1.3 Auctioning Off through a Broking System

All the above was about market making. Beside these purely bilateral deals—the successors to bilateral phone conversations—there nowadays are increasingly many semi-multilateral deals. If a trader actively wants to buy, or sell, she may enter a limit order into EBS's or Reuters's limit-order book rather than calling a number of market makers or waiting until someone else calls her. This is comparable to you offering, say, a used car for sale on eBay rather than calling various car dealers or posting a sign on your door and then waiting until someone rings your bell. For instance, bank A may have EUR 30m for

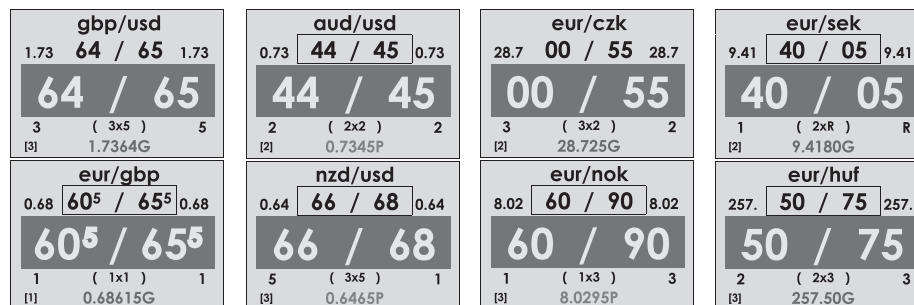


Figure 3.4. A panel of Reuters broking windows. The entries should by now be obvious, except the bottom line, which shows the last trade (quantity and price).

sale and want at least USD/EUR 1.3007 for them—an ask price. The bank posts this info, for instance, on Reuters’s “3000” system. Reuters’s window, at any moment, then shows the best bid across all “buy” limit orders, and the best ask among all “sell” limit orders outstanding at that moment. For instance, on Reuters’s 3000 screen a line EUR/USD 10–11 3xR means that the highest bid posted at that very moment is 10, the lowest ask 11, and that the quantities for these limit orders are, respectively, 3 and “a number exceeding 50” (= R).⁹ You see the EBS counterpart of Reuters 3000 in figure 3.2. Any party interested in one of these offers can then click on the quote they like (either the bid or ask) and specify the quantity taken. Or another bank may enter a limit order that is automatically matched, wholly or partly, with an already outstanding limit order. Reuters’s computer then informs the IT systems of both banks of the transactions that were concluded so that no more human intervention with “tickets” and telexes and faxes is needed (*straight through processing* (STP)).

The decision by an FX trader whether to use EBS or Reuters Dealing 3000 (also known as D2) is driven largely by currency pair. In practice, EBS is used mainly for EUR/USD, USD/JPY, EUR/JPY, USD/CHF, and EUR/CHF, and Reuters D2 is used for all other interbank currency pairs. Have a look at table 3.4 to see who leads where. In these multilateral electronic dealing systems, the spread for EUR/USD is typically one pip, that is, one hundredth of a USD cent. (Online currency brokers targeting private investors typically offer a two-pip spread; just feed “foreign exchange” into your Web search engine to find these brokers.) For other exchange rates spreads are often wider.

Note that the advent of these multilateral systems has made the market somewhat more like an organized market: there is centralization of buy and sell orders into one matching mechanism, there are membership rules (not anyone can log on to the program), rules about orders, etc. But the exchange market is still fully private, whereas many exchanges are semi-official institutions that are heavily regulated and need, at least, a license.

⁹The quotes are, again, “small numbers” and the quantities mean millions of dollars. Remember also that, for traders, EUR/USD means “the value of the euro in dollars.”

Table 3.4. EBS versus Reuters D2: who leads, who follows, who fails.

	Primary		Other cross rates		Cross against EUR	
	EBS	Reut.	EBS	Reut.	EBS	Reut.
EUR/USD	+	+/-	AUD/JPY	-	EUR/AUD	-
AUD/USD	+/-	+	AUD/NZD	+/-	EUR/CAD	-
GBP/USD	+/-	+	CHF/JPY	-	EUR/CHF	+ +/-
NZD/USD		+	GBP/JPY	-	EUR/CZK	- +
USD/CAD	+/-	+			EUR/DKK	- +
USD/CHF	+	+/-			EUR/GBP	+/- +
USD/CZK		-			EUR/HUF	- +
USD/DKK		-			EUR/ISK	- +
USD/HKG	-	+			EUR/JPY	+ +/-
USD/HUF		-			EUR/NOK	- +
USD/ILS		+			EUR/NZD	- +/-
USD/INR		+			EUR/PLN	- +
USD/ISK		-			EUR/RON	- -
USD/JPY	+	+/-			EUR/SEK	- +
USD/MXN	-	+			EUR/SKK	- +
USD/NOK		-			EUR/TRY	- -
USD/PLN	-	+			EUR/ZAR	- +/-
USD/RON		-				
USD/RUB		+				
USD/SAR		+				
USD/SEK		-				
USD/SGD	-	+				
USD/THB		+				
USD/TRY		+				
USD/ZAR		+				

+, Primary liquidity source; +/-, supported, but liquidity not good or not stable; -, supported but not used in practice. Rates are expressed following the "name" convention, not the dimensions. *Source:* www.londonfx.co.uk/autobrok.html, accessed February 2007.

3.2.1.4 Brokers

One last way of shopping around in foreign exchange markets is through currency brokers. In the telephone-market days, brokers used to do the go-between stuff that nowadays is handled via limit-order books: on behalf of a bank or company, the broker would call many market makers and identify the best counterpart. Roughly half of the transaction volume in the exchange market used to occur through brokers. Nowadays, brokers are mainly used for unusually large transactions, or "structured" deals involving, say, options next to spot and/or forward; for bread-and-butter deals their role is much reduced.

BIS is commonly described as the bank of the central banks. It was first set up after World War I to act as a payment agent distributing the German and Austrian war reparation payments. After World War II it ran the European Payment Union (EPU), serving as a netting institute for payments among EPU members. By netting the international payments, the volume of actual payments was reduced, which alleviated the problems of dollar shortages in the first years after the war. Currently, the BIS is still the bank of the central bankers: all central banks have accounts there, in various currencies, and can route their payments

to each other via the BIS. But nowadays the BIS mainly serves as a talking club for central bankers and regulators. One of its missions is to gather data on exchange markets, euro and OTC markets, new financial instruments, bank lending to sovereign borrowers, and so on. Another mission is to provide a forum where regulators coordinate the capital adequacy rules that they impose on financial institutions. The Basel I rules covered credit risk—in a crude way, perhaps, but it was a useful first step; the recent Basel II rules refine Basel I and add market-price risks; see the chapter on Value-at-Risk.

Panel 3.1. The Bank of International Settlements (BIS).

