

Using DFA for Modelling the Impact of Foreign Exchange Risks on Reinsurance Decisions

Peter Blum^{1,2}, Michel Dacorogna², Paul Embrechts¹, Tony Neghaiwi², Hubert Niggli²

Abstract

The fluctuations of the foreign exchange (FX) rates are a source of additional risk, but also an opportunity for further profits for internationally operating reinsurers. A DFA model that includes FX rates can be a means for measuring the potential impact of FX rate fluctuations on portfolios of ceded reinsurance and internationally invested assets. Moreover, such a DFA model can also be a means for testing different strategies for managing FX risk.

We start by surveying the empirical properties of FX rate data and introduce some simple approaches for generating FX scenarios. We then study the different ways in which reinsurers are exposed to FX risk, and possible strategies for managing these risks. The emphasis is on so-called on-balance sheet currency hedging techniques, i.e. on offsetting FX risks on reinsurance contracts by investing appropriate amounts of money in the respective foreign currency.

The theoretical reasonings are accompanied by a DFA study on an international loss portfolio. This study shows that FX fluctuations can have an impact on the reinsurer's business, and that the optimal strategy may not be a full hedge depending on the type of business and on the economic circumstances.

Acknowledgements

The authors wish to acknowledge the useful advice of Sholom Feldblum at the onset of this work and the valuable comments and reviews of their referee Gary Blumsohn on earlier versions of the manuscript.

¹ DEPARTMENT OF MATHEMATICS, ETH ZURICH, CH-8092 ZURICH, SWITZERLAND,

E-mail addresses: {blum|embrechts}@math.ethz.ch

² ZURICH INSURANCE COMPANY, REINSURANCE, PO BOX, CH-8022 ZURICH, SWITZERLAND,

E-mail addresses:

{peter.blum|michel.dacorogna|tony.neghaiwi|hubert.niggli}@zurichre.com

1. Introduction

Reinsurance business is becoming increasingly international. Therefore, the impact of foreign exchange risk (FX risk) has to be dealt with by both cedant and reinsurer when structuring reinsurance programs.

The standard approach to dealing with FX risks emanating from (re)insurance contracts denominated in foreign currencies is to take sufficient asset positions in the same currency and pay the liabilities from these positions, thus avoiding cash flows across the currency border at the time the losses occur (so-called currency hedging). Servicing losses at lowest possible cost is, however, only one part of the insurance business. The other part is that the company's assets should produce good investment returns. While the approach of currency hedging definitely reduces cash flows that are exposed to FX risk, it also leads to possibly suboptimal investment returns by binding assets in suboptimal positions. It should also be taken into account that FX risk has not only a downside, but is itself an opportunity for investment benefits which should not be left unused given that one has to take exposure in foreign currencies. For international financial asset management, corresponding results are given by Froot (1993) and Levich and Thomas (1993), and we will make some considerations for the insurance-specific context in Section 5.

Hence, in the context of integrated risk management, the management of FX risks should not only be considered from the point of view of servicing international liabilities, but also from the point of view of maximizing the return of an international investment portfolio. The big challenge is to find an optimal tradeoff between these two points of view, and DFA can be a tool for achieving this goal.

Therefore, a modern DFA system that is able to model FX risks has many practical applications for today's global players. Taking into account FX risks when doing DFA allows for more accurate and innovative structuring and pricing of ceded reinsurance arrangements, which is in the interest of both cedant and reinsurer.

This paper is accompanied by a practical DFA study in order to underline the statements being made. The setup of this study is explained in Section 2, and the study will be revisited in each of the subsequent sections. Section 3 investigates the properties of foreign exchange rates and related variables and proposes two simple approaches for scenario generation. Section 4 presents the ways in which a (re)insurance company is exposed to FX risk and gives indications how these exposures can be modelled in DFA, and Section 5 surveys methods for FX risk management. Section 6 presents the application of the previously introduced methods in a DFA study for structuring reinsurance, and Section 7 states some conclusions and directions for further research.

2. The example

We introduce here the setup of a multi-currency (re)insurance deal that we will use throughout the whole paper to illustrate our reasoning. Notice that this is – in principle – a real case, but the number of involved currencies has been reduced (for simplicity's sake), and the figures are modified such that the customer cannot be identified.

The insurance company (cedant) has its Head Office in the United States and consolidates and reports its business in USD. The lines of business (LOB) involved in our deal are, however, situated in other countries and denominated in the respective currency: LOB1 is in Switzerland (denominated in CHF), LOB2 is in the United Kingdom (denominated in GBP), and LOB3 is in Japan (denominated in JPY). Therefore, each transaction of funds between the Head Office and the foreign LOBs is subject to FX risk. Moreover, if the Head Office values the three foreign LOBs, all their figures have to be translated back into USD.

