KATHOLIEKE
UNIVERSITEIT
LEUVEN

DEPARTEMENT TOEGEPASTE ECONOMISCHE WETENSCHAPPEN

RESEARCH REPORT 0136

PESO EFFECTS IN THE FORWARD BIAS: EVIDENCE FROM THE PRIVATE ECU

by
P. SERCU
T. VINAIMONT

D/2001/2376/36

Peso Effects in the Forward Bias: Evidence from the Private ECU

by

Piet Sercu	Tom Vinaimont
Tel (+32) (0)16 326 618	Tel (+32) (0)16 326 758
Fax (+32) (0)16 326 620	Fax (+32) (0)16 326 732
piet.sercu@econ.kuleuven.ac.be	tom.vinaimont@econ.kuleuven.ac.be
Department of Applied Economics at K.U.Leuven	Interuniversitair College voor Managementswetenschappen

Naamsestraat 69, 3000 Leuven, Belgium

October 2000

Abstract. Forward rates of European currencies against the private and official ECU exhibit a bias similar to the one found in other data: the Cumby-Obstfeld-Fama (COF) regression coefficients are systematically below unity, and two thirds of them are negative. We use the discount of the private ECU relative to the official ECU as a measure of "diffidence", a term that may cover both sharply fluctuating risk premia as well as Peso risk. Peso risk, in this context, covers not only fears of realignments but also the risk of a meltdown of the private ECU relative to the official one (a notion that receives some support from a time-series analysis of these data). Dichotomizing the data on the basis of the size of the discount in the private ECU), we find that the COF beta strongly depends on the degree of diffidence and that the negative COF coefficients are generated by typically less than 20 percent of the data. If the diffidence factor contains a risk premium, then this risk premium is definitely not the one predicted by Bansal (1997). Nor is the diffidence factor proxying for Huisman *et al.* (1997)'s transaction-cost effects. Thus, Peso risk remains as a strong candidate explanation for the forward bias in this sample.

Acknowledgments. The paper has benefited from useful comments following presentations at KU Leuven, Norges Bank, BI Oslo, and the 1999 VVE, EFMA, and EFA Conferences. We especially thank Geert Bekaert, Gikas Hardouvelis, Christian Rydqvist, Eric Ongena, and Frans de Roon. The authors remain responsible for all remaining shortcomings.

Peso Effects in the Forward Bias: Evidence from the Private ECU

One empirical puzzle in International Finance is the size of the bias in the forward premium as a predictor of future exchange rate changes. The Unbiased Expectations Hypothesis (UEH) posits that, in the regression of exchange rate changes on beginning-of-period forward premia, the slope should equal unity. However, as shown by e.g. Cumby and Obstfeld (1984), Fama (1984), and many others after them—see Froot and Thaler (1990) for a review—the empirical coefficients are not only systematically below unity, but even predominantly negative. True, mainstream pricing models predict that risk premia could lower the slope, but not that much (Hodrick and Srivastava, 1986; Backus, Foresi, and Telmer, 1996; Hollifield and Uppal, 1997; see Bansal (1997) for a dissident view). The empirical results are all the more unexpected as, in unconditional tests taken over long periods, the cross-section of average interest differentials matches the average rates of appreciation quite well (Backus *et al.*, 1996). This contrast between the results of regression tests and means tests, coupled with the low R^2s of these regressions, suggests that the negative COF betas may be generated by a relatively small subset of data points.

Most of the regression evidence stems from tests of mainstream floating currencies against the USD. In this paper, we study rates of individual EMS currencies against the private ECU. As a side benefit, this choice of data allows us to externally validate the findings of other studies: our Cumby-Obstfeld-Fama (henceforth COF) style regressions on private-ECU data do as badly as usual. However, a replication of (by now) standard regressions on new data is not our prime motivation for studying rates against the ECU. What is primarily interesting about this market is the availability, next to the ECU as actually traded (the private ECU), of a theoretical ("official") ECU. Even though it was widely perceived that the private ECU would almost surely converge to the official one in due time, still the former often traded at a substantial discount relative to the official ECU; and this discount varied markedly over time, waxing and waning in line with the (mis)fortunes that befell the EMU's exchange rate mechanism and the prospects for eventual monetary unification. For this reason, the discount between the private and official ECUs provides a candidate proxy for general diffidence in the market, a term that may cover both sharply fluctuating risk premia as well as Peso risk. Peso risk, in this context, may cover not only fears of realignments but also the risk of a meltdown of the private ECU relative to the official one, a risk stemming from the ambiguity and uncertainty about the nature and legal status of the private ECU.

Risk premia and Peso risks are familiar candidate explanations of the COF puzzle. As is well known, a time-varying risk premium generates negative betas only if it exhibits a strong positive correlation with the expectation.[1] This feature is hard to explain theoretically, but the existence of models with the required features can of course never be ruled out a priori. What is appealing about the Peso-risk interpretation of the discount in the private ECU is that it does not require a specific model. Instead, we can rely on an empirical regularity: when bad news hits the market (as indicated by a widening spot discount in the private ECU), then forward rates for the private ECU typically predict that even worse is to come. If the bad news is of the Peso type, however, the actual outcome is never, or rarely ever, in the direction predicted by the forward rate. Thus, if the discount in the private ECU does reflect Peso risk, then in periods of deep discounts we should observe a negative relationship between the forward premium and the actual outcome. We do find that (i) deeper discounts in the private ECU *are* associated with more negative COF regression coefficients, and (ii) after dummying out less than 20 percent of the deepest-discount data, the COF coefficients become typically positive. All this is consistent with the Peso explanation of the forward puzzle.

Still, the discount in the private ECU may be due to other factors than Peso risk. We therefore add three more tests. In the first check, we consider the possibility that the discount in the private ECU is due to a risk-premium. To explain the patterns that we observe in our main tests, we need a risk premium that is not only positively correlated with the expected exchange rate change but also shows even higher positive correlations when diffidence is high. Mainstream (representative-investor) asset-pricing models, where the pricing kernel exhibits little variation over time, are unlikely to produce such patterns (see for example Hollifield and Uppal, 1997), but it is of course impossible to prove that no model of the risk premium can produce the required pattern. All we can do is test specific models. The one we consider is the Bansal (1997) asymmetric risk-premium model which, uniquely to date, does predict the required strong positive correlation between expectation and risk premium under some circumstances, and has obtained empirical support from USD-based rates. We find no traces of an asymmetric risk premium in our data: in the entire sample, the results flatly contradict Bansal's predictions, and even in a subsample with low discounts in the private ECU (that is, tentatively, with low Peso risk) there is little or no evidence in favor or against the asymmetric risk premium. Thus, our measure of market diffidence is surely not proxying for an asymmetric risk premium.

The second robustness check starts from the idea that Peso risk, if present, may have two components: the standard type, namely fears of re-alignments within the system, and the risk of a major confidence crisis (and the resulting loss of value) in the private ECU itself. If we can eliminate one of these risks, and if we then find that the patterns in the COF betas are weakened, then we have an indication that the eliminated Peso factor did contribute to the patterns we

---

[1] In line with the general approach in finance, we define the percentage risk premium as the expectation minus the forward premium, which is the negative of Fama (1983)'s definition. Using a short-hand notation E for the time series of expectations, F for the forward premia, and R for the risk premia, we have $F = E - R$, and therefore (assuming efficiency) $\beta = [\mathrm{var}(E) - \mathrm{cov}(E,R)]/[\mathrm{var}(E) - \mathrm{cov}(E,R) + \mathrm{var}(R)]$.

observed. The type of Peso risk that is easy to eliminate is the risk of a confidence crisis in the private ECU: just consider rates against the official ECU rather than against the private one. We find that, in these regressions, the patterns suggested by the Peso view are still present, but are weakened indeed. This suggests that the Peso risk of a meltdown of the private ECU has contributed to the patterns detected in the first-pass tests.

The third robustness check verifies whether the patterns we have seen are, in fact, due to transaction-cost effects. Huisman *et al.* (1997) argue that, conditional on forward premia being large, the COF beta should be much closer to unity. As periods of unrest are conceivably also periods with larger dispersion in forward premia, our measure of diffidence may simply have proxied for Huisman *et al.*'s transaction-cost effect. Our tests confirm that large forward premia are associated with algebraically larger betas. However, the patterns are much weaker than the ones we see in our main tests; and, worse, even the "large" betas remain negative. To sort out this mess, we then consider a subsample that is characterized by low discounts in the private ECU (that is, tentatively, little or no Peso risk). In this subsample, it turns out, there is no relation between the absolute size of the forward premia and the discount in the private ECU, so that contamination between Peso risk (if any) and transaction-cost effects (if any) cannot occur; also, when we further split up this subsample into high- and low-forward-premium sections, the difference between the regressor variances across these sections is unusually high, which should make a transaction-cost effect easy to detect. Still, we find no traces of a transaction-cost effect. Rather, it appears that the relation between COF betas and the size of forward premia, as observed in the entire sample, is spurious, caused by the latter proxying for Peso risk rather than the other way around.