3.2.2 Markets by Location and by Currency

Every three years, in April, the Bank of International Settlements (panel 3.1) makes a survey of the over-the-counter (OTC) markets, including forex. At the latest count, April 2007, the daily volume of trading on the exchange market and its satellites—futures, options, and swaps—was estimated at more than USD 3.2 trillion. This is over 45 times the daily volume of international trade in goods and services, 80 times the United States's daily GDP, 230 times Japan's GDP, and 400 times Germany's GDP, and 7,500 times the world's official development-aid budget.¹⁰ The major markets were, in order of importance, London, New York, Tokyo, and Frankfurt (the European Central Bank's home base). London leads clearly, easily beating even New York, Tokyo, and Singapore taken together, and still increasing its market share. Frankfurt is a fast riser but from a low base.

The most important markets, per currency, are the USD/EUR and the USD/JPY markets; together they represent almost half of the world's trading volume. Add in the GBP, and the transactions involving just the top four moneys represent two thirds of all business. The USD still leads: in 88% of transactions it takes one of the sides (down from 90% in 2004), while the EUR is one of the two currencies in less than 40% (up from 35%) of that volume—and the bulk of that is USD/EUR trade.

3.2.3 Markets by Delivery Date

The exchange market consists of two core segments: the spot exchange market and the forward exchange market.

The *spot* market is the exchange market for quasi-immediate payment (in home currency) and delivery (of foreign currency). For most of this text we shall denote this spot rate by S_t , with t referring to current time. In practice,

¹⁰All data are from the CIA Factbook. Trade and aid, 2004; GDP, early 2007 estimates for 2006.




Figure 3.5. Forex turnover, daily, billions of USD, and market shares of currency pairs. Source: BIS, *Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity in April 2007*, Preliminary Global Results, September 2007.

“quasi-immediate” means “right now” only when you buy or sell notes or coins. (This section of the market is marginal.) For electronic money (that is, money that will be at your disposal in some bank account), delivery is in two working days for most currencies (“ $t + 2$ ”), and one day between Canada and the United States or between Mexico and the United States (“ $t + 1$ ”). Thus, if you buy AUD 2m today, at AUD/EUR 2, the AUD 2m will be in your account two working days from now, and the EUR 1m will likewise be in the counterpart’s account two days from now. The two-day delay is largely a tradition from the past, when accounts were kept by hand. The hour of settlement depends on the country, but tends to be close to noon. Thus, the EUR side of an EUR/USD transaction is settled in Europe about six hours before the USD leg of the deal is settled in New York.¹¹

The *forward* market is the exchange market for payment and delivery of foreign currency at some future date, say, three months from now. For example, supposing today is January 3, you could ask your bank to quote you an exchange rate to sell dollars for pounds for a date in March, say March 5, and the transaction would be settled on that date in March, at the rate agreed upon on January 3 (irrespective of the spot rate prevailing on March 5). The forward market, in fact, consists of as many subsegments as there are delivery dates, and each subsegment has its own price. We shall denote this forward rate by $F_{t,T}$, with T referring to the future delivery date. (Forward rates and their uses will be discussed in great detail in chapters 4 and 5.)

The most active forward markets are for 30, 90, 180, 270, and 360 days, but nowadays bankers routinely quote rates up to ten years forward, and occasionally even beyond ten years. Note that months are indicated as 30 days.

¹¹ This leads to the risk that, between the two settlement times, one party may file for bankruptcy or be declared bankrupt. This is called “Herstatt risk,” after a small German bank that pulled off this feat on June 26, 1974. Nowadays, regulators close down banks *outside* working hours.

Table 3.5. Market shares. Percentage for foreign exchange trading.

	U.K.	U.S.	Japan	Singapore	Other
1998	32.5	17.9	6.9	7.1	35.6
2001	31.2	15.7	9.1	6.2	37.8
2004	31.3	19.2	8.3	5.2	36.0
2007	34.1	16.6	6.0	5.8	37.5

Source: BIS, *Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity in April 2007*, Preliminary Global Results, September 2007.

In principle, a 30-day contract is settled one month later than a spot contract, and a 180-day forward contract is settled six months later than a spot contract—each time including the two-day initial-delay convention.¹²

Example 3.9. A 180-day contract signed on March 2 works as follows. Assuming that March 4 is a working day, spot settlement would have been on March 4. For a 180-day forward deal, the settlement date would be moved by six months to, in principle, September 4, or the first working day thereafter if a holiday. The actual number of calendar days is at least (2+)184 days: there are four 31-day months in the March–September window.

The above holds for standard dates, But you can always obtain a price for a “*broken date*” (i.e., a nonstandard maturity), too. For instance, on April 20 you can stipulate settlement on November 19 or any other desired date.

Worldwide, spot transactions represent less than 50% of the total foreign-exchange market volume. The forward market together with the closely related swap market (see chapter 7) make up over 50% of the volume. About 3% of total trade consists of currency-futures contracts (a variant of forward contracts traded in secondary markets—see chapter 6) and currency options (see chapter 8).

After this digression on the meaning of exchange rates and their relation to real quantities, we now return to the operations of the spot exchange market. We want to introduce one of the cornerstones of finance theory: the law of one price.

3.3 The Law of One Price for Spot Exchange Quotes

In frictionless markets, two securities that have identical cash flows must have the same price. This is called the law of one price. There are two mechanisms that enforce this law. The first one is called *arbitrage* and the second one could be called *shopping around*. We explain these two concepts below.

Suppose that two assets or portfolios with identical cash flows do not have the same price. Then any holder of the overpriced asset could simultaneously

¹²Further details of settlement rules are provided in Grabbe (1995).

In a short sale you hope to be able to buy low and sell high, but with the selling preceding the buying, unlike in a long position. Thus, a short seller hopes to make money from falling prices rather than from rising prices.

In markets with delivery a few working days later, you can always go short for a few hours: sell “naked” in the morning, for instance, and then buy later within the same day so as to be able to deliver n days later.

For longer horizons one needs more. In the case of securities, short selling then requires borrowing a security for, say, a month and selling it now; at the end of the month you then

buy back the number of securities you borrowed and restate them to the asset lender, including dividends if any were paid out during that period.

For currencies, longer-term short selling can be done by just borrowing forex and selling it, hoping to be able to buy back the forex (including interest) later at a lower price. If there is a forward market, lastly, going short is even easier: promise to deliver on a future date at a price that is fixed now. If prices have dropped by then, as you hope, you will be able to close out (buy spot) cheaply and make money on the forward deal.

Panel 3.2. What is short selling?

sell this asset and buy the cheaper asset instead, thus netting the price difference without taking on any additional risk. If one does not hold the overpriced asset, one could still take advantage of this mispricing by short selling (panel 3.2) the overpriced asset and covering this with the purchase of the cheaper security. For example, you sell an overpriced asset at 1.2135 and buy a perfect substitute at 1.2133, netting 0.0002 per unit right now and no net cash flow at T . Such transactions are called arbitrage. These arbitrage transactions generate an excess supply of the overpriced asset and an excess demand for the underpriced asset, moving the prices of these two assets toward each other. In frictionless markets, this process stops only when the two prices are identical. Note that apart from the arbitrage gain, an arbitrage transaction does not lead to a change in the net position of the arbitrageur; that is, it yields a sure profit without requiring any additional investment.