The three LOBs are runoff business. Each LOB has an initial loss reserve for time 0 and a given stochastic payout pattern, and all claims payments are to be made in the currency of the respective LOB. These loss reserves are typical P/C loss reserves in that they bear risk in the timing as well as in the final amount of the payments. Therefore, from the point of view of the Head Office, we are faced with a superposition of risks on timing and amount (inherent to each LOB) and FX risk.

The Head Office wants to protect its (USD-denominated) results by taking out a stop-loss contract. There are different degrees of freedom in the design of this contract:

- One contract for the aggregate portfolio or single contracts for each LOB.
- Different retentions and limits.
- Currency in which to denominate the contract(s) and parameters; use of fixed exchange rates in the contracts yes/no.
- What is the additional exposure of the reinsurer in its layer due to FX volatility? This, in turn, will influence the price of the contract.
- Hedging issues for insurer and reinsurer: hedging yes /no, partial hedging, static or dynamic hedging, single currency or cross currency hedging.

The following figure shows the three basic loss processes involved in the example (without influence of any FX rates):

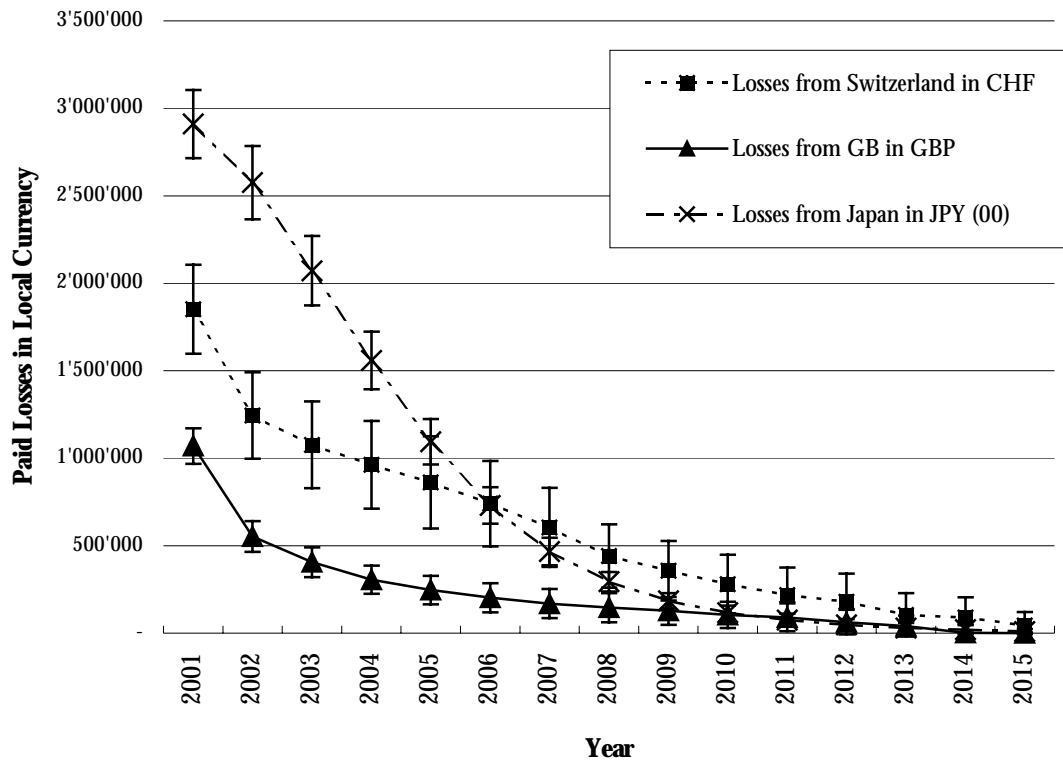


Figure 1: Losses per year and per LOB denominated in local currency. The intervals around the observations are plus/minus one standard deviation.

3. Scenario generation

3.1 Requirements

Using DFA for assessing FX risks requires the generation of FX scenarios. These scenarios must not be generated in isolation, but they must also reflect the relations between FX rates and other economic variables used in DFA studies, in particular interest rates and inflations in the involved countries. Moreover, our aim is not to optimize our generator so as to produce the most accurate level or interval forecasts, but credible and coherent scenarios for the time up to the defined horizon. The generator must be such that it can be calibrated with low amounts of data: The usual time resolution in DFA is one year, which means that – given the historical facts exhibited below – only about 25 data points are available. Even if we use monthly data, not more than about 300 data points will be available. The generator should be transparent in that each parameter has a practical meaning, such that manual adjustments by the user are possible: either for correcting the calibration or in order to bring in expectations of the future behavior that are not reflected by historical data. A word of warning: It is not our ambition to develop here a highly sophisticated econometric model of the FX markets. We only have the aim to propose a simple model that reflects the most important features of the data while being tractable for the average practitioner.

