Christiano
362, Winter 2002
January 16

Lecture #4: More on Exchange Rates

1. Review. Last time, we derived:

$$\text{nominal return} = \frac{D + P'}{P} = 1 + R,$$

where D is the dollar payment you get from holding the asset, and R is the (net) nominal return.

Also:

$$\text{real return} = \frac{(D + P')/P_c'}{P/P_c} = (1 + R)\frac{P_c}{P_c'} = \frac{1 + R}{1 + \pi} \approx 1 + R - \pi.$$

Finally, the US dollar return, R, on a European asset that generates a euro rate of return $R_{\mathfrak{E}}$, is:

$$1 + R = \frac{\left(D^{\mathfrak{E}} + P^{\mathfrak{E}\prime}\right) \times E'}{P^{\mathfrak{E}} \times E} = (1 + R_{\mathfrak{E}})\frac{E'}{E} \approx 1 + R_{\mathfrak{E}} + \frac{E' - E}{E},$$

where $\mathfrak{E}$ is the notation we'll adopt to signify the euro, the new European currency. Cancelling out the 1's from each side:

$$R = R_{\mathfrak{E}} + \frac{E' - E}{E}.$$

This says that the return, in domestic currency, on a foreign asset is the foreign denominated return on that asset, plus what you make from holding foreign currency for a while. Thus, if the foreign rate of interest is 5% and the depreciation of the domestic currency is 10%, then the return, in domestic currency units, of the foreign asset is 15%. This is actually only an approximation. But, it is a very good one. To see this, note that, from the exact formula:

$$R = (1 + R_{\mathfrak{E}})\frac{E'}{E} - 1 = 1.05 \times 1.1 - 1 = 1.155 - 1 = 0.155,$$

which is pretty close to 0.15.

When you're buying a foreign, versus a domestic asset, you're really doing two things. First, you are investing in the asset itself. That is, you are earning 5% on the foreign asset. Second, you are benefiting from any depreciation that might occur in your exchange rate. For example, suppose $E' = 1.10$ and $E = 1.00$, so that the currency depreciates by 10 percent. Then, if you buy one euro with one dollar, and then turn around later and sell that euro for a dollar again, you get E' dollars back. The gross return on this transaction is the dollars you get back, E', divided by the dollars you put in, E, i.e., $E'/E = 1.10$. The net rate of return is $E'/E - 1 = 0.10$. The round trip through the foreign exchange market earns you 10%.

The above formula says that if the domestic currency depreciates, this adds to the return earned by domestic residents on financial assets in that country. Similarly, if the domestic currency depreciates, then this subtracts from the return earned by domestic residents.

2. Uncovered interest parity. Let the nominal rate of return on a US asset be denoted by $R_{\$}$. In practice, when traders decide how many dollar and euro assets to hold, they may know the values of $R_{\$}$ and $R_{\mathfrak{E}}$ (at least as long as it's bonds we're talking about, and not equity), but E' is not known. Instead, they must form an expectation of what $E^{'}$ is likely to be. We denote this by E^e.

   The markets for foreign exchange are extremely active. Vast amounts of currency changes hands each day. Some people are buying and selling because they're hoping to profit from changes in exchange rates. Others participate because they have receipts in one currency, but pay out profits in another currency. With many people trading financial assets, we would expect assets in different countries to generate similar rates of return, when denominated in the same currency units. Thus, we expect this relationship between the nominal return on US dollar assets, $R_{\$}$, and the nominal return on European assets, $R_{\mathfrak{E}}$:

   $$R_{\$} = R_{\mathfrak{E}} + \frac{E^e - E}{E}.$$

   In words, this expression says that if US interest rates are higher than European interest rates ($R_{\$} > R_{\mathfrak{E}}$), then the US dollar must be expected to depreciate. Similarly, if US interest rates are lower than European interest rates, then the US dollar must be expected to appreciate.

   How might markets produce this equality, in practice? Suppose US interest rates were higher than the dollar returns one could earn on European assets, i.e., suppose

   $$R_{\$} > R_{\mathfrak{E}} + \frac{E^e - E}{E}.$$

What would happen? Presumably, traders in Europe would sell euros and buy US dollars, to take advantage of high US returns. There would be a rush for the foreign exchange market, as traders scrambled to sell euros and acquire dollars. But, this would have the effect of causing the US dollar to appreciate, i.e., of driving E down. The lower E, assuming E^e does not change, implies a higher anticipated depreciation of the currency, restoring equality to the above relation. A similar story explains how the markets would prevent US rates from being lower than foreign rates, when denominated in dollars.

The above equality is called the Uncovered Interest Parity (UIP) relation.

3. UIP, Risk and Liquidity. The rate of return, $R = R_{€} + (E^e - E)/E$, is the 'expected' rate of return, denominated in domestic currency units, of a foreign financial asset whose foreign currency return is $R_{€}$. When we compare R with $R_{\$}$, we are implicitly assuming that traders don't care about the risk characteristics of an asset. They only care about their expected returns. This leads to the UIP relationship, which says that expected returns on different assets should be the same. In practice, of course, traders do care about risk. An asset that is very risky may have to have a higher expected rate of return than an asset that has no risk, for traders to be willing to hold both of them.