The paper is structured as follows. Section I describes the relation between the private and official ECUs, conceptually as well as statistically. In Section II we provide the standard UEH tests of Cumby-Obstfeld and Fama, and we show the dismal results of these tests improve the more of the high-diffidence data points are dummied out. In Section III we provide the three robustness tests. Section IV concludes.

I. The ECU Market

In this section, we briefly review the private ECU and its relation to the official one. We then study the deviations between their two spot values, discuss two possible explanations (a regular covariance-risk premium or Peso risks), and provide some preliminary empirical tests of the two competing views.

I.1. The Private and official ECU: A Qualitative Discussion.

Since the early eighties, lenders and borrowers have been able to do money-market operations or forward contracts in private ECUs; soon afterwards, investors could also trade Eurobonds and UK treasury bills denominated in private ECU, futures contracts on ECU interest rates, and options on the private ECU itself. Despite the popularity of these products, there has been some confusion regarding the nature of the private ECU and its relation to the official ECU. In this section we review that relationship.

The official ECU was just a unit of account that was used by EC/EU institutions and its member governments when dealing with the Community. The value of one official ECU relative to a base currency b was defined to be equal to the value of a basket containing known quantities n_j of each of twelve currencies j:

$$S^E_{b,t} \equiv \sum_{j=1}^{12} n_j \, S^j_{b,t} \,. \tag{1}$$

where S^j_b refers to a spot value of currency j in units of an arbitrarily chosen base currency b, and the (upper-case) superscript E refers to the official ECU.

Being a pure unit of account, the official ECU was used only to determine how much was to be paid or received by EU institutions and governments; actual payments occurred in regular currency, at the prevailing exchange rate given by (1). Stated differently, the official ECU did not exist as a separate entity: individuals or firms could not create, buy, or sell official ECUs directly (that is, as distinct from the basket of currencies). In these respects, the official ECU differed from the private one as it existed since the late eighties.

Although, from the beginning, the private ECU has enjoyed the moral support of the EC/EU Commission, it was created by private banks rather than by European institutions. Originally, the nature of this private ECU and its relation to the official basket was quite clear. In the early days, a bank that accepted an ECU deposit actually bought a strip of deposits in the different currencies, with the weights of the deposits being designed so as to pay out one unit of the official basket at the time of expiration. In addition, the ECU clearing banks' policy was, quite logically, to set the spot price per ECU at (or close to) the value of the official ECU. The equilibrium ECU interest rate was then computed as a (harmonic) weighted average of the rates on the component currencies, using a Law-of-One-Price argument like the one spelled out in Appendix A. In those days, also the accounting treatment of the private ECU reflected the then-prevailing nature of the private ECU. Specifically, since in the beginning private ECU deposits or loans were really combinations of money market positions in each of the component currencies, and since the ECU was not a recognized currency for reporting or prudential purposes, any private-ECU deposit or loan was originally shown, in the bank's books, as a series of N deposits or loans.

The private ECU did become a truly separate entity only in the mid-eighties when, following an initiative from the EC Commission, ECU deposits and loans were allowed to be booked as one single accounting entry in a category of its own (as if the ECU were an actual currency), not as a bundle of twelve accounting entries in separate currencies. As of then, ECUs could be bought and sold directly between investors and banks, or among investors, without any need to (un)bundle the deposit or loan from (or into) the component currencies. This innovation made eminent sense from a transaction-cost point of view, but potentially impaired the link between the values of the private ECU and the basket. For various reasons that are beyond the scope of this paper,[2] around 1987 all ECU clearing banks except Kredietbank discontinued their efforts to keep the value of the private ECU close to par. In November 1988, also Kredietbank threw in the towel, as its long position in ECU (and its prudentially required matching short position in the basket) had become unacceptably high. As of then, premia and discounts emerged that could no longer be explained by transaction costs. In the next section, we describe these discounts and look for possible explanations.

I.2. A History of the discount in the Private ECU

Our data are weekly spot rates and 3-month LIBOR interest rates,[3] 17/11/89-to-6/2/98 (418 observations), all taken from Datastream. Figure 1 provides a time-series plot of the week-by-week private-to-official (P/O) premium. For symmetry, this premium is measured as a continuously compounded percentage deviation:

$$P_t = \ln \frac{S^e_{b,t}}{S^E_{b,t}} . \tag{2}$$

In (2), the lower-case superscript (e) refers to the private ECU, the upper-case superscript E refers to the official ECU, and the subscript b refers to an arbitrary base currency.

<Insert Figures 1 and 2 about here>

From the graph we see that there was only one period of pronounced positive deviations. This ECUphoric period, starting mid-1990 and culminating in a +0.8% premium in January 1991, was the time of the emerging consensus that eventually produced the Maastricht Treaty. A few months before the actual signing, end 1991, of the Treaty, the market had already become more level-headed, and during the last weeks of 1991 uncertainty grew when Denmark and France announced referenda. Things turned decidedly ugly when, in June 1992, the Danish referendum rejected Maastricht, and especially in September 1992, when successful speculation against the FIM and SEK spilled over into the Exchange Rate Mechanism (ERM), with the familiar disastrous consequences. Investor diffidence showed up in a 1.6% discount in the private-ECU spot rate. Calm returned soon, but this regained optimism soon proved to be premature. During the second wave of distrust with respect to the ERM—in the summer of 1993, culminating in the

---

[2]See e.g. Sercu (1999) for a review of the institutional background.

[3]The three-month horizon is the only one for which interest rates are available for all twelve currencies that form the ECU.

virtual suspension of the ERM—the spot ECU dropped 1.8% below par at its worst, and remained about 1 to 1.5 percent below par until October 1993. Early 1994, when by and large the intra-ERM cross rates had returned to their pre-1992 levels, also ECU spot rates returned close to par.

Again, the interlude turned out to be brief. As of February 1994, the private ECU started a long (but irregular) decline, culminating in an unprecedented 3% discount relative to the basket in March 1996. Unlike the 1992 and 1993 episodes, this period did not coincide with marked turmoil in the ERM exchange markets. Rather, 1994-1996 corresponds to the period of growing doubts about the Maastricht scenario in general. Among the public, the EC was perceived to be overly bureaucratic and meddlesome, and suffering from a "democratic deficit", and many Germans opposed the common currency. The tide could be stemmed only slowly, with the EC (or, later on, the EU) re-emphasizing the subsidiarity principle and the Kohl-Chirac tandem relentlessly supporting the EMU. Additional factors, especially as of 1996, were the strong convergence among ERM-members' interest and inflation rates and the increasing confidence in the eventual advent of the common currency. In the last quarter of 1997, the private ECU even traded slightly above par, a situation that had not occurred for four years.

The relation between the size of the discount and the problems in the EMU suggests that the P/O discount can be interpreted as a general indicator of investor diffidence. Such mistrust could be related to realignment risk (in member currencies' rates against the official ECU); however, the discount could also indicate risk of a collapse of the private ECU relative to the official one. We address this issue in the next section.

I.3. What may have driven the deviations from par?

We just saw that most deviations (and all of the major ones) were discounts. To detect the possible reasons for a discount, we consider a standard pricing equation,

$$S^e_{b,t} = \frac{(1+r_{e,t,T})\, E_t(S^e_{b,T})}{1 + E_t(R^e_{b,t,T})}, \tag{3}$$

where $r_{e,t,T}$ denotes the private-ECU riskfree rate, and $E_t(R^e_{b,t,T})$ the required return. This equation points out three factors that, each in itself, could have explained a discount. First, the private-ECU riskfree rate, $r_{e,t,T}$, may typically have been too low to ensure parity. Second, the expected future value of the private ECU, $E_t(S^e_{b,T})$, may have been below par, even for a distant horizon T. Third, a standard risk premium may have driven the required return, $E_t(R^e_{b,t,T})$, above the private-ECU riskfree rate. We consider each of these three explanations in turn.

The first potential explanation of the discounts is that there may have been chronic shortfalls in interest rates on the private ECU relative to the theoretical interest rate that would have been consistent with permanent parity (see Appendix A). There is no support for this in our interest-rate data. Figure 2 shows that, as of the second quarter of 1994 (the end of the turmoil around the suspension of the ERM), three-month interest rates stayed quite close to the values consistent with permanent parity.[4] For the preceding period, November 1989 to end March 1994, the graph actually reveals a strongly *negative* correlation (–0.667) between spot-rate and interest-rate deviations: the deeper the ECU spot value dropped, the more the actual rate rose above the theoretical level.[5] Thus, the discounts in the private ECU were surely not caused by shortfalls in the interest earned on ECUs relative to the theoretical rate; rather, they occurred in spite of above-normal interest rates.