3.2 Definition of Terms

Unless otherwise stated, we will always use the *spot exchange rate* with respect to the US Dollar (USD), i.e. we denote by S_t^{XXX} the USD price for one unit of currency XXX at time t . The *cross rate*, i.e. the price of one unit of currency YYY in units of currency XXX, can then be calculated as $S_t^{XXX/YYY} = S_t^{YYY} / S_t^{XXX}$. For the analytical treatment, we will use the logarithm of the spot exchange rate, i.e. $s_t^{XXX} = \log(S_t^{XXX})$. By I_t^{XXX} and R_t^{XXX} we denote the rate of inflation and the risk free short-term interest rate, respectively, for country (currency) XXX. For analytical treatment we will also use some sort of logarithmic transformation $i_t^{XXX} = \log(1 + I_t^{XXX} / 100)$ and $r_t^{XXX} = \log(1 + R_t^{XXX} / 100)$ that put these percentage rates on equal footing with the FX rates, which are essentially prices. By \mathbf{x}_t we denote the value of a multivariate time series at time t . Rather than in the level \mathbf{x}_t , we may be interested in analyzing the changes of the time series in a time interval, i.e. $\Delta \mathbf{x}_t = \mathbf{x}_t - \mathbf{x}_{t-1}$. For the analysis of one currency pair (always with the USD as the reference), we will consider the time series $\mathbf{x}_t = (x_t^1, \dots, x_t^5)^T = (r_t^{USD}, i_t^{USD}, r_t^{XXX}, i_t^{XXX}, s_t^{XXX})^T$. The extension to more currencies is obvious.

3.3 Properties and stylized facts of FX data

FX rates can be governed by different types of regimes: A currency is *pegged* if its exchange rate with respect to some reference currency is fixed. A currency is *semi-pegged* if its exchange rate with respect to a reference currency is only allowed to float within a narrow band. Maintenance of

pegged and semi-pegged relations requires suitable interventions from the involved governments. If no formal restrictions are imposed on the exchange rate, then it is called *free-floating*. From the end of World War II up to 1971, the so-called Bretton-Woods system (semi-) pegged all currencies to the US Dollar within a band of 1%. Free floating of the most traded currencies only emerged after the breakdown of the Bretton-Woods system in 1971. Therefore, FX data from before 1972 should not be used for the calibration of scenario generators that model free-floating currencies. Pegging and semi-pegging relations still exist nowadays for some (mostly less traded) currencies. A special case is the European Monetary Union (EMU): Whereas the Euro is allowed to float freely against all outside currencies, the domestic currencies of the members are pegged to the Euro. This state will disappear by the end of 2001, when the Euro will fully replace all domestic currencies of EMU countries (e.g. the Deutsche Mark). See Luca (2000) for more information on historical and political backgrounds, as well as for a thorough treatment of the modern FX markets.

There are two basic ways to analyze and model the interplay of FX rates with the other economic variables involved: fundamental (economic) and technical (statistical) analysis. Economists have explored relations of FX rates with a high number of macro-economic variables (e.g. money supply, export balance, inflation, etc.); see MacDonald (2000) or Clostermann and Schnatz (2000) for easy-to-read treatments. The ubiquitous concept in these treatments is the *Purchasing Power Parity* (PPP), which states that – under the assumption of efficient markets and without transportation costs – goods should have effectively the same price in the two countries, i.e. $P^{YYY} = S^{XXX/YYY} \cdot P^{XXX}$, where P denotes the prices in the respective currencies. While this relation is clearly too strict to hold in practice, there exist more elaborate versions. One that takes into account those variables that we need in DFA anyway is:

$$s_t^{XXX/YYY} = \alpha_0 + \alpha_1(i_t^{XXX} - i_t^{YYY}) + \alpha_2(r_t^{XXX} - r_t^{YYY}) + I_t \quad (1)$$

where α_0 , α_1 and α_2 are constants, and I is a stationary stochastic process (e.g. auto-regressive). This equation relates the FX rate to the inflation difference and to the interest rate difference, which can be seen as two key incentives for taking or leaving positions in some currency. An alternative is to relate the FX rate to the real rates of return in the two countries, i.e.

$$s_t^{XXX/YYY} = \alpha_0 + \alpha_1(r_t^{XXX} - i_t^{XXX}) + \alpha_2(r_t^{YYY} - i_t^{YYY}) + I_t. \quad (2)$$

Clearly, possible investment benefits are not the only cause governing the demand (and hence the price) for a foreign currency. Therefore, we must not expect the above relations to explain all of the dynamics of the FX rates. Relations of the above type, which postulate that some linear combination of components of an otherwise non-stationary time series are stationary, are called *cointegration relations*, and they are very popular in statistics and econometrics, see Clements and Hendry (1998) for more information.

