Christiano 362, Winter 2003 January 15

Lecture #4: Towards a Simple Model of Exchange Rates and Interest Rates

- 1. The model we will construct today has two 'legs': one it the Uncovered Interest Parity (UIP) theory described last time. The other is a model of the money market (MM). We start with a restatement of the UIP, followed by a discussion of how well it fits the data. Similarly a statement of MM is presented, and then it is evaluated in light of the data. In each case, the model seems to capture long-run patterns in the data. The UIP and MM models do less well on the shorter run movements. The forces summarized by those models seem to be the dominant ones in the long run, but there appear to be more things going on in the short run. For us, this is fine. This is a course about the basics, and we abstract from the rest.
- 2. A basic assumption about the UIP is that in comparing different financial assets, they only look at expected returns. We abstract from the real-world fact that traders also worry about liquidity and risk. For example, under the UIP a trader is indifferent between an asset which generates a return of 5% with probability 1 and an another asset which generates a return of 7% and 3% with probability 1/2 each. Later in this course we will depart a little from UIP by bringing risk into the analysis. However, it turns out that we can go a long way if we abstract from it.

The return, in dollar terms, on a foreign asset whose return in foreign nominal terms is $R_{\mathfrak{E}}$, is:

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

where E^e is the expected value of the exchange rate in the future. If the nominal return on the domestic asset is R, then efficient markets, together with the assumption that traders only care about the expected return on an asset, implies:

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

The assumption of efficient markets is what guarantees that this equation (we'll call it 'the UIP equation') holds. The notion is that if it did not hold, attempts by people to shift from one currency to another would move E until equality was restored.

To see how the UIP works, consider the following example. Suppose $R_{\mathfrak{E}}$ and E^{e} are just given and fixed for now. Suppose that in some particular point in time 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 UIP holds. With the fall in $R_{\$}$, but before any change in E (remember, we're holding E^e and $R_{\mathfrak{E}}$ constant from beginning to end of this experiment), European assets will look much more attractive than American assets to everyone (at least, to people who care only about expected return). So, people will attempt to sell US dollars and buy Euros to take advantage of the higher rates there. This process will drive down the value of a dollar, sending E up and, therefore, $(E^e - E)/E$ down. That is, the instantaneous 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 UIP holds again. In this way, additional anticipated dollar appreciation compensates people for the now relatively low nominal return on US assets.

3. Evaluating the UIP. At one level, this is tough to do. To evaluate UIP we'd like to compare $R - R_{\mathfrak{E}}$ with $(E^e - E)/E$, to see if the the UIP's prediction that these two terms are equal to each other is true in the data. But, we don't observe E^e , which is people's psychological expectation of what value the exchange rate will take on in the future. What we can do is compare $R - R_{\mathfrak{E}}$ with (E' - E)/E, where E' the actual value of the exchange rate in the future. In any one period, of course, E' and E^e will not be the same. However, over time, they should be similar on average. For example, suppose that from the point of view of any period, next period's exchange rate is the result of a coin toss were if the result is heads, the exchange rate is 1.5 and if it's tails the exchange rate is 0.5. In this case, in each period it would be reasonable to think that $E^e = 1$, which is the expected value of next period's exchange rate. Now, in this case, it can never happen that $E^e = E'$. But, if we looked at a series of observations on E^e over time and a series of observations of E', we expect that the average of the two would be the same (the average of E^e would of course be unity, and the average of E' would be roughly unity, depending on how many periods you average over.)

Consider the first figure below. It is a scatter plot of $R - R_{\mathfrak{E}}$ (horizontal axis) versus (EI - E)/E (vertical axis), where $R_{\mathfrak{E}}$ is a nominal German Mark denominated interest rate, R is a US dollar interest rate, and E is dollars per West German Mark. Each dot represents an observation on $R - R_{\mathfrak{E}}$ and (EI - E)/E, for a different date. The straight line is a regression line computed through the data. It is the straight line with the property that the vertical distance to each dot is as small as

possible. Note that the line is positively sloped, with slope roughly unity. This indicates that 'on average' if $R - R_{\mathfrak{E}}$ goes up by one percentage point, then (EI - E)/E goes up by one percentage point too. Put differently, when (EI - E)/E goes up by one percentage point, it is just as likely that $R - R_{\mathfrak{E}}$ goes up as down. This is consistent with the UIP.¹ Note that the actual data in many cases like quite far from the line. This is consistent with UIP because there is no limit to how big forecast errors can be.

The results in the figure are actually only consistent with UIP holding in the 'long run'. This is because they obtained using 5 year US and German interest rates, and (EI - E)/E is computed over the same 5 year period. When the same calculations are done using shorter term rates and exchange rate changes over correspondingly shorter horizons, then the results are much less supportive of the UIP. In this case, the regression line actually has a *negative* slope.² This is one sense in which the UIP seem to apply to the longer run, while it applies less well to the short run.