A second possible explanation of this discount is that it may reflect a regular risk premium, stemming from covariance with the representative investor's marginal utility. If risk premia are the sole sources of deviations, then the hypothesis must be that the market knows that the two ECUs will ultimately converge, but most of the time the private ECU needs a higher expected return than the official one. The problem with this view is that there are no obvious theoretical grounds why the private ECU should have had higher covariance risks most of the time. True, given such a higher-risk assumption, any setback in the monetary unification process would imply a longer time-to-convergence of both ECUs; but since the difference in covariance-with-a-kernel risk cannot be very large, one would need disconcertingly sharp changes in the expected waiting time to explain the fluctuations we observe in the P/O premium.

Lastly, the discount may have reflected the Peso-type risk of a major confidence crisis in the private ECU, a currency that for most of its life had a rather uncertain legal status. The sources of legal uncertainty were many. First and foremost, against the expectations at the time the currency was created, as of the late 80s the ECU clearing banks proved unable to maintain parity between both ECUs. Thus, investors suddenly woke up to the fact that the relation with the official ECU was (at best) based on trust. True, it was always perceived that, if and when the common currency would be introduced, the private ECU would become assimilated with the common currency (the ECU, later renamed into Euro). In addition, it always seemed quite likely that, if the common-currency plans would fail, the EC/EU would still avoid the embarrassment of a meltdown of the private ECU—after all, the instrument was created at the instigation of the Commission, and the EC institutions were ostentatious early investors in that market. But, however likely the eventual re-convergence to parity, for a long time there was no iron-clad guarantee. For instance, until 1994 the terms and conditions of private-ECU bonds did not even include explicit provisions for conversion, at par, of private-ECU liabilities into the common currency if and when the currency would be introduced. Likewise, the assimilation of the private ECU with the common currency was formally laid down as late as 1996, the time when the new

---

[4]Not surprisingly, this was also the case prior to October 1989, the period of intervention in the private-ECU market (not shown in the graph).

[5]The finding of correlation between the levels is robust to possible unit-root properties: while for first differences the correlation is a mere –0.04, there is a cross-correlation of –0.41 when the change in the P/O premium is lagged once.

name (Euro) was agreed upon. And even then doubts still existed whether the Euro would ever materialize, or, if so, whether, say, US courts would accept the conversion of private ECUs (and other Euroland currencies) into Euros. In short, a catastrophic confidence crisis and loss of value, however improbable, could never be quite ruled out. In actual fact, of course, such a meltdown has not occurred—but this is ex post. If there was a Peso-type risk indeed, these hindsight observations would just reflect the most likely outcome rather than the entire ex-ante distribution.

Two qualifying remarks are in order. First, while the distinction between risk premia and Peso risks is clear-cut in a representative-investor model (where risks and their price change only slowly), that distinction becomes blurred in more complicated models, where risk premia are potentially much more time-varying. Thus, it is impossible to say, at this stage, whether the P/O premium captures Peso risk, or unorthodoxly sharp fluctuations in the risk premium, or both. For this reason, we initially use the more agnostic term "diffidence", covering both Peso risk and markedly changing risk premia. Qualitative evidence in favor of a "diffidence" interpretation of the discount in the private ECU is provided by its strong relation with setbacks in the monetary unification process—the Danish No-vote, the ERM crises in 1992 and 1993, and the 1995-6 period of general disaffection with the Maastricht plans. These events are logically associated with higher Peso probabilities and/or higher risk premia. A second caveat is that exchange rates for the private ECU contain not just the private-ECU factor, but also factors related to each of the base currencies. As deep discounts in the private ECU also coincided with periods of distrust in the ERM, we need to keep in mind that Peso risk, if any, is quite likely to pick up also the standard type of Peso event, realignment-risk. In short, market diffidence (as proxied for by the P/O premium) may cover sharply rising risk premia, the risk of an (unlikely) meltdown of the private ECU, and realignment risk. In Section II we proceed under the assumption that Peso risk is the main factor behind the P/O premium. Section III provides additional support for this view.

II. Main Tests for Peso Effects

II.1. Standard Tests of the Unbiased Expectations Hypothesis

To set the stage for our tests of the risk-premium and Peso theories, we first provide the results for the standard Cumby-Obstfeld-Fama (COF) regression test of the unbiased expectations hypothesis,

$$s^e_{b,t,t+\Delta} = \alpha_b + \beta_b \, FP^e_{b,t,t+\Delta} + \eta_{b,t,t+\Delta} \,. \qquad (4)$$

The observations are weekly, as before. As, for reasons of data availability, our interest rates are three-monthly, we use. in our OLS regressions, the Hansen and Hodrick's (1980) standard deviations to cope with the overlap in the holding-period returns. As we have many equations in this paper, each estimated (often rather imprecisely) over eleven separate currencies[6] and four periods, we have condensed the estimation results of each equation into a few key statistics: the mean slope coefficient across all eleven currencies, the number of estimates that exceed unity, the number above [below] zero, and the number of significantly positive [negative] coefficients. The original equation-by-equation results are available on request.

Table 1. Summary of main regression results (I)

$$s^e_{b,t,t+\Delta} = \alpha_b + \beta_b\, FP^e_{b,t,t+\Delta} + \eta_{b,t,t+\Delta}$$

Sample Period	mean β	# >1	# >0	whereof signif 5%	# <0	whereof signif 5%
Total (Oct 98 to Feb 98)	–0.65	1	3	1	8	0
I (Oct 98 to Aug 92)	–0.31	3	5	1	6	2
II (Sept 92 to Dec 94)	0.68	4	6	3	5	0
III (Jan 95 to Feb 98)	–0.70	4	4	1	7	1

Key to Table 1. The table summarizes results from 11 regressions, one for each base currency b (BEF, DEM, NLG, PTE, ITL, IEP, GRD, DKK, ESP, GBP, FRF). The regressand is the three–month percentage spot-rate change for the private ECU, $s^e_{b,t,t+\Delta}$, the regressor is the beginning–of–period three-month forward premium, $FP^e_{b,t,t+\Delta}$. In the regressions we follow Hansen and Hodrick (1980), that is, OLS with standard deviations that take into account the overlap in the holding periods.

The summary statistics for equation (4) are shown in Table 1. In terms of their mean, the slope coefficients in Table 1 are hardly better than the meta-average (–0.80) reported by Froot and Thaler: across the eleven β coefficients estimated from our total sample, the average is –0.65 for private-ECU rates. Among the regression coefficients for rates against the private ECU, only one individual coefficient exceeds unity, and eight are even below zero (although not significantly so). The subperiod results are similar. In terms of the mean, the only noteworthy exception occurs in the (relatively short and noisy) second period. However, even for that period the numbers of coefficients exceeding unity or zero are hardly better; and the more recent data (subperiod III) again display the familiar picture of predominantly negative coefficients.

In their review of the literature, Froot and Thaler (1990) discuss three candidate explanations of the forward-bias puzzle that would be consistent with rational expectations: (i) risk premia, (ii) the need to learn about changing economic circumstances, and (iii) Peso problems. While it is known that a risk premium should obscure the relationship between exchange rate changes and forward premia, the observed bias is much larger than conventional risk theories predict. To our knowledge, only Bansal (1997)'s model has achieved some success in this area. We test this model in Section III. The need to learn about a change in, for instance,

---

[6]As the BEF and LUF always trade at par, we treat them as one currency.

monetary policy, should be a temporary phenomenon. More in general, while errors in expectations may explain coefficients below unity, it is hard to see how they could have induced pervasively negative coefficients across all currencies and time periods. Thus, we initially focus on Peso problems.

II.2 Testing for Peso Risk in Rates against the Private ECU

Assume that there is a Peso risk (here, a sudden loss of confidence in the private ECU and/or realignment risk). Such a feature can explain the negative coefficients that we have observed if, sufficiently often, the forward rate points into the direction of a low-probability but major event that is never, or hardly ever, observed in smallish samples. In general, the problem with this view is that it invokes a factor—deviations between the true expectation and the small-sample mean—that is as hard to observe as an unspecified risk premium. To lend credibility to this Peso-explanation, we need an independent indicator of the degree of Peso risk, so that we can verify whether the negative relationship between prediction and outcome actually is more likely when Peso risk is higher. The advantage of our data set is that a plausible candidate indicator is readily available, namely the discount in the private ECU.