The second mechanism that enforces the law of one price is shopping around. Here, in contrast to arbitrage, investors do intend to make particular changes in their portfolios. Shopping around has to do with the fact that, when choosing between different ways of making given investments, clever investors choose the most advantageous way of doing so. Therefore, when choosing between assets with identical cash flows, investors buy the underpriced assets rather than the more expensive ones. Likewise, when choosing which assets to sell, investors sell the overpriced ones rather than the ones that are relatively cheap. This demand for the underpriced assets and supply of the overpriced ones again leads to a reduction in the difference between the prices of these two securities.

Although the arbitrage and shopping-around mechanisms both tend to enforce the law of one price, there are two differences between these mechanisms.

- First, an arbitrage transaction is a round-trip transaction. That is, you buy and sell, thus ending up with the same position with which you started.

As arbitrage requires a two-way transaction, its influence stops as soon as the price difference is down to the *sum* of the transaction costs (buying and selling). In contrast, in shopping around one wishes to make a particular transaction, and the issue is which of the two assets is cheaper to trade.¹³ As a result, the influence of shopping around can go on as long as the price difference exceeds the *difference* of the two transactions costs.¹⁴

- Second, arbitrage is a strong force because it does not require any capital and can, therefore, generate enormous volumes. In contrast, shopping around can be a price-equalizing mechanism only if there are investors who wish to make that particular transaction. This exogenously triggered volume, if any, is always finite and may be exhausted before it has actually equalized the prices.

In this section, we apply these arguments to spot rates quoted for the same currencies by different market makers. In a perfect exchange market with zero spreads, arbitrage implies that the rate quoted by bank X must equal the rate quoted by bank Y: there can be only one price for a given currency—otherwise, there is an arbitrage opportunity.

Example 3.10. If Citibank quotes DEM/USD 1.6500, while Morgan Chase quotes DEM/USD 1.6501, both at zero spreads, then there are two possibilities:

- There is an arbitrage opportunity. You can buy cheap USD from Citibank and immediately sell to Morgan Chase, netting DEM 0.0001 per USD. You will, of course, make as many USD transactions as you can. So will everybody else. The effect of this massive trading is that either Citibank or Morgan Chase, or both, will have to change their quotes so as to stop the rapid accumulation of long or short positions. That is, situations with arbitrage profits are inconsistent with equilibrium, and are eliminated very rapidly.
- There is also a shopping-around pressure. All buyers of USD will buy from Citibank, and all sellers will deal with Morgan Chase.

The only way to avoid such trading imbalances is if both banks quote the same rate.¹⁵

What we now want to figure out is how arbitrage works when there are bid-ask spreads. The point is that, because of arbitrage, the rates cannot be

¹³Accordingly, Deardorff (1979) refers to standard arbitrage as *two-way arbitrage* and to shopping around as *one-way arbitrage*.

¹⁴Denote by P_U and k_U the price and transaction cost when dealing in the underpriced asset, and denote by P_O and k_O the counterparts for the overpriced asset. The advantage of buying the cheap asset rather than the expensive one remains positive as long as $P_U + k_U < P_O + k_O$; that is, as long as $P_O - P_U > k_U - k_O$. In contrast, the advantage of buying the cheap asset and selling the expensive one remains positive as long as $P_O - k_O - (P_U + k_U) > 0$, that is, as long as $P_O - P_U > k_O + k_U$: you pay both costs instead of replacing one by another.

¹⁵This is often put as “by arbitrage, the quotes must be the same” or “arbitrage means that the quotes must be the same.” Phrases like these actually mean that to rule out arbitrage opportunities, the quotes must be the same.




Figure 3.6. Arbitrage and shopping-around opportunities across market makers.

systematically different; and if the quotes do differ temporarily, they cannot differ by too much.

3.3.1 Arbitrage across Competing Market Makers

Suppose bank X quotes you INR/NZD 20.150–20.158 while bank Y quotes INR/NZD 20.160–20.168. If you see such quotes, you can make money easily: just buy NZD from bank X at INR 20.158, immediately resell it to bank Y at INR 20.160, and pocket a profit worth INR 0.002 for each NZD. Note two crucial ingredients: (1) you are not taking any risk, and (2) you are not investing any capital since the purchase is immediately reversed and both transactions are settled on the same day. The fact that you immediately reverse the transaction explains why this is called arbitrage.

If such quotes are found in the exchange market (or elsewhere, for that matter), large trades by a few alert dealers would immediately force prices back into line. The original quotes would not be equilibrium quotes. In equilibrium, the arbitrage argument says that you cannot make money without investing capital and without taking risk. Graphically, any empty space between the two quotes would correspond to an arbitrage profit. Thus, the no-arbitrage condition says that any two banks' quotes should not be separated by empty space; that is, they should overlap by at least one point, like the quotes X' and Y in figure 3.6.

3.3.2 Shopping Around across Competing Market Makers

Shopping-around activity implies that small differences like those between the pair (X', Y) in figure 3.6 will not persist for very long. Rather, the two quotes will sometimes be the same, and, if at other times one bank is more expensive, then this would say very little about what the situation will be five minutes later. To see this, suppose that bank X' quotes INR/NZD 20.55–20.63, while bank Y quotes INR/NZD 20.60–20.68. In such a situation, all buyers of NZD will, of course, prefer to deal with bank X', which has the lower ask rate (20.63 instead of 20.68), while all sellers will now deal with bank Y, which has the better bid rate (20.60 instead of 20.55). It is conceivable that these banks actually want

this to happen, for instance, if bank X' has an excess of foreign currency (long), and bank Y is short of forex and wants to replenish its FC inventory. But we would not expect this to be a long-run phenomenon. It is true that very often a bank may want one type of transaction only, but situations like that must change very rapidly because otherwise that bank's position would become unacceptably large and risky.

Example 3.11. Suppose you see five banks quoting EUR against USD, as follows:

Citibank	USD/EUR	1.3450-52
Bank of America	USD/EUR	1.3450-52
Continental Bank	USD/EUR	1.34 <u>51</u> -53
Deutsche Bank	USD/EUR	1.3450-52
Banca da Roma	USD/EUR	1.3449- <u>51</u>

Q1. Which bank(s) is (are) keen on buying EUR? Keen on selling EUR? Not interested in dealing?

A1. Continental, with its high bid, is quite attractive to sellers, so this trader clearly wants to buy—for example, to fill a short position or because she expects a rise. Roma, in contrast, judging by its low ask, is quite attractive to buyers, so their trader clearly wants to sell—maybe to move an unwanted long position, or in anticipation of a fall in the rate. The others are just twiddling thumbs: as things stand, they are unwilling to match Continental's or Roma's rates, and they hope that things will soon be better.