In statistical analysis we investigate the multivariate time series (x_t) , of variables of interest by means of statistical data analysis techniques, thus trying to explore properties and relations in the data with the final aim of finding and calibrating a model that reflects the actual properties of the data. Here is a short survey of relevant properties. The results were obtained from investigations of the authors with yearly and monthly observations of several countries with free-floating exchange rates (USA, UK, Japan, Switzerland) over the past two decades, and they go in line with the findings of MacDonald (2000), Clostermann and Schnatz (2000), and Clements and Hendry (1998). See Figure 1, graphs (a) to (f), for some illustrative examples; tables and charts of all involved data series can be found in Appendix A.

- The multivariate series x_t as a whole shows clear signs of non-stationarity (i.e. drifts and non-constant volatility, see (a)), whereas the differenced series Δx_t appears stationary. This is not fully obvious from Graph (b), but statistical tests on monthly and higher-frequency observations underpin the hypothesis of stationarity in differences.
- The values of interest rates, FX rates and inflation show clear dependence on their own past-year values (autocorrelation) as well as on present-year and past-year values of other variables (cross-correlations). However, the cross-correlation relations are different from country to country. As an example, Swiss inflation is correlated with the USD/CHF exchange rate, whereas US inflation is largely unaffected by any single FX rate. Differences in one year are, moreover, often related to the levels of previous years, thus indicating level-dependent volatility and – if the relation has negative sign – mean reversion.
- The above-mentioned cointegration relations (1) or (2) can be observed in the data after removing autoregressive components. Whether (1) or (2) is more significant depends on the currency pair. See (c) for the USD/JPY rate: The solid line is the prediction based on formula (2), the dotted line represents the true values; the residuals (dotted line minus solid line) have zero mean, i.e. the model is unbiased. In monthly data neither (1) nor (2) are very significant. This goes in line with the above-mentioned econometric sources, which state that these relations only hold over the longer term³, but large deviations are possible in the shorter term. Our observations with real data suggest that one year is already a sufficiently long time span for the involved variables to revert to the equilibrium relations, at least for freely traded currency pairs and liquid markets.
- Moreover, on yearly data, the inflation could be shown to be very significantly correlated with the interest rate of the respective country. The model $i_t^{xxx} = \beta_0 + \beta_1 r_t^{xxx} + I_t$ was found to fit to the data very well. Graph (d) shows predicted and actual inflation for the US in the same manner as described above. This does, however, not mean that we postulate that

³ Notice: In the econometric literature, time horizons of one year and longer are usually referred to as “long term”, whereas one day up to one quarter are referred to as “short term”.

the inflation causally depends on the interest rate; the real causation is likely to be bi-directional. For modelling purposes, however, the above formula is useful, and its use is justified by the fact that it fits well to the data. The converse modelling approach (interest rates as a function of the inflation) can also be taken, see e.g. Daykin et al. (1994).

- Most of the time, the log-transformed variables show Gaussian behaviour. There exist, however, a few extreme moves in the sample that cannot be reconciled with the hypothesis of Gaussianity. Graphs (e) and (f) show two QQ-normal plots⁴ of residuals: The residuals of Swiss inflation (Graph (e)) show clearly non-Gaussian behaviour, whereas the residuals of the US inflation (Graph (f)) look clearly Gaussian. Therefore, even on this very high level of aggregation, we must be aware of extreme, non-Gaussian movements.