We have already mentioned circumstantial evidence that such a Peso-mechanism, if present in the data, might be associated with the discount in the private ECU. Recall that there is a strong negative correlation between the P/O premium and the difference between the interest rate on private ECUs and the average interest rate of the component currencies (Section I.3). Thus, when negative news about the private ECU reaches the market, not only the spot rate falls, but also the private-ECU interest rate rises relative to other interest rates—that is, the forward premium typically predicts a further drop of the private ECU. Now bring in Peso risk. If the expectation imbedded in this higher forward premium is systematically negated by the subsequent return, then the forward discount is typically followed by an appreciation of the private ECU, and the beta is negative.

To formally verify whether the negative empirical relationship is indeed due to episodes of deep P/O discounts—that is, tentatively, episodes of high Peso risk—we condition the betas on the degree of discount on the private ECU. The technique is borrowed from Bilson (1981) and Huisman *et al.* (1997), who hypothesize that the regression relation may be different conditional on a large forward premium (Bilson) or a large variance in the day's cross-section of forward premia (Huisman *et al.*). Accordingly, they estimate the regression of the type

$$s^{e}_{b,t,t+\Delta} = \alpha_b + \beta_{b1}\,[I_{b,t} \times FP^{e}_{b,t,t+\Delta}] + \beta_{b0}\,[(1-I_{b,t}) \times FP^{e}_{b,t,t+\Delta}] + \eta_{b,t,t+\Delta}\,, \qquad (5)$$

where $I_{b,t}$ is set equal to unity if observation t is characterized by a high premium (Bilson) or a high cross-sectional dispersion (Huisman *et al.*). Analogously, we estimate (5) and set $I_{b,t}$ equal to unity if the day-t P/O premium is among the N percent most negative of all ranked P/O premia, where the percentage N is consecutively set at 1%, 5%, 10%, 20%, and 40%. If the P/O premium is related to Peso risk, then observations where the private ECU is closer to par should be less affected by the Peso effect, while the deep-discount data should be more strongly affected. Thus, Peso-risk (as proxied for by the size of the P/O premium) has the following testable implications. First, for any given cut-off criterion (1%, 5%, etc.) the beta for the deep-discount data points should be algebraically smaller (for instance, more negative) than the beta of the closer-to-par observations, because these deep-discount data have higher Peso risks. Second, when we make the definition of a deep discount more stringent (going, for instance, from a 40-percent percentile cut-off criterion to a one-percent percentile criterion), both sets of betas should drop steadily.[7] The reason is that, in either sample, Peso risk gets more weight: the higher-risk sample becomes less and less diluted by medium-risk data, while the smaller-risk sample contains more and more medium-risk data.

<Insert Table 2 about here>

Panel A of Table 2, also illustrated graphically in the left-hand-side part of Figure 3, bears out these hypotheses. First, for any given cut-off criterion the average large-discount betas are invariably below the small-discount betas and the number of negative slopes is systematically larger in the deep-discount sample. Second, when going down the table (that is, making the criterion of a "large" discount more severe) we see that both sets of average betas drop steadily, reaching a highly negative –7.56 for the one-percent deepest-discount observations. For the deep-discount sample, where all but one estimate are already negative with the weakest cut-off value (40%), we of course see no trend in the number of negative betas; but such a trend is very clear in the small-discount sample, where the number of negative slopes steadily rises from four to eight as we tighten the high-risk cut-off criterion. All this is consistent with the Peso-risk hypothesis. Equally interesting, for both the 20- and 40-percent cutoff criteria the small-discount betas are, on average, positive. That is, if we classify observations on the basis of the P/O deviation, then the negative slopes in the standard COF regression turn out to be generated by less than 20 percent of the data points. This is what one would expect in light of the results from unconditional tests.

---

[7]Thus, the high- and low-risk betas are not necessarily centered around the unconditional beta. The unconditional beta is a weighted average of (i) the conditional betas *and* (ii) the beta of the conditional means, with as respective weights the contributions of the conditional variances and the variance of the conditional means to the total variance of the regressor. If the conditional means differ across the subsamples, then a fall of one of the conditional betas is not necessarily accompanied by a rise in the other one. Conversely, if both conditional betas fall, then by implication the beta of the conditional means must rise.

An alternative test design might have been Non-Linear LS where the beta as a function of the P/O premium. However, the potential non-stationarity of the P/O premium (Section II.3) would make the t-tests hard to interpret.

Figure 3: Average regression coefficients for deep-discount versus small-discount observations

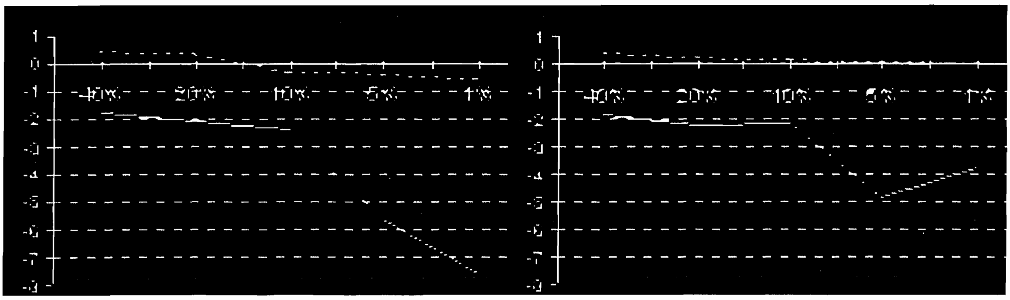

rates against the private ECU rates against the private ECU

— mean coefficients for deep-discount data

- - - mean coefficients for smaller-discount data

Key to Figure 3: We define a data point to be deep-discount if it belongs to the 1% (5, 10, 20, 40%) most negative discounts in the sample. The bold line shows the average beta for deep-discount data points, the regular line the average beta for smaller-discount data, in the regressions from Table 2.

III. Three Corroborating Tests

While the above evidence is consistent with the P/O premium proxying for a Peso-type risk, also other factors may have explained the empirical results. In this section we critically discuss some alternative explanations. Before turning to theories that require additional tests, we first dispell the notion that our empirical results are due to a few observations, or to a particular time period.

The critical reader may have noted, from the time series plots of spot and interest rates, that there are two outlier observations that superbly conform to the Peso theory. In both the fall of 92 and the summer of 93, the private ECU dipped dramatically and ECU interest rates soared above the weighted average of the EMS rates—that is, forward rates predicted worse to come. In actual fact, both times the spot rate sharply returned towards par. Could these two outlier episodes dominate the results? We argue they cannot. First, if the two outliers were behind the results, then imposing stricter and stricter definitions of a "large" FP would not have affected the small-FP betas. That is, the fact that both the small- and large-FP betas drop when we raise the cut-off point means that the observations between, say, the 60th and 80th percentile value, seem to be more risky than those below the 60th percentile and less risky than those above the 80th. Second, if the 93-93 two outliers would dominate the sample, they should have had a serious impact on the COF regressions in the middle subperiod results. Table 1 however show that this is not at all the case; in fact, in that subperiod the COF regressions do somewhat better than average. A similar argument can be invoked against the view that the deep P/O-discount sample may just proxy for a particular time period where an unidentified factor may have driven the results. If so, this subperiod would be the third subperiod where, as Graph 2 shows, the private

ECU traded at systematically low prices. Yet, the COF tests summarized in Table 1 reveal no extraordinary behavior for that subperiod.

We now turn to other theories, like asymmetric risk premia or transaction costs, that require additional tests.

III.1 Is the P/O Premium proxying for an Asymmetric Risk Premium?

In the preceding sections, we have ignored the possibility that changes in the P/O premium may capture sharply fluctuating risk premia. Thus, risk premia lead to negative COF betas provided that they are strongly positively correlated with expectations (see footnote 1). As Fama (1983) already remarks, such a phenomenon is hard to explain in general. In view of our test results, the risk-premium hypothesis becomes even harder to understand: we now need a theory that explains why this positive correlation mainly exists when the P/O premium is high.

Of course, the risk-premium hypothesis is, at its most general level, irrefutable; all one can do is test specific models. For example, Hollifield and Uppal (1997) reject that the COF betas are generated by a model of the Dumas (1991) type, a general-equilibrium model with an endogenous risk premium. In this paper, we consider a model, due to Bansal (1997), which does predict strong variation in the risk premium and has received empirical support in USD-based data. Starting from pricing-kernel interest-rate theory, Bansal argues that the risk premium may be approximately quadratic in the interest rate differential, thus producing negative COF coefficients for positive forward premia and vice versa. He accordingly conditions the COF regression coefficient on the sign of the forward premium rather than its absolute size, and does find that negative forward premia go together with positive betas and vice versa.