Q2. Why does Continental raise both its bid and its ask, rather than just its bid?

A2. Apparently it wants not just to attract sellers but also to scare off buyers. Similarly, Roma does not just fancy buyers, but does not want any sellers at all.

Q3. If we were to look at these banks' quotes every five minutes, would we always expect to see the same pattern, i.e., Continental quoting higher and Roma lower than the majority?

A3. Of course not: as soon as their desired positions are reached, they will return to the fold. Thus, the top and bottom positions are picked by a particular bank for only a brief period, and move randomly across the list of banks.

3.3.3 Triangular Arbitrage

Now that we know how exchange rates are quoted and what arbitrage means, let us look at the relationships that exist between spot rates quoted in various currencies. The forces that support these linkages are again arbitrage and shopping around. For our purposes, we can ignore the many market makers: when we talk about bid and ask, we now mean the *market quote*, that is, the best bid across all market makers, and the best ask. The new issue is how these market quotes in various currencies are linked.




Figure 3.7. (a) Triangular arbitrage (do I make money doing this?) and (b) triangular shopping around (which of the two gives me the best price?).

- Someone engaging in triangular *arbitrage* tries to make money by sequentially buying and selling various currencies, ending with the original currency. For instance, you could convert AUD into USD, and then immediately convert the USD into GBP and the GBP back into AUD, with the hope of ending up with more AUD than you started out with. The no-arbitrage condition says that you should not make a profit from such activities. Actually, when there are transaction costs or commissions, you are likely to end up with a loss. The potential loss is due to commissions, notably the bid-ask spread. Thus, in this context, arbitrage implies that the set of exchange rates quoted against various base currencies should be such that you cannot make any risk-free instantaneous profits after paying transaction costs.
- *Shopping around* is the search for the best way to achieve a desired conversion. For instance, an Australian investor who wants to buy GBP may buy directly, or may first convert AUD into USD and then convert these USD into GBP. Shopping around implies that the direct AUD/GBP market can survive only if its quotes are no worse than the implied rates from the indirect transaction.

In the case of perfect markets, the regular arbitrage and shopping-around arguments lead to the same conclusion. We illustrate this in the following example (and figure 3.7 shows the difference between arbitrage and shopping around).

Example 3.12. Suppose 1 GBP buys USD 1.5, while 1 USD buys AUD 1.6; therefore, if we directly convert GBP into AUD, 1 GBP should buy $1.5 \times 1.6 = 2.4$ AUD. With this AUD/GBP rate and assuming a zero spread,

- nobody can make a free-lunch profit by any sequence of transactions, and
- everyone is indifferent between direct conversions between two currencies and indirect, triangular transactions.

Below, we see what the implications of arbitrage and shopping around are when there are bid-ask spreads. In order to simplify matters, we shall first

show how to compute the implied rates from an indirect route. We shall call these implied rates the *synthetic* rates. Having identified these synthetic rates, we can then invoke the same mechanisms that enforce the law of one price as when we studied the relationship between the quotes made by various market makers.

3.3.3.1 Computing Synthetic Cross Rates

In general, a synthetic version of a contract is a combination of two or more other transactions that achieves the same objective as the original contract. That is, the combination of the two or more contracts *replicates* the outcome of the original contract. We shall use the notion of replication repeatedly in this textbook. For now, consider a simple spot transaction: a Japanese investor wants to convert JPY into GBP.

- The investor can use the direct market and buy GBP against JPY. We will call this the original contract.
- Alternatively, the investor can first buy USD with JPY, and then immediately exchange the USD for GBP. This is a combination of two contracts. It replicates the original contract since, by combining the two transactions, the investor initially pays JPY, and ultimately ends up with GBP. Thus, this is a synthetic way of achieving the original transaction.

Note that the synthetic contract may be the more efficient way to deal, since the USD market has a lot of volume (or depth) in every country, and therefore has smaller spreads. (This is why the USD is involved in 90% of the trades.) Let us see how the synthetic JPY/GBP rates can be computed.

Example 3.13. What are the synthetic JPY/GBP rates, bid and ask, if the quotes are JPY/USD 101.07–101.20 and USD/GBP 1.3840–1.3850?

Step 1: multiply or divide? The dimension of the rate we are looking for is JPY/GBP. Because the dimensions of the two quotes given to us are USD/GBP and JPY/USD, the way to obtain the synthetic rate is to multiply the rates, as follows:

$$[\text{JPY/GBP}] = [\text{JPY/USD}] \times [\text{USD/GBP}]. \quad (3.4)$$

Note that on the right-hand side of the equation, the USD in the denominator of the first quote cancels out the USD in the numerator of the second quote, leaving us with the desired JPY/GBP number.

Step 2: bids or asks? The first quote is the natural quote for a Japanese agent, the second one takes the USD as the base. Consider the synthetic ask (relevant for buying GBP from a JPY position). Starting from JPY we buy USD, so we need the ask; and with the USD we buy GBP, so we again need ask. Thus,

$$\begin{aligned} \text{Synthetic } S_{t,\text{ask}}^{\text{JPY/GBP}} &= S_{t,\text{ask}}^{\text{JPY/USD}} \times S_{t,\text{ask}}^{\text{USD/GBP}} \\ &= 101.20 \times 1.3850 = 140.16. \end{aligned} \quad (3.5)$$

By a similar argument, we can obtain the rate at which we can synthetically convert GBP into USD and these into JPY:

$$\begin{aligned} \text{Synthetic } S_{t,\text{bid}}^{\text{JPY/GBP}} &= S_{t,\text{bid}}^{\text{JPY/USD}} \times S_{t,\text{bid}}^{\text{USD/GBP}} \\ &= 101.07 \times 1.3840 = 139.88. \end{aligned} \quad (3.6)$$

This example is the first instance of the law of the worst possible combination, or the rip-off rule. You already know that, for any single transaction, the bank gives you the worst rate from your point of view (this is how the bank makes money). It follows that if you make a sequence of transactions, you will inevitably get the worst possible cumulative outcome. This law of the worst possible combination is the first fundamental law of real-world capital markets. In our example, this law works as follows:

- Note that we are computing a product. The synthetic ask rate for the GBP (the higher rate, the one at which you buy) turns out to be the highest possible product of the two exchange rates: we multiply the two high rates, ask times ask. Note, finally, that if the purpose is to buy forex, then a high rate is also an unfavorable rate. In short, we buy at the worst rate, the highest possible combined rate.
- We see that, likewise, the synthetic bid rate for the GBP (the lower rate, the one at which you sell) turns out to be the lowest possible product of the two exchange rates: we multiply the two low rates, bid times bid. Note also that if the purpose is to sell forex, then a low rate is also an unfavorable rate. In short, we sell at the worst rate, the lowest possible combined rate.

Let us look at another example. The data are the same except that the British quotes are now direct and not indirect.