   There is another reason why focusing exclusively on expected returns oversimplifies things. Different assets have different liquidity characteristics. A highly liquid asset is one for which it is easier to find a buyer in case you need to sell it. For example, US government debt is highly liquid. The market for that is so highly developed and there are so many people in it all the time, that US government debt is as easy to dump in case you have to, as it is to dump regular currency. The IOU I gave to my colleague yesterday in exchange for lunch money is completely illiquid. So, the UIP relationship also implicitly abstracts from the different liquidity characteristics of different assets.

   The upshot is that there is no reason for UIP to hold exactly. At best, it can only be expected to hold only as an approximation. Consistent with UIP, we do find that countries with low interest rates generally have an appreciating currency, at least over long periods of time. An example is Japan, whose currency has appreciated on average relative to the dollar and whose interest rates are lower than US interest rates on average. UIP tends not to hold over shorter periods of time.[1]

[1] For a recent paper that documents this, see 'Long Horizon Uncovered Interest Rate Parity,' by Guy Meredith and Menzie Chinn. This is a November 1998 working paper available as NBER working paper 6797 at http://www.nber.org/papers/w6797. This paper uses simple econometric techniques. It is *not* required reading for the class.

It is interesting to think a little more about the ways in which risk considerations creep into comparisons of different assets. Consider, for example, US and German government debt. Both are very liquid. Both are essentially risk free, when denominated in their own currencies. Still, the risk characteristics of the two types of debt, when the returns are denominated in common units, *are* different. To see this, suppose $R_{\$}$ and $R_{€}$ are the return on US and German government debt, respectively. If the assets are held until maturity, then their risk when the returns are *denominated in their own currency*, is roughly zero. At the time you buy US government debt that you plan to hold onto until maturity, $R_{\$}$ is known for sure. The same is true for $R_{€}$. However, what is not known at the time German government debt is purchased is its *dollar denominated* return. This is not known because the value of the exchange rate when the debt matures is not known. This is what makes German debt riskier than US debt, to an American. Of course, the opposite is true from the perspective of a German. A German will find US government debt riskier than German government debt because the former involves exchange risk, while the latter does not.

As noted above, these issues of risk are likely in practice to prevent the UIP from holding exactly. Consider the following example. Suppose $E^e/E = 1.05$, so that a 5 percent dollar depreciation is expected. Suppose that $R_{\$} = R_{€} = 0.05$, i.e., the interest rates in both countries is five percent. Now, the uncovered interest parity relationship says that no one should be holding American denominated assets. Any American holding US government debt should sell it and buy German government debt. Why might an American hold on to US government debt anyway? The American may well agree with the assessment that E^e/E is 1.05. And, if E' turned out to equal E^e, the American would regret not having sold his or her US government debt and bought German government debt instead. But, the fact is that E' is uncertain, and therefore it could end up higher than E^e, or lower. The American may be especially concerned about the latter. For example, he or she may be worried that E'/E might turn out to be 0.95, say, where the US dollar appreciates by 5 percent. In this case, the American would *lose* money by holding German government debt.

Often, people hold government debt without planning to hold it to maturity. For example, you may want to buy 30-year government debt and only plan to hold onto it for one year. Then, even the own currency return on government debt is risky. The value of government debt that is sold before it matures is determined in the market, and is a random variable when you first buy it. Thus, there are at least two sources of risk to consider in comparing rates of return across countries: one that stems from uncertainty in the local currency denominated return and the other that stems from uncertainty in the exchange rate. Both of these forms of risk, if important enough, could lead to UIP not holding.

In particular, if there were evidence that UIP did not hold in the data, that would not constitute evidence of irrationality on the part of portfolio managers. That's because, in deriving UIP, we abstracted from risk and liquidity considerations. As it happens, UIP tends *not* to fit the data very well when we consider assets with short-term maturities. It does better on assets with longer term maturities. In developing our theory of exchange rates, we will make heavy use of UIP. This is because it probably does a good job in capturing the primary channel linking changes in interest rates and expected future exchange rates to the current exchange rate. A complete understanding of exchange rates requires also knowing how interest rate and expected exchange rate changes impact on the current exchange rate via their impact on risk. This channel is less well understood, and, in any case, well beyond the scope of this course.

4. Covered Interest Parity. This is a relationship which does hold in the data. Let F denote the exchange rate in the forward market. This is known for sure at the time you buy German or US government debt. The return on German denominated debt, denominated in dollars, assuming the forward market is used, is

$$\frac{F(1+R_{\mathfrak{C}})}{E} \approx 1 + R_{\mathfrak{C}} + \frac{F-E}{E}.$$

The covered interest parity relation implies:

$$R_{\$} = R_{\mathfrak{C}} + \frac{F-E}{E}.$$

If this did not hold, then sure profits could be made simply by selling one of the assets and buying the other. In efficient markets, sure profits, or arbitrage opportunities, don't exist. Or, if they do they are quickly exploited until they disappear.