We first replicate the Bansal tests on our data. Panel B of Table 2 contains the results. On the bright side (from Bansal's perspective), the average β for positive premia is doing particularly badly indeed (–3.52). However, upon closer inspection of the individual slopes (not shown) this figure turns out to be driven by one extreme outlier, –34.33 for the DEM, a coefficient generated by a mere six observations (out of 417); deleting this outlier actually switches the mean to a positive 1.46, which is against Bansal's predictions. Equally against the theory, the average β for negative-forward-premia observations is below zero rather than above zero (–0.92; –0.80 without the DEM). The tallies of negative individual coefficients in each subsample suggests the same pattern: when forward premia are negative, betas should be positive, but in actual fact this is observed in two cases only—which is even less than the number obtained when premia are positive. In short, we find no support for the Bansal hypothesis.

This finding suggests that the results of our tests for Peso-effects, in Section II.2, are quite unlikely to be caused by P/O premia proxying for Bansal-type risk premia. However, the opposite may still be true: Peso effects may be obscuring the Bansal risk premia. To verify whether Bansal's view fares better when we control for Peso risk, we do a four-way split of the sample: we first dichotomize the data on the basis of the P/O premium (using the 40th-percentile criterion, so as to have enough data in each subsample), and we then subdivide each sample on the basis of the sign of the forward premium. We are particularly interested in the cells with smaller P/O discounts, where peso risk (if any) should be weaker and the transaction-cost effect more visible. The results in Panel A of Table 3 show that, once Peso risk is removed, the contradictions with Bansal's result are lessened. In accordance with the Bansal effect we see that, given that the P/O is small, positive FPs are associated with negative average betas and vice versa. The average coefficient for negative forward premia in the absence of Peso risk, around 0.90, would definitely please supporters of the UEH view. However, in this test the numbers of positive and negative betas do not confirm the evidence from the means: behind the positive average there are only four positive individual estimates, which is less than the number of positive betas behind the negative average that we observe when FPs are positive. Thus, even after controlling for (tentatively) Peso risk there is no strong evidence of an asymmetric risk premium.

To sum up: the evidence of our first-pass tests provide flatly contradicts the Bansal prediction. While, after controlling for the P/O premium, these contradictions disappear, there still is no firm evidence of an asymmetric risk premium. Thus, we reject the notion that the P/O premium may have proxied for a Bansal-type risk premium. Pending the arrival of other risk-premium models that predict a strong positive correlation between expectation and risk premium, we tend to retain the Peso-risk explanation.[8]

III.2. Testing the Effect of Eliminating One Type of Peso Risk

In the literature, Peso risk stands for the small risk of large changes, like realignments. As stated before, the private ECU might have been subject to another Peso-type risk, a possible collapse of the private ECU. In this section we first provide empirical indications that the latter type of risk has existed. We then test whether this type of risk has contributed to the general beta-patterns we have observed in Section II. If so, this would further lend credibility to our contention that the patterns have to do with Peso risk.

To obtain a clue as to whether the risk of a collapse of the private ECU market was entirely absent, or at least one of the factors in the P/O premium, we can resort to unit-root tests. The complete absence of a Peso-type melt-down risk means that eventual convergence between the private and official ECUs is a certainty. Thus, in that view the discount would be a temporary phenomenon and cannot have a unit root. Conversely, the presence of a unit root must have meant that there was a Peso risk in the private ECU relative to the official one. Figure 1 already

---

[8]The discrepancy between our and Bansal's test results could of course be due to a different sample, but also to the fact that Bansal's theory refers to two distinct economies. In our case, the (relative) risk of the private ECU cannot be related to the (relative) risk of a particular economy. At most, the private ECU is related to the official one, which, in turn, refers to a group of economies.

suggests that the P/O premium exhibits strong persistence indeed. Table 4 provides Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) test statistics on the P/O premium. The PP test rejects the unit-root hypothesis, but only at the 10% level (and very marginally so). Judging by the ADF statistic, in contrast, the hypothesis of a unit root is quite acceptable. We conclude that there are credible, but not quite conclusive, indications of one type of Peso-type risk, a potentially disastrous loss of confidence in the private ECU.

Table 4: ADF and PP tests for unit roots in the P/O Premium

model	ADF t–test	PP t–test	10% critical value
Constant, no trend	–2.1110	–2.7884	–2.75
Constant, trend	–2.2692	–3.1557	–3.13

Key to Table 4. The table shows Augmented Dickey–Fuller and Phillips–Perron tests for the null hypothesis of a unit root in the percentage difference between the values of the private and official ECUs. The data are 418 weekly observations from October 1989 to February 1998.

If the risk of a collapse of the private ECU was one source of Peso risk (next to regular realignment risk), it should have affected also the modified COF regressions of Section II.2. To verify this, we rerun the modified COF regressions (5) using, this time, exchange rates and forward premia against the official ECU.[9] Exchange rates against the official ECU are subject to realignment risk, as are rates against the private ECU. But unlike the private ECU, the official ECU can change value *only* because of realignments and is not otherwise subject to possible crises of confidence.

Panel C of Table 2, shown graphically in the right-hand-side part of Figure 3, summarizes the results of regressions with the official-ECU data. Although the broad picture is similar, the effects are weaker and somewhat more blurred. As expected under the Peso hypothesis, for any cut-off criterion the average smaller-discount beta is algebraically below the deep-discount slope coefficient, and the smaller-discount betas do drop steadily as we go down the table. For the deep-discount betas, however, which should be most sensitive to Peso risk, the drop in the average coefficients is less steady, and does not nearly go as far as what we found in the case of exchange rates against the private ECU. Also the trends in the numbers of negative coefficients are weaker in Panel B than in Panel A. All this again suggests that the difference between both sets of data—risk about the future P/O deviation—may have added a Peso risk of its own.

[9] Forward rates for the official ECU are not directly observable, but they follow from the component currencies' forward rates using the no-arbitrage condition (A.3) in the Appendix.

III. 3 Is the P/O Premium proxying for Transaction-Cost Effects?

To put our results into perspective, we also compare them to results from similar tests that sort the data on either the absolute size of the forward premium (rather than on the P/O premium), as in Huisman *et al.* (1997). This test starts from the fact that real-world markets are subject to friction. As a result of such transaction and information costs, uncovered interest arbitrage cannot perfectly align expected exchange rates and forward premia. Most of the time, expectations of exchange rate changes are, moreover, so small that this friction-induced noise between expectations and premia largely obscures the potential relation between the two. However, there may be occasions where the market does expect large changes; and if the impact of friction is essentially unaffected by the size of the expected change, then in these instances the signal-to-noise ratio is relatively large. Highly positive or negative forward premia should, therefore, be better predictors than small premia. Cast in familiar statistical terms: the COF regression suffers from a standard errors-in-the-regressor bias towards zero, and for a given variance of the noise term this bias can be reduced by constructing a sample where the variance of the regressor is larger.

To test this view, we again run regression (5), this time setting $I_{b,t}$ equal to unity whenever, in terms of the absolute value of the forward premium, the observation exceeds the 1% (or 5, 10, 20, 40%) upper percentile. The Huisman *et al.* predictions are similar to our hypotheses *re* Peso-risk: for any given cut-off value the large-FP coefficient should be closer to unity than the smaller-FP one; and the stricter the cut-off value, the closer to unity the large-FP coefficients should be. In addition, the large-FP beta should be close to unity, or at least positive.

Panel D of Table 2 contains the results. These are not encouraging: while average large-premium coefficients are systematically larger, algebraically, than smaller-premium ones, the effect is not nearly as strong as the one we documented for a sort on the size of the P/O discount. Also, the negative trending in the average coefficients is less pronounced, and the non-parametric support for such negative trending, from the numbers of negative estimates, is weaker (for the large-FP samples) or even entirely absent (for the small-FP samples). The most damning information, however, is that almost all average coefficients remain negative. The lonely exception, for the one-percent largest forward premia, remains far from unity (0.17), and behind this mean value there still are six negative coefficients.[10] Equally disconcerting, when we sort on the basis of the size of the forward premium we need to discard more than 95 percent of the data points before we get a positive coefficient. In light of the unconditional cross-sectional tests, where there is a strong positive relation between mean exchange rate changes and forward premia, this finding would implausibly mean that the cross-sectional evidence is due to very few

---

[10]Huisman *et al.* (1997) do obtain coefficients close to unity for very strict definitions of high premia; for the extreme fractiles, their coefficient even exceeds unity. However, from Nissen (1996), these results disappear when their numeraire-invariance constraint is dropped; and that constraint is rejected by the data.

data points. In contrast, when sorting on the basis of the P/O discount we needed to discard less than 20 percent of the observations to find back that positive relationship.