DIY Problem 3.1. The JPY/GBP synthetic bid and ask rates, if the quotes are JPY/USD 101.07–101.20 and GBP/USD 0.722 02–0.722 54, are

$$\left. \begin{aligned} \text{Synthetic } S_{t,\text{bid}}^{\text{JPY/GBP}} &= \frac{S_{t,\text{bid}}^{\text{JPY/USD}}}{S_{t,\text{ask}}^{\text{GBP/USD}}} = 139.88, \\ \text{Synthetic } S_{t,\text{ask}}^{\text{JPY/GBP}} &= \frac{S_{t,\text{ask}}^{\text{JPY/USD}}}{S_{t,\text{bid}}^{\text{GBP/USD}}} = 140.16. \end{aligned} \right\} \quad (3.7)$$

- Derive this solution from the previous one, invoking our earlier results on inverse rates, equations (3.2) and (3.3).
- Verify that you get the above answer also if you first think of the dimensions and then apply the law of the worst possible combination.

Figure 3.8 shows the spreadsheet set up by one particular trader to help him shop around.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V
2	EUR	USD	1,4868															EUR/NOK	USD/NOK			
3	EUR	JPY	157,9300	1,4869														8,0745	8,0765	5,4321	5,4304	Via EUR
4	EUR	GBP	0,7477	1,4868														8,0758	8,0768	5,4303	5,4323	Reuters
5	EUR	CHF	1,6050	1,4868														EUR/DKK	USD/DKK			
6	EUR	AUD	1,6673	1,6053														7,4527	7,4530	5,0122	5,0128	Via EUR
7	EUR	CAD	1,4846	1,4868														7,4533	7,4535	5,0123	5,0128	Reuters
8	EUR	CZK	26,062	1,4868														EUR/CAD	USD/CAD			
9	EUR	DKK	7,4527	1,4868														1,4846	1,4851	0,9985	0,9988	RD2002
10	EUR	EEK	15,645	1,4868														1,4846	1,4853	0,9986	1,6539	Reuters
11	EUR	HKD	11,5940	1,4868														EUR/AUD	AUD/USD			
12	EUR	HUF	259,41	1,4868														1,6676	1,6681	0,8914	0,8916	RD2002
13	EUR	IDR	13877	1,4868														1,6673	1,6683	0,8914	0,8916	Reuters
14	EUR	ISK	96,75	1,4868														EUR/INZ	NZD/USD			
15	EUR	LTL	3,4526	1,4868														1,8955	1,8963	0,7841	0,7844	RD2002
16	EUR	LVL	0,6973	1,4868														1,8950	1,8968	0,7837	0,7847	Reuters
17	EUR	MYR	4,8086	1,4868														EUR/SGD	USD/SGD			
18	EUR	NOK	8,0745	1,4868														2,1074	2,1077	1,4174	1,4175	RD2002
19	EUR	NZD	1,8950	1,4868														Via USD		1,4172	1,4177	Reuters
20	EUR	PHP	60,19	1,4868														koers in kleine pips		26,5		

Figure 3.8. A dealer's shopping-around spreadsheet. Courtesy of Paul Goossens, dealer at KBC Brussels. Paul's spreadsheet shows the best quotes from EBS's broking screens, from Reuters Dealing 2002, and the indirect quotes (via USD or EUR). The latter are obviously rounded. Check how the indirect quotes are always wider at one side at least. (With only two pips between the best direct quotes, and with rounding of the synthetic quotes, one side must always seem to match.) The wider quotes labeled Reuters are the indicative, nonbinding ones from the Reuters FX/FX pages; they mean nothing except that some banks are willing to quote. See how Paul's sheet gets the EUR/USD quote from EBS into the darker parts of the spreadsheet. Cell I1 is selected; spot the underlying command =RtGet("IDN": "EUR=EBS" : "BID") in the enter function box above the spreadsheet. From the imported data in the black part, synthetic rates are computed.




Figure 3.9. Triangular arbitrage and shopping around.

3.3.3.2 Triangular Arbitrage with Transactions Costs

Now that we understand synthetic quotes, we can derive bounds imposed by arbitrage and shopping around on quotes in the wholesale market. Just think of the direct quotes as the quotes from bank X, and think of the synthetic quotes as the quotes from bank Y.

- *Arbitrage* then says that the two bid-ask quotes should overlap by at least one point; otherwise, you can buy cheap in the direct market and sell at a profit in the synthetic market or vice versa.
- *Shopping around* implies that if a bank skews its quotes so as to be (very) attractive at (only) one side, then it will attract a lot of business very fast; thus, this skewing cannot be persistent. But when we talk about market quotes (the best bid, and the best ask, across all market makers) rather than the quotes by an individual dealer, the force is even stronger. Individually, a market maker may very well want to make one of its quotes unappealing for some time, as we saw. But if there are many market makers it would be quite unlikely that, across all market makers, even the best direct quote would still be unappealing against the synthetic one, for that would mean that among all the competing market makers there is not a single one that is interested in that particular type of deal. Thus, instances where a direct quote is dominated by a synthetic one at one side should be rare and short-lived, and the more so the higher the number of market makers.
- The above assumes that the direct market has enough volume. Indeed, with a very thin market, the spread required to make market making sustainable may be too wide to allow the direct market to compete on both sides with the synthetic market via a heavily traded vehicle currency (like the USD or the EUR). The volume and depth of the wholesale market for dollars relative to almost any other currency is so large (and the spreads, therefore, so small) that a substantial part of the nondollar transactions are, in fact, still executed by way of the dollar. Direct cross deals emerged as of the mid 1980s only, and are still confined to heavy-volume currency pairs.

As a final note, in the retail markets most customers have no direct access to cross rates, and bank clerks occasionally compute cross rates even where the actual transaction could be executed very differently. A Japanese bank, for instance, would post quotes for JPY/GBP and JPY/EUR rates for its retail customers, but typically not for GBP/EUR. Should a retail customer sell EUR and buy GBP, the clerk would actually compute the synthetic rates we have just derived, as if the customer first went from EUR to JPY and then to GBP, even if in the bank's trading room the actual conversion may be done directly from EUR into GBP. Unless you have an account with a Euroland or U.K. bank, or enough clout with your home bank, you would have little choice but to accept the large spread implied by such synthetic rates.

This finishes our tour of the workings of the exchange markets. We continue the chapter with some wise advice on the merits and shortcomings of using exchange rates to translate foreign amounts of money. This brings us to the twin concepts of "PPP" and "real" exchange rates, key issues for understanding the relevance of currency risk.

3.4 Translating FC Figures: Nominal Rates, PPP Rates, and Deviations from PPP

Obviously, when you exchange an FC amount into HC or vice versa, you will use the exchange rate relevant at the moment. But actual transactions like this are not the sole conceivable purpose for such a conversion; rather, the purpose may just be *translation*, that is, to have an idea what an FC amount means in a unit that you are more familiar with, the HC. For instance, if a resident of Vanuatu tells you she is making 1m vatus a month, most people would not have a clue whether they should be impressed or not. In a case like this we do not want to actually exchange any vatus into our own HC; we would simply like to translate an FC number into a unit that is more meaningful to us.