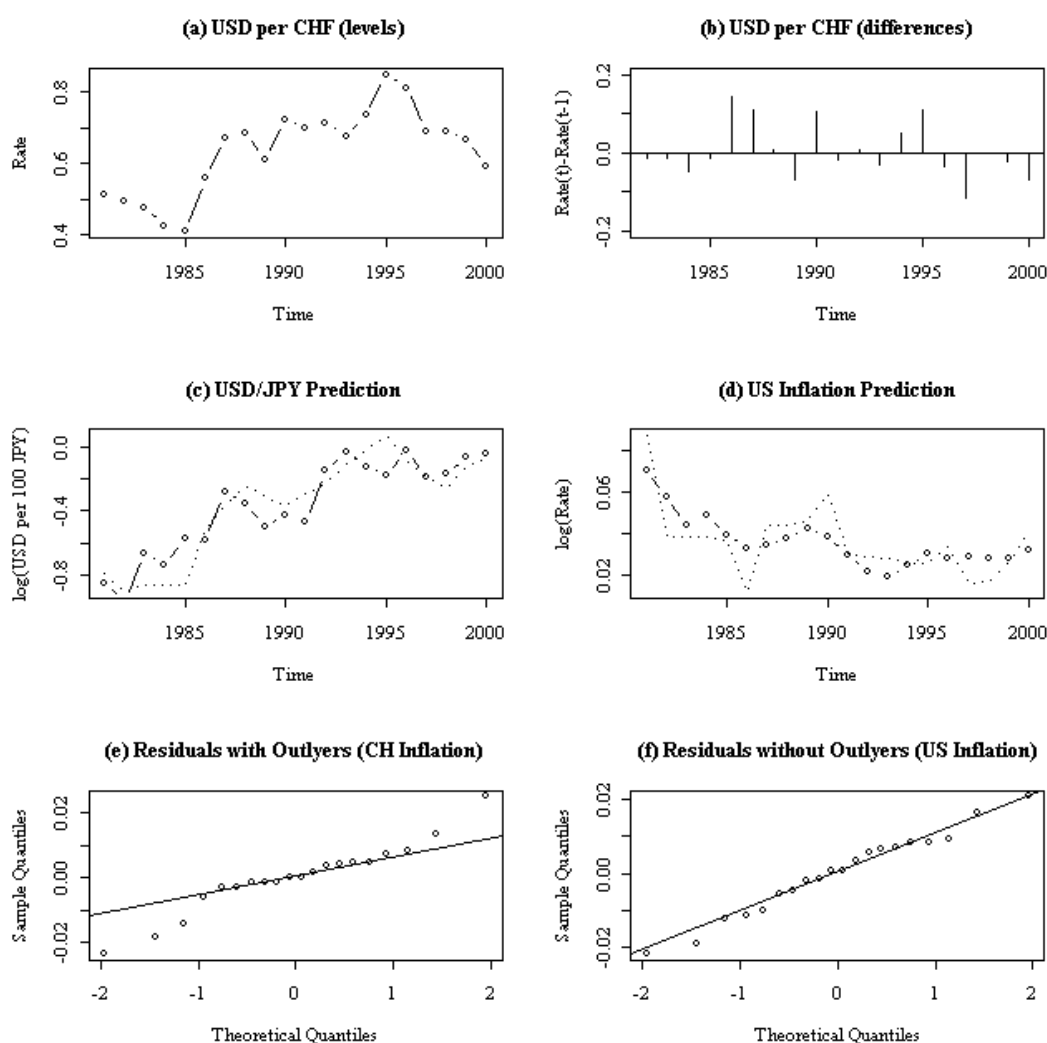


Figure 2: Some examples from data analysis (*Source: Datastream*)

⁴ In a Quantile-Quantile-normal plot, the sample quantiles of the observed residuals are plotted against the theoretical quantiles of the Normal distribution. If the points are tightly grouped around the diagonal line (which represents the theoretical Normal distribution), then the sample can be assumed to have Gaussian distribution. See any textbook on statistical data analysis for further information.

- If one looks at monthly instead of yearly data, the following can be observed: Clearly the autoregressive relations extend to several time steps backwards. Moreover, particularly the inflation shows a strong intra-yearly seasonality, which can be overcome by using seasonally adjusted data (available for most countries). FX rates and interest rates have negligible seasonal effects. Finally, there are relatively more extreme (non-Gaussian) movements due to the reduced smoothing caused by the lower time aggregation.

Based on the summarized statistical analysis, it is possible to formulate some general rules, but there are also differences between the currencies. When setting up and calibrating a FX scenario generator, it is essential to do dedicated statistical analysis on the involved currencies, and the above explanations may serve as a guideline.