The prevalence of negative coefficients, and the absence of coefficients close to unity, is inconsistent with the pure errors-in-the-regressor argument; rather, there must have been another factor at work, like Peso risk. Peso risk may have affected the results of our transaction-cost tests in two ways. First, there may have been an transaction-cost effect on top of the Peso-bias. Such an addition of Peso risk on top of a transaction-cost effect would explain why the betas in Panel D of Table 2 are too low but seem to have, otherwise, the predicted pattern. The second possible view is that the patterns reported in Panel C are spurious because forward premia simply proxy for P/O discounts.[11]

To gain some insight into the *a priori* credibility of the second explanation, we compute the Spearman rank correlation (SRC) between the P/O premium and |FP|. Over the entire sample, the SRC is typically weak (the average is –0.04, with five positive individual coefficients), but the correlations also seem to vary substantially depending on the size of the P/O premium. To preserve enough observations in the subsamples to be formed in later tests, we again use the 40% cutoff value for the discount in the private ECU to form a high- and a low-risk subsample, and we compute SRCs within each subsample. We find that, conditional on a deep discount the average SRC is nontrivial and negative (–0.29, with eight negative individual values), while in the smaller-discount sample there is virtually no link between the rankings (the average SRC is just 0.04, with five negative coefficients). This correlation pattern implies the following testable implications:

- If the Peso effect is genuine and the transaction effect entirely spurious (that is, merely caused by correlation between both ranking criteria), then if we sort on the basis of forward premia we should see virtually no symptoms of a transaction-cost effect within the low-Peso-risk sample; and the transaction-effect symptoms within the high-Peso-risk sample should be weaker than the raw Peso effect of Table 2A because the proxying is far from perfect.
- In contrast, if there is a genuine transaction effect on top of the Peso-effect, then we should observe symptoms of a transaction-cost effect within both the high- and low-Peso-risk samples.[12]

---

[11] A third possible view would be that the symptoms of Peso risk are the spurious ones, and the transaction-effect genuine. This view is discarded. First, the negative signs in Table 2A and 3D cannot be explained by errors in variables, but could be due to a Peso effect. Second, the symptoms of Peso effects are stronger than those of transaction cost effects, which is implausible if the ranking on P/O is an imperfect proxy for the ranking on |FP|.

[12] Ceteris paribus, this hypothesis also predicts stronger symptoms in the high-risk sample, because the genuine transaction-cost effect is reinforced by a proxy Peso-risk bias. However, we argue, below, that the power to detect transaction-cost effects in the high-risk sample is lower, which tends to invalidate this additional test criterion.

Table 5. Comparison of Regressor Variances Across Cells

	(a) average var(FP ; |FP| > c) and var(FP ; |FP| ≤ c), 10^{-4} (b) average of the currency-by-currency ratio of large- and small-FP variances			
criterion	5th |FP| percentile	10th |FP| percentile	20th |FP| percentile	40th |FP| percentile
small peso risk	(a) 4.46 — 0.11 (b) 18.85	(a) 2.83 — 0.09 (b) 19.96	(a) 1.55 — 0.07 (b) 21.08	(a) 0.83 — 0.05 (b) 26.85
large peso risk	(a) 0.34 — 0.08 (b) 6.09	(a) 0.25 — 0.06 (b) 5.74	(a) 0.31 — 0.05 (b) 7.84	(a) 0.21 — 0.03 (b) 10.14

Key to Table 5. The table shows the variance of the regressor, FP, in each of the cells, according to various ranking criteria. "Large Peso Risk" indicates that the observation belongs to the set with the 40% deepest P/O discounts. Data are also ranked on the basis of the absolute size of the forward premium, using the 5% (10, 20, 40%) percentiles as alternative cutoff values. Within each cell defined by a P/O and |FP| criterion, we show the variance (10^{-4}) of the forward premia conditional on a large and on a small |FP|, respectively. We also show the average of the currency-by-currency ratio of the high- and low-FP variances.

For regression tests of the above hypotheses, we form four subsamples on the basis of two sets of dummies—the first one indicating whether or not, for observation t, the P/O premium belongs to the 40% deepest discounts, and the other dummy indicating whether or not the time-t forward premium belongs to the 40% (20, 10, 5%) largest[13]—and we allow the regression slope to differ across the four cells. Prior to discussing the results, we provide some information on the variance of the regressor within each cell. If, for instance in the low-Peso-risk sample, the variability of the forward premia would be low, then there would not be much of a difference between the subsamples of high- and low-FP data, and we would not expect to see a strong transaction-cost effect. Actually, as shown in Table 5, for any cut-off criterion for |FP| the variance of the regressor is systematically larger in the small-Peso-risk sample than in the high-risk one, and so are the differences between the large-FP- and small-FP-samples' variances. Thus, any transaction effect present in the data should be easier to detect within the small-Peso-risk sample.

Given this information, we now discuss the regression results in Table 3B. First consider the subsamples with smaller P/O discounts. In this subsample, peso risk should be weaker; and the transaction-cost effect, if any, cannot be spurious and should be easiest to detect. Table 3B reveals no trace of any transaction-cost effect within the small-discount data. Against the predictions of the transaction-cost view the average large-FP betas are actually always somewhat below the small-FP ones, and they do not trend downward when the definition of "large FP" is tightened. Likewise, in these low-Peso-risk cells the large-FP betas are less often positive than the small-FP ones, and there is no trend in the number of positive signs when the definition of "large FP" is made stricter. In short, within the subset of small-Peso-risk observations, we do not

[13]Thus, the rankings are made independently. That is, we did not rank forward premia separately within the deep- and small-discount samples, respectively. Our way of splitting the data has the advantage that the definition of a large forward premium remains comparable to the one used in Panel D of Table 2.

detect any of the symptoms of bias that could be caused by transaction costs. In contrast, in the high-Peso-risk subsample the large-FP betas are less negative than small-FP ones; there is some downward trending in the coefficients for stricter definition of "large FP", especially in the first cell, where discounts and FPs are both large; but the differences between large- and small-FP betas as well as the downward trending are much weaker than in Table 2A. Recall, lastly, that on the basis of the variabilities in the forward premia any transaction-cost effect present in the data should be hardest to detect in the high-risk sample. In light of all these considerations, we reject the view that the anomalies in Table 2D are driven by a combination of transaction-cost effects and Peso-biases, and we retain the hypothesis that the symptoms of transaction-cost bias are spurious.

IV. Conclusion

We find that the forward rates against the private ECU exhibit a similar bias as other data: the Cumby-Obstfeld-Fama (COF) regression coefficients are systematically below unity, and mostly negative. To test whether Peso risks explain these results, we would like to have an indicator of Peso risk—here not only the risk of realignments but potentially also the (tiny) chance of a major crisis of confidence in the private ECU. The discount of the private ECU relative to the official one may provide such a measure. To gain some insight, we verify whether the size of the premium affects the COF coefficients in the way predicted by the Peso view: a deeper P/O discount should be associated with a lower (more negative) beta. The empirical results are in line with the latter view. When we consider rates against the official ECU, the effect is weaker, which suggests that the risk of a collapse of the private ECU may have been a source of Peso risk, next to standard realignment risk.[14] Of course, one can never reject the hypothesis that the P/O premium may have proxied for sharply rising risk premia, too; all one can test is whether specific risk-premia models are associated with the P/O premium. We find no evidence or Bansal's (1997) asymmetric risk premium, a theory that does predict negative COF slopes. Even though, after controlling for Peso risk, the data no longer flatly contradict his hypothesis, there are no grounds for the hypothesis that the P/O effect is just a proxy for an asymmetric risk premium. Also the transaction-cost model (Huisman *et al.*, 1997) is particularly bad at generating positive betas, and the other indications in favor of this alternative explanation of the forward puzzle appear largely spurious: the size of the forward premia is proxying for the depth of the P/O discount.

[14]The tests for a unit root in the P/O premium, a characteristic that would be incompatible with the view that there is no Peso risk, provide no clear-cut evidence as to whether there were fears of a massive confidence crisis in the private ECU.

Appendix: The interest rate on a pure basket ECU

We denote the private-ECU exchange rate (observed at time t) and the private-ECU risk-free rate of return (on money market operations that start at t and expire at T) by $S^{e}_{b,t}$ and $r_{e,t,T}$, respectively, and we add asterisks to denote a pure basket ECU, defined as one that always trades at par (as was the case until November 1989). Thus, by definition,

$$S^{e*}_{b,t} = \sum_{j=1}^{N} n_j\, S^{j}_{b,t}\,, \tag{A.1}$$

$$S^{e*}_{b,T} = \sum_{j=1}^{N} n_j\, S^{j}_{b,T}\,. \tag{A.2}$$

Let $F_{t,T}$ denote a forward rate set at time t for delivery at time T. From (A.2) and interest rate parity, respectively, it then follows that

$$F^{e*}_{b,t,T} = \sum_{j=1}^{N} n_j\, F^{j}_{b,t,T}\,. \tag{A.3}$$

$$S^{e*}_{b,t}\,\frac{1+r_{b,t,T}}{1+r^{*}_{e,t,T}} = \sum_{j=1}^{N} n_j\, S^{j}_{b,t}\,\frac{1+r_{b,t,T}}{1+r_{j,t,T}}\,. \tag{A.4}$$

Dividing through both sides by $S^{e*}_{b,r}\,(1+r_{b,t,T})$ and using (A.1), we obtain

$$\frac{1}{1+r^{*}_{e,t,T}} = \sum_{j=1}^{N} \frac{S^{j}_{b,t}\, n_j}{\sum_{j=1}^{N} S^{j}_{b}\, n_j}\,\frac{1}{1+r_{j,t,T}}\,.^{15} \tag{A.5}$$

Thus, for a pure–basket currency, the effective gross risk-free rate is a harmonic weighted mean of the effective gross rates on the component currencies, not a simple weighted mean as is sometimes thought. In practice, the difference between the simple and harmonic weighted mean is minor.