The most commonly used solution is to resort to the market exchange rate to make the translation. The result is an improvement on the FC amount in the sense that you know what you would be able to do with this converted amount if you consumed it here, at home. But your objective may be to have a feel for what the FC amount would mean to a resident of the foreign country, that is, if the money is consumed there, not here. Both questions—the purchasing power of some amount of money in your home country, and the purchasing power abroad—provide the same answer if prices abroad and at home are on average the same once they have been converted into the same currency. This situation is known as (absolute) purchasing power parity (APPP). As we will illustrate below, APPP does not hold in reality, with deviations becoming more important the more different the two countries are in terms of location or economic development.

3.4.1 The PPP Rate

To have a more reliable feel for what a given amount of foreign money really means locally, one needs for each country a number called the price level, which we denote by Π (at home, and in HC), and Π^* (abroad, and in FC), respectively. A price level is an absolute amount of currency—not an index number—needed to buy a standard consumption bundle. Computing price levels for different countries makes sense only if the consumption bundle whose cost is being measured is the same across countries. In a simple economy in which fast food is the only commodity, the bundle may be one soda, one burger, one fries (medium), a salad, and a coffee—let’s call this the BigMeal. We simply jot down the prices of the components abroad and at home, and tot them up into price levels for BigMeals.

Any differences in price levels, after conversion into a common currency, would make a simple conversion of an FC amount into the HC rather misleading if translated price levels are very different.

Example 3.14. You often chat with a friend living in the Republic of Freedonia, where, since the presidency of Groucho Marx, the currency is the Freedonian crown (FDK). Let $S_t = \text{USD}/\text{FDK} = 0.010$. You earn USD 50 per unit of time, your Freedonian friend 2,000 FDK. What does that income really mean if the standard consumption bundle, our BigMeal, costs USD 5 here and FDK 250 in Freedonia?

- At the spot rate of USD/FDK 0.010, your friend seems to earn only $2,000 \times 0.010 = \text{USD } 20$, suggesting that she is 60% worse off than you.
- But this ignores price differences. What you “really” earn is $50/5 = 10$ BigMeals, while your friend makes $2000/250 = 8$ BigMeals. That is, your friend is “really” almost as well off as you are.

3.4.1.1 What Is the PPP Rate?

To buy eight BigMeals at home, you would need $8 \times 5 = \text{USD } 40$. So one way to summarize the situation is that FDK 2,000 means as much to your friend abroad as USD 40 means here, to you. The USD 40 is called the translation of FDK 2,000 using the PPP rate rather than the nominal rate, and the implied PPP rate is the $40/2,000 = 0.020$ USD per FDK, the ratio of the two price levels.

Let us generalize. Suppose you want to have a feel for what an FC amount Y^* “really” means to a foreigner. The question can be made more precise as follows. Give me an HC amount \hat{Y} such that its purchasing power here, \hat{Y}/Π_t , equals the purchasing power abroad of the original amount, Y^*/Π_t^* :

$$\begin{aligned} \text{Find } \hat{Y} \text{ such that } \frac{\hat{Y}}{\Pi_t} &= \frac{Y^*}{\Pi_t^*} \\ \implies \hat{Y} &= \frac{\Pi_t}{\Pi_t^*} Y^* = \hat{S}_t^{\text{PPP}} Y^*, \quad \text{where } \hat{S}_t^{\text{PPP}} \stackrel{\text{def}}{=} \frac{\Pi_t}{\Pi_t^*}. \end{aligned} \quad (3.8)$$

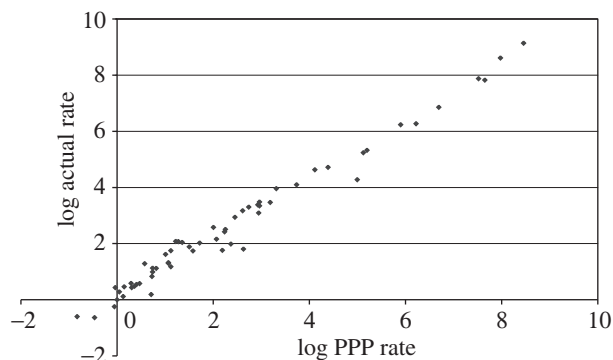


Figure 3.10. Log PPP versus log actual rates, HC/USD.
 Source: Based on data from the *Economist*, May 26, 2006.

So we can always compute the PPP rate as the ratio of the two price levels. For example, your friend's foreign amount (FDK 2,000) could have been translated at the PPP rate, $5/250 = 0.020$, which would have told you immediately that her income buys as much (in Freedonia) as USD 40 buys here.

Example 3.15. At the end of 2006, the CIA Factbook assessed Russia's 2005 GDP at 1.589 USD trillion using the PPP rate, and at 740.7 billion using the nominal official rate (www.cia.gov/cia/publications/factbook/geos/rs.html). What is the explanation for this? Are prices lower in Russia than in the United States, or is it the reverse?

For China, the figures were USD 8.883 trillion (PPP) and USD 2.225 trillion (official exchange rate), a ratio of about four to one instead of Russia's two to one. Which country, then, has the lower price level?

DIY Problem 3.2. Check that the PPP rate has dimension HC/FC.

The IMF and the World Bank, for instance, often use PPP rates rather than the regular ("nominal") rate to translate foreign GDPs or incomes or government budgets. Newspapers and magazines have also begun to adopt this approach. Lastly, the PPP rate also serves as a benchmark for the nominal rate. Many economists feel that, in the long run, the nominal rate for two similar economies should loosely fluctuate around the PPP rate, and never wander off very far above or below it. Let us see how well this theory fares, empirically.

3.4.1.2 PPP in Reality

In table 3.6 we take the *Economist's* favorite consumption bundle, the Big Mac, and we compute PPP rates for 59 countries—once in USD (a New Yorker should get 0.295 dollars to be as happy as a Beijinger with one extra yuan) and once in non-U.S. currency (a Beijinger should get 3.39 yuan to be as happy as a New Yorker with one extra dollar). You see that countries where the Big Mac has a high local price have, of course, low PPP rates but also tend to have low actual

Table 3.6. PPP rates based on Big Mac prices from the *Economist*, May 26, 2006.