3.3 The generator and its calibration

As one possibility for modelling FX rates together with their constituents relevant for DFA (i.e. interest rates and inflation) on a yearly basis we propose a *bottom up modular approach*. This approach permits to reuse already-existing models for interest rates and inflation in one country. The outputs of these models are then used to compute the exchange rates by using a cointegration relation, i.e.

$$s_t^{XXX/YYY} = \alpha_0^{XY} + \alpha_1^{XY} r_t^{XXX} + \alpha_2^{XY} i_t^{XXX} + \alpha_3^{XY} r_t^{YYY} + \alpha_4^{XY} i_t^{YYY} + \varepsilon_t^{XY}$$

where:

s, r, i : logarithmic FX, interest, and inflation rates as introduced above,

$\alpha_0^{XY}, \dots, \alpha_4^{XY}$: constant regression parameters specific to each currency pair XXX/YYY,

ε_t^{XY} : an iid series of random variables with zero mean,

either $N(0, \sigma_{XY}^2)$ if residuals are Gaussian,

or with Student's t distribution⁵ if residuals show non-Gaussian extremes.

If we let $\alpha_1^{XY} = -\alpha_3^{XY}$ and $\alpha_2^{XY} = -\alpha_4^{XY}$ this model is equivalent to the relation (1). Similarly, one can also achieve equivalence with relation (2). In both cases we have a reduction of the number of parameters to be estimated, which is an advantage in view of the low amounts of data available for estimation.

In the light of our findings in the previous section, the inflation for a given currency can be modelled by a simple linear regression on the respective short-term interest rate:

$$i_t^{XXX} = \beta_0^X + \beta_1^X r_t^{XXX} + \varepsilon_t^{iX} \quad \text{and} \quad i_t^{YYY} = \beta_0^Y + \beta_1^Y r_t^{YYY} + \varepsilon_t^{iY} \quad (3)$$

where:

r, i : logarithmic interest and inflation rates as above,

$\beta_0^X, \dots, \beta_1^Y$: constant regression parameters specific to currencies XXX and YYY,

$\varepsilon_t^X, \varepsilon_t^Y$: iid series of random variables with zero mean as above.

⁵ See Embrechts et al. (1997) for more information on the modelling of extremal events.

For the short-term interest rates, we take here the well-known Cox-Ingersoll-Ross (CIR) model, which is described in Chan et al. (1992), and which can be extended to generate a full yield curve without extra effort for calibration:

$$\Delta r_t^{XXX} = a^X (b^X - r_{t-1}^{XXX}) + s^X \sqrt{r_{t-1}^{XXX}} \varepsilon_t^{rX} \quad \text{and} \quad \Delta r_t^{YYY} = a^Y (b^Y - r_{t-1}^{YYY}) + s^Y \sqrt{r_{t-1}^{YYY}} \varepsilon_t^{rY}$$

where:

r : logarithmic interest rates as above

a, b, s : parameters of CIR model as specified in Chan et al. (1992)

$\varepsilon_t^X, \varepsilon_t^Y$: iid series of Gaussian random variables with zero mean and unit variance

In order to account for cross-country dependences, we have to introduce some dependence structure between the innovations of the two interest rate processes:

- either by letting $\varepsilon_t^r := (\varepsilon_t^{rX}, \varepsilon_t^{rY})^T$ be an iid series of bivariate Gaussian random vectors with zero mean and covariance matrix Λ ,
- or by linking the random variables ε_t^{rX} and ε_t^{rY} via a so-called copula. The concept of copulas is described in Embrechts et al. (1999). This is an advanced concept that allows to account for stronger dependence in the tails of the involved random variables.

The model introduced here for the inflation and interest rates is the same as used in Dynamo, see D'Arcy et al. (1997), but any other model for short-term interest rates and inflation can be used instead. It must only be extended by formula (3) for the exchange rate and by a dependence structure for the innovations of the basis variables.

Calibration of this model is straightforward:

- i. Calibrate the interest rate model for each currency (for CIR: see Chan et al. (1992)).
- ii. Investigate the joint distribution of the residuals of the interest rate processes. If there are signs for stronger correlation in the tails, use a copula, otherwise take the estimated covariance matrix.
- iii. For each currency, estimate the regression coefficients for the inflation with respect to the interest rate. If the residuals show non-Gaussian outliers, fit a heavier-tailed distribution.
- iv. For each currency pair, estimate the coefficients of formula (3) by linear regression. If the residuals show non-Gaussian behaviour, fit a heavier-tailed distribution.

The actual parameters estimated for the setup of our example can be found in Appendix B.