Figure 2 displays a graph of $R - R_{\mathfrak{E}}$ (bottom panel), where R corresponds to the nominal interest rate in the US and $R_{\mathfrak{E}}$ corresponds to the nominal, Canadian dollar denominated interest rate in Canada. Note how the Canadian interest rate is typically higher than the US interest rate. Consistent with UIP, the US dollar has appreciated on average throughout this period. Figures 3-6 report the analogous data for 4 other countries. In each case, the exchange rate is measured as dollars

$$R_t - R_{\mathfrak{E},t} = \frac{E'_t - E_t}{E_t} + \varepsilon_t,$$

where $E'_t = E_{t+\tau}$ if the interest rate is a τ period interest rate. It is reasonable to suppose that ε_t is uncorrelated with $(E'_t - E_t)/E_t$, given that it is a forecast error (this assumption about trader's forecast or expectation of the future value of the exchange rate is called 'rational expectations'). Thus, the UIP predicts that if you run a regression of $R_t - R_{\mathfrak{E},t}$ on $(E'_t - E_t)/E_t$, the coefficient on the right hand variable should be unity. Interestingly, this is what you find with the German data.

²Explaining why this is so is one of the important, unresolved puzzles in International Finance, although most people have a hunch that it has something to do with fluctuations in people's perceptions of risk. Results for shorter term maturities are presented in the Chinn and Meredith article mentioned in the previous footnote.

¹The figure is taken from 'Testing Uncovered Interest Parity at Short and Long Horizons during the Post-Bretton Woods Era', by Menzie D. Chinn and Guy Meredith, June 10, 2002 unpublished manuscript, University of California at Santa Cruz. To think about the regression line more carefully, let the forecast error in forecasting the future exchange rate be $\varepsilon = (E' - E^e)/E$. Then, UIP implies:

per unit of foreign currency and the interest rate differential pertains to interest rates denominated in the indicated country's currency.

Figure 3 displays results for Japan. Note how the US dollar has depreciated relative to the Yen on average since the 1970s. Consistent with UIP, the US interest rate has been consistently higher than the corresponding Japanese exchange rate. Interestingly, the two interest rates were virtually equal to each other in 1991, despite the underlying trend in the exchange rate. A puzzle is, why did anyone invest in US assets at this time, when nominal returns were the same in each country's currency, but the trend appreciation in the Yen promised an extra payoff to anyone holding Yen-denominated assets? The solution to the puzzle may be that investors were afraid of the Japanese currency because they were worried about the possibility that it would depreciate, rather than continue on its trend appreciation. At this time one can imagine that there was an unusually great amount of uncertainty in Japan because of the stock and real estate market crashes that had recently occurred. Moreover, as the graph indicates, there had been a bit of a depreciation in the Yen just recently. Investors might have been concerned that spelled the end of the underlying trend. So, this is a period when more seemed to be going on than what is captured in the expected return considerations incorporated into UIP.

Another period in Japan seems puzzling from the point of view of UIP. From the late 1970s to 1985, the US dollar appreciated slightly relative to the Yen. Surely, foreign exchange traders in this period would have started thinking of the Yen as a depreciating currency. And yet, US interest rates were much higher than Japanese interest rates. Why didn't they all move into the US currency, where they not only would have benefitted from the high interest rate, but also would have benefited from a high expected return from holding dollars? Again, a quantitatively convincing answer doesn't exist. However, it probably has something to do with risk. Even though people noticed the continuing dollar appreciation, something must have made them worried about the possibility that the dollar would depreciate. They would have had to have had this concern, despite the fact that the dollar didn't actually start to depreciate until 1985.

Figure 4 reports results for Switzerland. Note that the US interest rate was higher than the Swiss rate in the 1970s and 1980s. At the same time, the US dollar depreciated on average. After this, the interest rate difference is quite small, while the exchange rate was more or less trendless. Again, these longer-run features of the data seem consistent with UIP. At the same time, there are subperiods when UIP does not seem to work so well. For example, as in Japan, the US dollar appreciated from the late 1970s to 1985, yet US interest rates were high during this period. One would think that traders would have started to expect a continuation in the appreciation of the dollar. But if so, then US assets would have looked doubly blessed: they generated a high return, plus anyone in the rest of the world who bought them would have earned an extra payoff in the roundtrip through the foreign exchange market. Moreover, any American investing in Swiss (as well as Japanese) assets would have looked doubly dammed: those assets generated a relatively low return in their own currency, plus the American would have expected to lose money in the foreign exchange market. So, what was going on in people's minds that there was not a wholesale abandonement of Swiss assets? It probably had something to do with risk.

Figure 5 displays results for the British pound. Note how in the early period the US dollar generally appreciated. At the same time, the US interest rate was higher than the foreign interest rates, consistent with the UIP. After this, the interest differential becomes quite small and, consistent with UIP, the exchange rate is relatively trendless. Figure 6 displays results for Thailand. Note that in the 1970s and 1980s the Thai baht generally appreciated. Consistent with the UIP, the US interest rate was higher, at this time, than the Japanese interest rate. After this, the interest rate differentiable became smaller, and the currency a stable exchange rate.

In sum, the UIP seems to capture long-run features of the data. Additional factors not incorporated into the UIP seem to also be at work over short horizons. This reflects that UIP captures the 'fundamentals', and so seems like a good starting point.

4. 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:

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:

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. We can think of velocity as 'how hard money works in transactions'. If all transactions were financed by a single dollar bill flying around quickly from hand to hand, then velocity would be astronomical. In fact, velocity is much lower. When the interest rate is high, the theory says that people will economize on cash balances and velocity will be high.

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 (in the shorter run), the relationship seems less tight. 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.

- 5. 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 $R_{\mathfrak{E}}$ and results in (US) currency appreciation, E goes down.



Figure 3: Deutschemark/U.S. dollar depreciation against the interest differential, 5 year horizon.



















-