	Currency	Local price	Actual value of \$	PPP rate of \$	Real rate of \$	Actual value in \$	PPP rate in \$	Real rate in \$
China	yuan	10.5	8.03	3.39	2.371	0.125	0.295	0.422
Macau	pacata	11.1	7.99	3.58	2.231	0.125	0.279	0.448
Malaysia	ringgit	5.5	3.63	1.77	2.046	0.275	0.564	0.489
Hong Kong	dollar	12	7.75	3.87	2.002	0.129	0.258	0.499
Indonesia	rupia	14600	9325	4709.68	1.980	0.000	0.000	0.505
Philippines	peso	85	52.6	27.42	1.918	0.019	0.036	0.521
Paraguay	guarani	9000	5505	2903.23	1.896	0.000	0.000	0.527
Egypt	pound	9.5	5.77	3.06	1.883	0.173	0.326	0.531
Ukraine	hryvna	8.5	5.05	2.74	1.842	0.198	0.365	0.543
Moldava	leu	23	13.2	7.42	1.779	0.076	0.135	0.562
Uruguay	peso	42.3	23.9	13.65	1.752	0.042	0.073	0.571
Russia	ruble	48	27.1	15.48	1.750	0.037	0.065	0.571
Dominican Rep	peso	60	32.6	19.35	1.684	0.031	0.052	0.594
Sri Lanka	rupee	190	103	61.29	1.681	0.010	0.016	0.595
Honduras	lempira	35.95	18.9	11.60	1.630	0.053	0.086	0.614
Bulgaria	lev	2.99	1.54	0.96	1.597	0.649	1.037	0.626
Slovakia	koruna	58	29.5	18.71	1.577	0.034	0.053	0.634
Poland	zloty	6.5	3.1	2.10	1.478	0.323	0.477	0.676
Thailand	baht	60	28.4	19.35	1.467	0.035	0.052	0.682
South Africa	rand	13.95	6.6	4.50	1.467	0.152	0.222	0.682
Pakistan	rupee	130	60.1	41.94	1.433	0.017	0.024	0.698
Venezuela	bolivar	5701	2630	1839.03	1.430	0.000	0.001	0.699
Costa Rica	colon	1130	510	364.52	1.399	0.002	0.003	0.715
Japan	yen	250	112	80.65	1.389	0.009	0.012	0.720
Singapore	dollar	3.6	1.59	1.16	1.369	0.629	0.861	0.730
Guatemala	quetzal	17.25	7.59	5.56	1.364	0.132	0.180	0.733
Argentina	peso	7	3.06	2.26	1.355	0.327	0.443	0.738
Georgia	lari	4.15	1.8	1.34	1.345	0.556	0.747	0.744
Taiwan	dollar	75	32.1	24.19	1.327	0.031	0.041	0.754
Estonia	kroon	29.5	12.3	9.52	1.293	0.081	0.105	0.774
Saudi Arabia	riyal	9	3.75	2.90	1.292	0.267	0.344	0.774
Lithuania	litas	6.5	2.69	2.10	1.283	0.372	0.477	0.779
Australia	dollar	3.25	1.33	1.05	1.269	0.752	0.954	0.788
UAE	dirham	9	3.67	2.90	1.264	0.272	0.344	0.791
Latvia	lats	1.35	0.55	0.44	1.263	1.818	2.296	0.792
Mexico	peso	29	11.3	9.35	1.208	0.088	0.107	0.828
Colombia	peso	6500	2504	2096.77	1.194	0.000	0.000	0.837
Croatia	kuna	15	5.72	4.84	1.182	0.175	0.207	0.846
South Korea	won	2500	952	806.45	1.180	0.001	0.001	0.847

exchange rates. Figure 3.10 shows this graphically. To “shrink” the outliers and give the smaller numbers more space, we plot the log of the actual against the log of the PPP rate. (This explains why there are negative rates: numbers below unity produce negative logs.) There is obviously a very strong link.

DIY Problem 3.3. Knowing that the Big Mac costs 3.10 in the United States and 155 in Freedonia, and that the spot rate is 100 crowns per dollar, complete Freedonia’s PPP rates in the table:

	Currency	Local price	Actual value of \$	PPP rate of \$	Real rate of \$	Actual value in \$	PPP rate in \$	Real rate in \$
Freedonia	korona	155	100					

Table 3.6. Continued.

	Currency	Local price	Actual value of \$	PPP rate of \$	Real rate of \$	Actual value in \$	PPP rate in \$	Real rate in \$
Czech Rep	koruna	59.05	22.1	19.05	1.160	0.045	0.052	0.862
Fiji	dollar	4.65	1.73	1.50	1.153	0.578	0.667	0.867
Hungary	forint	560	206	180.65	1.140	0.005	0.006	0.877
Turkey	lire	4.2	1.54	1.35	1.137	0.649	0.738	0.880
New Zealand	dollar	4.45	1.62	1.44	1.129	0.617	0.697	0.886
Slovenia	tolar	520	189	167.74	1.127	0.005	0.006	0.888
Aruba	florin	4.95	1.79	1.60	1.121	0.559	0.626	0.892
Brazil	real	6.4	2.3	2.06	1.114	0.435	0.484	0.898
Morocco	dirham	24.5	8.71	7.90	1.102	0.115	0.127	0.907
Peru	new sol	9.5	3.26	3.06	1.064	0.307	0.326	0.940
Chile	peso	1560	530	503.23	1.053	0.002	0.002	0.949
United states	dollar	3.1	1	1.00	1.000	1.000	1.000	1.000
Canada	dollar	3.52	1.12	1.14	0.986	0.893	0.881	1.014
Britain	pound	1.94	0.532	0.63	0.850	1.880	1.598	1.176
Euroland	euro	2.94	0.781	0.95	0.824	1.280	1.054	1.214
Sweden	krona	33	7.28	10.65	0.684	0.137	0.094	1.462
Denmark	kroner	27.75	5.82	8.95	0.650	0.172	0.112	1.538
Switzerland	franc	6.3	1.21	2.03	0.595	0.826	0.492	1.680
Iceland	kronur	459	72	148.06	0.486	0.014	0.007	2.056
Norway	kroner	43	6.1	13.87	0.440	0.164	0.072	2.274




Figure 3.11. PPP versus actual rates, HC/USD.
Source: Based on data from the *Economist*, May 26, 2006.

But a closer look at the table reveals big relative deviations, which are hard to spot from a log graph dominated by outliers. Kicking out the twenty highest cases so as to be able to forgo logs, this time, we get figure 3.11. Note how the observations tend to be above the equality line (where actual = PPP): the dollar tends to be too expensive, by Big Mac PPP standards. Yet there are also important deviations below the 45° line, where the slope of the ray through the dot is even below 0.5 in one case. The slope of this ray is called the *real exchange rate*, to which we now turn.

3.4.2 Commodity Price Parity

A concept used in textbooks is commodity price parity (CPP). It is said to hold when translated prices for an individual good are equalized across two countries:

$$\text{CPP holds if } P_{j,t} = S_t \times P_{j,t}^*, \quad (3.9)$$

with j referring to an individual good, and P_j (P_j^*) referring to its price at home, in HC (abroad, in FC). In fact, all the Big Mac evidence shown thus far is about CPP rather than PPP, a distinction that the *Economist* tends to gloss over.

CPP would hold if trading were costless and instantaneous. Obviously, in reality it does not work across the board; for commodities it is not too bad an approximation (within the bounds created by transportation costs and the like), but for consumer goods it is essentially a joke.