In the light of the findings of the data analysis in the previous section, this approach is well suited for modelling on a yearly basis. Moreover, it has the advantage that already-used models for interest rates and inflation can be reused, and that the parameters have clear practical interpretations. On a quarterly or monthly basis, FX rates also depend on past values of inflation and interest rates. A careful data analysis along the lines presented in the previous section should be done in this case before selecting and calibrating a model. A possible closed-form alternative

would be the use of a purely statistical multivariate time series model for the process $\Delta \mathbf{x}_t$, e.g. the *Vector Autoregressive (VAR)* model:

$$\Delta \mathbf{x}_t = \boldsymbol{\tau}_t + \sum_{k=1}^p \boldsymbol{\Gamma}_k \Delta \mathbf{x}_{t-k} + \boldsymbol{\varepsilon}_t$$

where:

$\Delta \mathbf{x}$: as introduced above,

$\boldsymbol{\tau}_t$: vector of deterministic drift terms,

$\boldsymbol{\Gamma}_k$: matrices of autoregression coefficients,

p : maximum order of autoregression,

$\boldsymbol{\varepsilon}_t$: iid series of Gaussian random vectors with zero mean and covariance matrix Λ .

The main advantage of this closed-form model is that it allows more easily to simulate conditional scenarios where the values of a part of the variables are fixed (“what-if analysis”). It is also more open in that it also allows dependences that are not economically derived. Disadvantages include the (relatively) high number of parameters and the simple dependence structure (only linear dependence on past values) that is not able to render certain characteristics of the data such as the level-dependent volatility of interest rates. When fitted to yearly data, the VAR(1) model was unbiased, but the goodness-of-fit was not as good as for the modular model.

The models were tested for quality by doing analysis of the residuals, i.e. the empirical values of the various $\boldsymbol{\varepsilon}$'s in the formulas. The model is *unbiased* if the mean of the residuals is zero, and if the time series of the residuals shows no signs of serial correlation. This means that the model captures the behaviour of the historical data without systematical deviations. The model has a high goodness-of-fit if the variance of the residuals is considerably smaller than the total variance of the historical data, and if the empirical distribution of the residuals can be reconciled with the one given by the model (e.g. Gaussian). This means that a relevant part of the variability of the data is explained by the model and not by the residual noise. If one uses a statistics software (in our case R, the freeware version of S-Plus), one can also obtain test statistics that give more elaborate indications for the goodness-of-fit. The ultimate means for measuring model quality would be to use a statistical information criterion, the most classical one being AIC (Akaike's Information Criterion). The problem with these criteria is, however, that they are relatively difficult to compute in practice.

When fitted to yearly data, both models (i.e. the bottom-up modular and VAR(1)) turned out to be significantly unbiased. The goodness-of-fit was better for the bottom-up modular model than for the VAR(1), due to the lower number of parameters to be estimated with the same amount of data. On the absolute scale the goodness-of-fit is not very high, but we have at least for all variables that the major part of the variability is explained by the model and not by the residual variance. The practical meaning of this statement is that simulated scenarios for the future will behave according to (roughly) the same stochastic discipline as the historical data, i.e. our model

is a good prediction of the future provided that the future behaviour of the modelled economies is fundamentally the same as in the preceding years. If expectations for the future are different, the estimated coefficients must be modified accordingly, which is easy due to the simplicity of both models.

We do not claim that our proposed models are the best possible generators for FX scenarios in DFA. But we have chosen them because they are able to reflect the relevant statistical and economic properties of the respective multivariate time series while being relatively simple, transparent, and easy to calibrate. Another alternative – if available – would be to use scenarios generated externally by one of the commercially available economic scenario generators, such as LongRun from RiskMetrics, see Kim et al. (1999), or mark-to-future from Algorithmics, see Dembo et al. (2000). For a general survey of models for strategic long-term financial risks see Kaufmann and Patie (2000) and the references therein.

To conclude, we show here the results of a Monte Carlo evaluation of the scenario generator for the bottom-up modular model. The study consisted of generating 1000 realizations of the USD/CHF exchange rate and its constituents for the time interval from 1991 up to 2000 by using the bottom-up modular model calibrated with data from 1980 to 2000. This is called an *in-sample* evaluation. In general, *out-of-sample* evaluations (i.e. with distinct time intervals for calibration data and reference data) would be preferable, but they were not done here because of lack of data. For each simulated year, mean and standard deviation of the 1000 realizations for each simulated variable were generated. The results are shown below: the points represent the true values of the respective variable, the solid line is the time series of the mean values of the forecasts, the dashed lines represent the interval mean \pm one standard deviation for each time step, and the dotted lines represent the interval mean \pm two standard deviations. A word of warning: the power of this evaluation is very limited; the graphs only give a hint on whether the scenarios are able to capture the behaviour of the real *past* series *on the average*. Whether one wants to adjust the parameters or not depends on the expectations for the future. The true values of the time series are in almost all cases comprised within one standard deviation from the mean, and always within two standard deviations. Swiss interest rates (and also inflation) are systematically overpredicted for the second half of the nineties. This is due to the fact that these rates were much higher in the eighties than in the nineties, with the respective influence on the parameters of the model. The average interest rate predictions of the CIR model (and hence also inflation and FX rate) show relatively strong mean reversion on the long run. This is not very realistic, but maybe the best guess for a distant future with the accordingly high uncertainty.