---

[15] An alternative way to obtain (A.5) is to realize that, when (A.2) holds, then a strategy of investing amounts $n_j/(1+r_{j,t,T})$ of each currency will produce one ECU at time T. The cost of buying this strip of deposits, measured in ECUs at par (use (A.1)), is given by the right-hand side of (A.4). Lastly, by the law of One Price, the time-t cost, in terms of ECUs, of one synthetic time-T ECU must equal the left-hand side.

References

Backus, D., S. Foresi, and C. Telmer, 1996, Affine Models of Currency pricing, *NBER Working Paper* 5623

Baldwin, E. R., 1990, Reinterpreting the Failure of the Foreign Exchange Market Efficiency Tests: Small Transaction Costs, *CEPR Discussion Papers*, 407

Bansal, R., 1997, An Explorartion of the forward Premium Puzzle in Currency Markets, *Review of Financial Studies* 10, 369-403

Bilson, J.F.O., 1981, The Speculative Efficiency Hypothesis, *Journal of Business* (54), 435-451

Cumby, R. and M. Obstfeld, 1984, International Interest Rate and Price Level Linkages under Flexible Exchange Rates: A Review of Recent Evidence, in J. Bilson and R. Marston eds. *Exchange Rate Theory and Practice*, University of Chicago Press, Chicago, IL.

Dooley, M.P., and J. Schafer, 1983, Analysis of Short–Run Exchange Rate behavior: March 1973 to November 1981, in Bigman, D. and T. Taya (eds.), *Exchange Rate Instability: Causes, Consequences and Remedies*, Ballinger, Cambridge, MA

Fama, E., 1984, Forward and Spot Exchange Rates, *Journal of Monetary Economics* 14, 319–338

Froot, K. and R. Thaler, 1990, Anomalies: Foreign Exchange, *Journal of Economic Perspectives* 4, 179–192

Hansen, L. P. and R. J. Hodrick, 1980, Forward Exhange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis, *Journal of Political Economy* **88**(5), 829-853

Hodrick, R. J. and S. Srivastava, 1986, The Covariation of Risk Premiums and Expected Future Spot Exchange Rates, *Journal of International Money and Finance* 5, 5-22

Hollifield, B., and R. Uppal, 1997, An Examination of Uncovered Interest Parity in Segmented International Commodity Markets, *Journal of Finance* 52, 2145–2170

Huisman, R., K. Koedijk, C. Kool, F. Nissen, 1998, Extreme Support for Uncovered Interest Parity, *Journal for International Money and Finance* **17**, 211-228.

Nissen, F., 1997, *International Financial Markets Dynamics. An Empirical Investigation of Exchange Rates, Interest Rates, and Stock returns*, Doctoral Dissertation, Universiteit Maastricht

Sercu, P. 1999, From ECUphoria to ECUphobia and Back Again: the Lives and Loves of a Strange Currency, in Cajot, J. (ed), 1999, *Bedrijfsbeheer en Taalbedrijf, Jubileumboek Dertig Jaar Vlekho*, Averbode, Belgium.

Table 2. One-by-One Tests of the Competing Hypotheses on the Forward Bias

Summary statistics on β_{b1} and β_{b0} in $s_{b,t,t+\Delta} = \alpha_b + \beta_{b1}\,[I_{b,t} \times FP_{b,t,t+\Delta}] + \beta_{b0}\,[(1-I_{b,t}) \times FP_{b,t,t+\Delta}] + \eta_{b,t,t+\Delta}$

prediction about β and criterion for $I_{b,t} = 1$	mean β_{b1}	# >1	# >0	whereof signif 5%	# <0	whereof signif 5%	mean β_{b0}	# >1	# >0	whereof signif 5%	# <0	whereof signif 5%
Peso risk I: β should be worse when P/O discount is deep	Panel A: Exchange rates against the Private ECU; $I_{b,t}$ indicates deep discount											
	β for deep P/O discounts						β for small P/O discounts					
P/O < 40% percentile	−1.76	0	1	0	10	3	0.40	2	7	1	4	0
P/O < 20% percentile	−2.05	0	1	0	10	2	0.31	2	6	2	5	0
P/O < 10% percentile	−2.35	1	1	0	10	1	−0.37	2	3	1	8	0
P/O < 5% percentile	−5.60	0	1	0	10	4	−0.42	1	4	1	7	0
P/O < 1% percentile	−7.56	2	3	0	8	4	−0.61	1	3	1	8	0
Asymmetric Risk Premium: β > (<) 0 when FP < (>) 0	Panel B: Exchange rates against the Private ECU; $I_{b,t}$ indicates a positive premium											
	β for positive forward premia						β for negative forward premia					
FP > 0	−3.52	1	4	1	7	2	−0.92	1	2	0	9	1
same, without DEM	1.46	1	4	1	6	1	−0.80	1	2	0	8	0
Peso risk II: β should be worse when P/O discount is deep	Panel C: Exchange rates against the Official ECU; $I_{b,t}$ indicates deep discount											
	β for deep P/O discounts						β for small P/O discounts					
P/O < 40% percentile	−1.86	0	2	0	9	5	0.23	2	6	2	5	0
P/O < 20% percentile	−2.20	0	1	1	10	3	0.12	2	5	2	6	0
P/O < 10% percentile	−2.10	0	2	0	9	1	−0.36	2	5	1	6	0
P/O < 5% percentile	−4.85	0	1	0	10	2	−0.50	1	5	1	6	0
P/O < 1% percentile	−3.74	2	2	0	9	1	−0.65	1	3	1	8	0
Friction: β should be worse when |FP| is small	Panel D: Exchange rates against the Private ECU; $I_{b,t}$ indicates a large premium											
	β for small FP						β for large FP					
|FP| > 99% percentile	−0.78	1	3	1	8	0	0.17	3	5	2	6	0
|FP| > 95% percentile	−1.15	0	3	0	8	0	−0.34	2	5	1	6	0
|FP| > 90% percentile	−1.09	0	4	1	7	0	−0.49	1	4	1	7	0
|FP| > 80% percentile	−1.03	1	4	0	7	1	−0.59	1	2	1	9	0
|FP| > 60% percentile	−1.59	0	4	0	7	1	−0.78	1	3	1	8	0

Key to Table 2. Panels A to D each summarize results from 11 regressions, one for each base currency b (BEF, DEM, NLG, PTE, ITL, IEP, GRD, DKK, ESP, GBP, FRF). In Panels A, B, and D the regressand , $s_{b,t,t+\Delta}$, is the three–month percentage spot-rate change for the private ECU while in Panel C the regressand is the change in the value of the Official ECU. The regressor is the corresponding beginning–of–period three-month forward premium, $FP_{b,t,t+\Delta}$. The data are weekly observations from October 1989 to February 1998, and in the regressions we follow Hansen and Hodrick (1980) to take into account the overlap in the holding periods. In each panel, the observations are dichotomized into two subsets: in Panels A and C on the basis of the percentage discount in the private ECU relative to the official one ("P/O"), in Panel D on the basis of the absolute size of the forward premium (|FP|), and in Panel B on the basis of the sign of the forward premium. Each set has its own coefficient for the forward premium. In each of the panels, β_{b0} should be closer to unity than β_{b1}, and upper lines should do better than lower ones.