PPP in the true sense—i.e., for a bundle of goods—would clearly hold if CPP held for every individual good, or if deviations from CPP washed out after averaging across many goods. As we have seen, this is not really the case; apparently, too many deviations from CPP turn out to be in the same direction, suggesting there is a common force behind them. Forget CPP.

3.4.3 The Real Exchange Rate and (Deviations from) Absolute PPP

The real exchange rate (RER) is a measure of how far the nominal rate differs from the PPP rate: it is simply the nominal exchange rate divided by the PPP counterpart.

Example 3.16. In our Freedomian story, the nominal rate was 0.010 USD/FDK while the PPP rate was 0.020 USD/FDK; thus, the real rate was 0.5, which is a large deviation from unity, but not uncommon between two very different economies.

The real rate is a dimensionless number: [HC/FC] divided by [HC/FC]. In a way, it simply tells us what the ratio is of the *translated* price levels:

$$\text{RER}_t \stackrel{\text{def}}{=} \frac{S_t}{\hat{S}_t^{\text{PPP}}} \quad (3.10)$$

$$= \frac{S_t \times \Pi_t^*}{\Pi_t}, \quad \text{from (3.8)}. \quad (3.11)$$

Again, in the example one can find the RER for the FDK against the USD by translating into USD the foreign price of the BigMeal, $\text{FDK } 250 \times 0.010 = \text{USD } 2.5$, and divide it by the domestic price level, 5, which gets us $2.5/5 = 0.5$. Thus, the RER rate tells you how much cheaper (if $\text{RER} < 1$) or more expensive (if $\text{RER} > 1$) the foreign country is. A country with a below-unity real rate would be a nice place to spend your domestic income, or could be an attractive base to export from, but may not be the best place to export to. These are very different questions than the one answered by the PPP rate.




Figure 3.12. Relative prices of the Big Mac across the world, based on data from the *Economist*, May 26, 2006.

Obviously, if the real rate equals unity, both countries have the same price level. If that is true, Absolute PPP is said to hold:

$$\text{Absolute PPP holds if } RER_t = 1 \Leftrightarrow S_t = S_t^{\text{PPP}} \Leftrightarrow S_t \times \Pi_t^* = \Pi_t. \quad (3.12)$$

In figure 3.12, and in table 3.6 the countries have been ranked on the basis of the real rate. Two observations stand out. First, there is a five-to-one ratio between the most and least expensive countries, Norway and China. So deviations from PPP are big. Second, there is a system to it, to some extent: undervalued currencies tend to be developing ones, and overvalued ones developed. (The fact that thus USD is not top is anomalous, in this view. The long-lasting deficit in its current account may be one reason.) The (imperfect but strong) relation between real rate and degree of economic development is discussed in chapter 10.

DIY Problem 3.4. Norway is most expensive. In figure 3.11, identify the dot that corresponds to Norway.

3.4.4 The Change in the Real Rate and (Deviations from) Relative PPP

For most of the time since the 1980s, Japan has had a real rate above unity: it was a more expensive place to spend a dollar than the United States or Europe. Sometimes one is interested in whether a country's situation has worsened or improved. That is, has the real rate increased or decreased (as distinct from the issue of whether its level is above unity or not)?

To measure this, one can simply compute the RER's percentage change. Not surprisingly, the percentage change in the RER is determined by the percentage changes in the spot rate and the price levels—the inflation rates.

Example 3.17. Q. Suppose that five years ago the FDK traded at USD/FDK 0.012, and the price levels were USD 4 in the United States and FDK 250 abroad. (So, with current price levels being 5 and 250, respectively, inflation was 25% in the United States and zero in Freedonia.) Recalling that the current RER is $250 \times 0.010/5 = 0.5$, how much has the RER changed since then?

A. The old RER was $250 \times 0.012/4 = 0.75$; so the rate changed by $(0.50 - 0.75)/0.75 = -0.33$, that is, -33% . There was real depreciation of the crown—that is, Freedonia became cheaper over time—because the FDK went down *and* because inflation in Freedonia was lower than in the United States.

Below, we first show the general relation between the percentage change in the RER and the changes in the nominal rate, and then a first-order approximation that is occasionally used:

$$\text{Percentage change in the RER} = (1 + s_{t_0,t}) \frac{1 + \text{infl}_{t_0,t}^*}{1 + \text{infl}_{t_0,t}} - 1 \quad (3.13)$$

$$\approx s_{t_0,t} + [\text{infl}_{t_0,t}^* - \text{infl}_{t_0,t}], \quad (3.14)$$

where $s_{t_0,t}$ is the simple percentage change in the spot rate S between times t_0 and t while $\text{infl}_{t_0,t}$ and $\text{infl}_{t_0,t}^*$ denote inflation at home and abroad, respectively, over the same time window. The first-order approximation works well if both inflation rates are low. This is not the case in our Freedonian example.

Example 3.18. In our above story, foreign inflation was zero, U.S. inflation 25%, and the exchange-rate changed by minus one sixth; so the RER changes by

$$\left(1 - \frac{1}{6}\right) \frac{1 + 0.00}{1 + 0.25} - 1 = 0.66667 - 1 = -\frac{1}{3},$$

as computed directly before. In contrast, the first-order approximation would have predicted a change of $-\frac{1}{6} - 0.25 = -41.67\%$ rather than -33.33% . The error is nontrivial because, in this example, two numbers—the U.S. inflation rate and the change in the exchange rate—are far from zero.

If the RER is constant, whatever the level, then relative PPP (RPPP) is said to hold; and the percentage change in the RER is a standard measure of deviations from RPPP. An RPPP deviation is most often resorted to if the RER itself cannot be computed because price-level data are missing. If, indeed, absolute price levels for identical bundles are not available, there is no way of computing which of the two countries is the cheaper. But one can still have an idea whether the RER went up or down if one estimates the inflation rates from the standard consumption price indices (CPIs) rather than the price levels. A CPI is a relative number vis-à-vis a base period, and the consumption bundle is




Figure 3.14. RPPP versus actual rates against USD, 1965 = 1.00. (a) JPY RPPP versus actual; (b) SAR RPPP versus actual. *Source:* Underlying data are from DataStream.

For this reason, deviations from RPPP are important. Are they large? Figure 3.13 shows time-series data, taking January 1965 as the base period, on relative real rates against USD, for the DEM-EUR, JPY, GBP, SAR, and THB. We note four facts.

- (i) First, there are huge swings in the medium run, with the real rate appreciating by 50% and then going back—and occasionally even doubling or halving—in a matter of years, not decades. Imagine being caught in this as an exporter.
- (ii) Second, in the short run there is lots of inertia: once the rate is above its mean, it tends to stay there for years. Statistical analysis shows that the average half-life is three to five years, meaning that it takes three to five years, on average, for a deviation to shrink to half its original size. Thus, when you get into a bad patch, you can expect that this will be a matter of years rather than weeks or months.