5. Exchange rate determination in the Short Run. The interest parity condition gives us a way to think about how the exchange rate is determined in the short run. Suppose $R_{\mathfrak{C}}$ and E^e are just given for now. Suppose the US monetary authorities cut the US interest rate, $R_{\$}$. What will happen to the current exchange rate, E? Suppose we start in a situation where covered interest rate parity holds. With the fall in $R_{\$}$, but *before* any change in E (I'm holding E^e and $R_{\mathfrak{C}}$ constant from beginning to end of this experiment), German assets will look much more attractive than American assets to everyone. So, people will sell US dollars and buy Marks to take advantage of the higher rates there.

This process will drive down the value of a dollar, sending E down and therefore, driving $(E^e - E)/E$ down. That is, the depreciation of the US dollar will (given that E^e is being held constant) create an anticipated appreciation of the dollar. This will happen up to the point where covered interest parity holds again. Thus, anticipated dollar appreciation will make up for the now relatively low nominal return on US assets.

6. Money demand and money supply. The book explains quite nicely, the following money demand relation:

$$\left(\frac{M}{P}\right)^{demand} = L(R, Y),$$

where L is decreasing in R and increasing in Y. In practice, this expression is sometimes assumed to have the following special form: $L(R, Y) = f(R)Y$, where f is a decreasing function of R. With this specification, the percent increase in the demand for real money balances resulting from a one percent increase in income, Y, (the *income elasticity of money demand*) is unity (i.e., one). We can test this view by looking at data on the velocity of money:

$$\text{Money velocity} = \frac{PY}{M}.$$

According to the money demand relation which imposes unit income elasticity, velocity should have the following relationship to the interest rate:

$$\text{Money velocity} = \frac{1}{f(R)}.$$

That is, as income changes, money velocity should not change, and velocity should move up and down in the same direction as the rate of interest.

To see what velocity actually does, look at the attached figure. The relatively smooth line is velocity (left scale) and the choppier line is the rate of interest (right scale).

There are several things worth noting in the figure. First, consider the velocity - interest relationship. At the low frequency level, they move together. Broadly, velocity moves up until 1980, whereupon it turns around and comes down again. The interest rate follows the same broad pattern. At a higher frequency, the relationship seems to change. In the first half of the sample, velocity does not respond much to the higher frequency movements in the interest rate, and in the second half it does. Second, consider the velocity - income relationship. Note that interest rates in the end of the sample are nearly where they were in the beginning. Yet, velocity is not back to where it was before. Instead, velocity seems to be somewhat higher. That is, as PY has gone up, M has not quite kept up. This suggests that the income elasticity of demand for money is a little less than unity. That is, a one percent jump in Y induces less than a one percent rise in M^{demand}. Another possibility is that all the technical and legal innovations that have occurred in the past decades (spurred in part by the high interest rates of the 70s and early 80s) have allowed people to economize on cash balances. Now that they are in place (ATM machines, information technology that makes credit card purchases easy, etc.), they will not be reversed and we can expect velocity to stay up for a while.

We can actually use the data in the figure to 'estimate' the money demand equation. Let's posit the following money demand equation:

$$\frac{M}{P} = f(R) \times Y^{\gamma},$$

where γ is a parameter, whose value we will estimate. In the previous lecture we talked about the version of this equation that is commonly used, the one in which $\gamma = 1$. The parameter, γ, is the elasticity of demand for real balances with respect to an increase in income, holding R fixed. This statement reflects two things. First, the elasticity of demand for M/P with respect to Y is defined as the percent increase in M/P demanded, when Y rises by one percent. Second, with the above equation, the percent increase in M/P with a one percent increase in Y is approximately γ.

We can estimate γ in the following way. The attached figure indicates that velocity now is around 2.3, and it was around 2 in 1967. Thus, it increased by 15 percent. At the same time, output (after inflation) has increased 130 percent over the same period. How can we use this information to estimate γ?

Recall the definition of V, velocity. It is $V = Y/(M/P)$. Rewriting the above equation, we find,

$$V = \frac{1}{f(R)} \times Y^{1-\gamma}.$$

Approximately,

$$\hat{V} = (1 - \gamma)\hat{Y},$$

where the hat over a variable means 'percent change'. Plugging in the numbers from above, we get that $1 - \gamma$ is 15/130, or that γ is 0.88.

Later, we'll find that this number is useful for figuring out what money growth rate will hit a given target inflation rate.

7. The Short Run.

    (a) Combine UIP and the model of the money market, and assume E^e, Y, P are fixed.

        Rationale for fixed P assumption:

        i. a lot of prices are fixed by contract. In addition, a lot of costs (like wages), which go into determining prices, are fixed by contract too.
        ii. prices move very little from one month to the next, compared to exchange rates (see Fig 14-11 in KO).

        Rationale for fixed Y assumption: increasing production requires a lot of advanced planning and takes time.

    (b) Experiments: increase in US money supply drives down $R_{\$}$ and results in currency depreciation, E goes up; increase in German money supply drives down $R_{€}$ and results in (US) currency appreciation, E goes down.

Money velocity (MZM from St Louis Fed) and Opportunity Cost