Table 3. Tests of Interactions Between the Competing Hypotheses on the Forward Bias

Summary statistics on $s_{b,t,t+\Delta} = \alpha_b + \beta_{b11}[I_{F,b,t} \times I_{P,b,t} \times FP_{b,t,t+\Delta}] + \beta_{b10}[I_{F,b,t} \times (1-I_{P,b,t}) \times FP_{b,t,t+\Delta}]$

$+ \beta_{b01}[(1-I_{F,b,t}) \times I_{b,t} \times FP_{b,t,t+\Delta}] + \beta_{b00}[(1-I_{F,b,t}) \times (1-I_{b,t}) \times FP_{b,t,t+\Delta}] + \eta_{b,t,t+\Delta}$

	Panel A: interaction of P/O discounts and sign of forward premium: regression coeficients											
	deep P/O discounts (P/O < 40% percentile)						small P/O discounts (P/O ≥ 40% percentile)					
	positive FP			negative FP			positive FP			negative FP		
criterion for $I_{b,t} = 1$	mean β_{b1}	# >0 (signif)	# <0 (signif)	mean β_{b0}	# >0 (signif)	# <0 (signif)	mean β_{b1}	# >0 (signif)	# <0 (signif)	mean β_{b0}	# >0 (signif)	# <0 (signif)
FP > 0	–1.76	4(1)	6(0)	–2.04	2(0)	4(1)	–3.78	5(1)	3(2)	0.86	4(0)	3(0)
same, without DEM	–1.57	4(1)	5(0)	–2.04	2(0)	4(1)	–0.43	5(1)	2(1)	0.90	4(0)	2(0)

Panel B: interaction of P/O discounts and size of forward premium

	Panel B1: Spearman Rank Correlation Coefficient between the P/O premium and the absolute size of the forward premium		
	average SRC	positive	negative
Total sample	–0.04	5	6
deep P/O discounts (< 40 %ile)	–0.29	3	8
small P/O discounts (> 40 %ile)	0.04	6	5

	Panel B2: COF regression coefficients											
	deep P/O discounts (P/O < 40% percentile)						small P/O discounts (P/O ≥ 40% percentile)					
	large FP			small FP			large FP			small FP		
criterion for $I_{b,t} = 1$	mean β_{b11}	# >0 (signif)	# <0 (signif)	mean β_{b11}	# >0 (signif)	# <0 (signif)	mean β_{b11}	# >0 (signif)	# <0 (signif)	mean β_{b11}	# >0 (signif)	# <0 (signif)
• |FP| > 95% percentile	–0.72	4(0)	7(0)	–2.20	2(0)	9(4)	0.61	5(2)	6(0)	0.71	8(3)	3(0)
• |FP| > 90% percentile	–0.91	2(0)	9(0)	–2.37	2(0)	9(2)	0.61	6(2)	5(0)	0.91	8(1)	3(0)
• |FP| > 80% percentile	–1.05	2(0)	9(0)	–2.55	2(0)	9(2)	0.58	5(2)	6(0)	0.67	8(1)	3(0)
• |FP| > 60% percentile	–1.83	1(0)	10(3)	–2.50	3(0)	8(3)	0.53	7(2)	4(0)	–0.44	5(0)	6(0)

Key to Table 3. Each line in Panels A and B2 summarizes the results from 11 regressions, one for each base currency b (BEF, DEM, NLG, PTE, ITL, IEP, GRD, DKK, ESP, GBP, FRF). The regressand , $s_{b,t,t+\Delta}$, is the three–month percentage spot-rate change for the private ECU ,the regressor is the corresponding beginning–of–period three-month forward premium, $FP_{b,t,t+\Delta}$. The data are weekly observations from October 1989 to February 1998, and in the regressions we follow Hansen and Hodrick (1980) to take into account the overlap in the holding periods. In each panel, the observations are first grouped into two main subsamples on the basis of the percentage discount in the private ECU relative to the official one ("P/O"). Each main subsample is then further split up: in Panel A on the basis of the sign of the forward premium, and in Panel B2 on the basis of the absolute size of the forward premium. Each cell has its own coefficient for the forward premium. Panel B1, lastly, shows the average Spearman rank correlation coefficient between the P/O premium and the absolute size of the forward premium, in the total sample as well as in each of the main subsamples.

Figure 1. Time-Series Plot of the P/O Discount

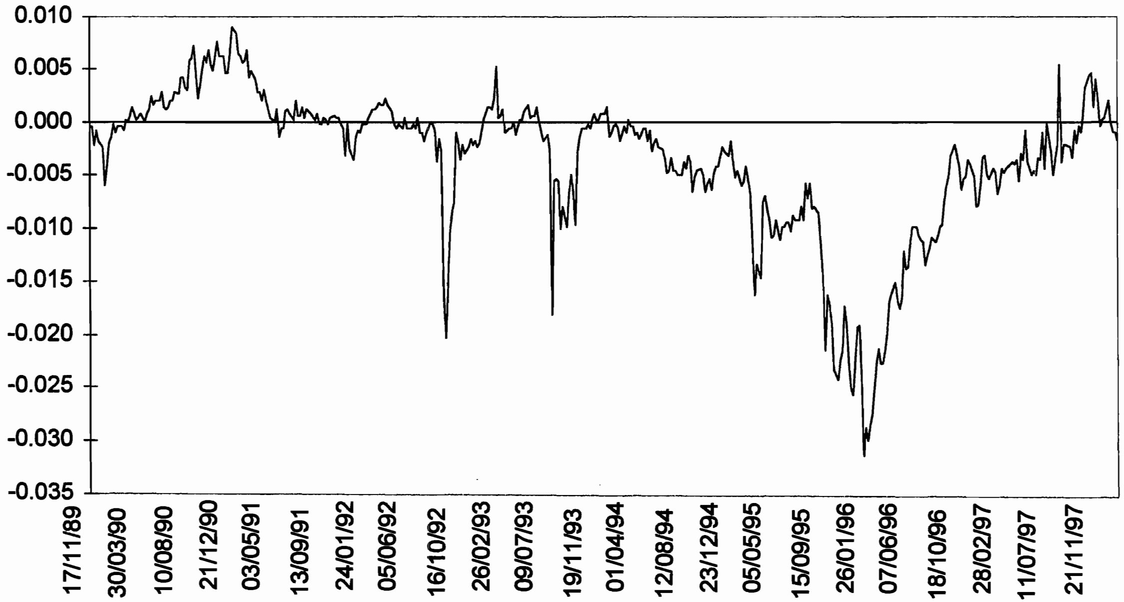

Key to Figure 1. The P/O premium is defined as the log of the ratio of the spot values of the private and official ECU. The data are weekly, Oct 17, 1989 to Feb 15, 1998.

Figure 2. Time-Series Plot of the Interest-Rate Deviation, Private versus Official

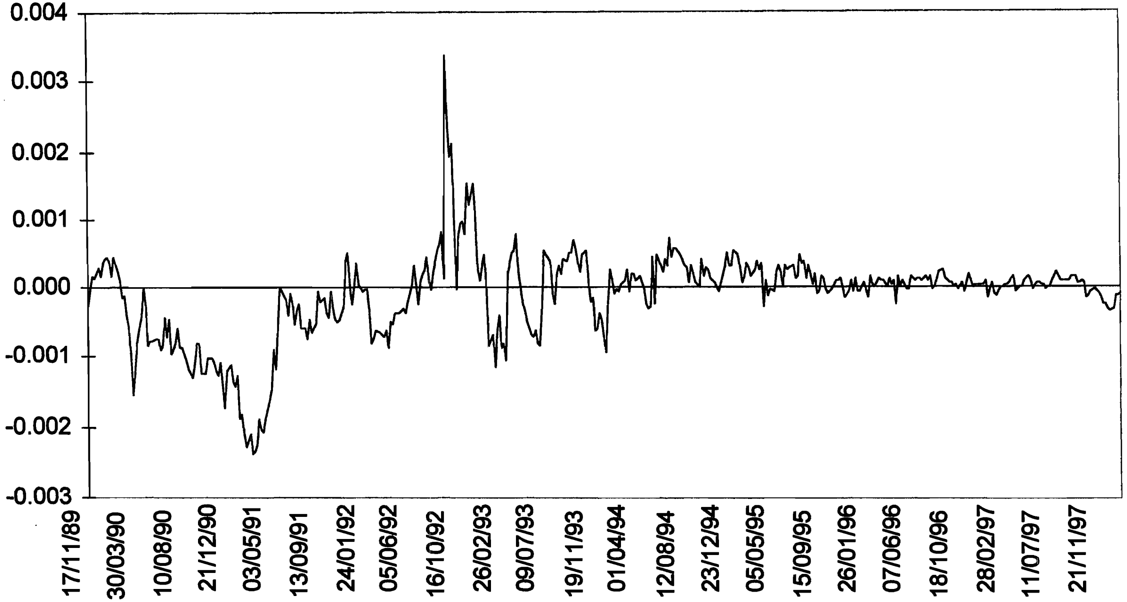

Key to Figure 2. The interest-rate deviation is defined as the log of the ratio of the gross three-month risk-free rates of return on private and official ECU. The gross three-month rate of return on official ECU is the harmonic weighted average of the gross returns on the component currencies (see Appendix). The data are weekly, Oct 17, 1989 to Feb 15, 1998.