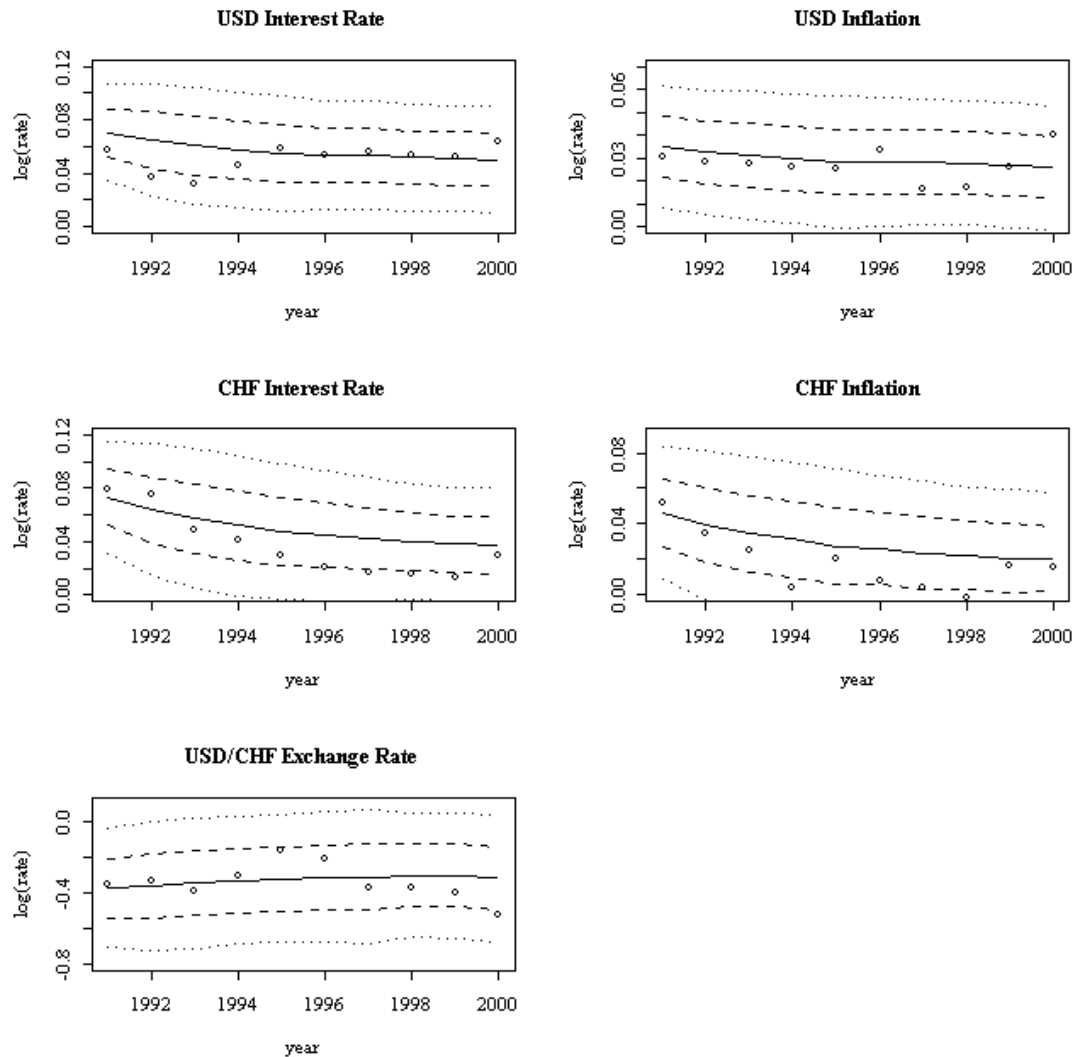


Figure 3: Results of Monte Carlo evaluation

3.4 Discussion of the limits of the modelling approach

We conclude this section by assessing two important problems related to the simulation of economic variables such as FX rates, interest rates, or inflation in the context of DFA. The first problem is the lack of sufficient amounts of data for the calibration of the model. With a time resolution of one year, we only have about 25 years of observations at hand. The model must therefore be kept very simple, i.e. with a low number of parameters to be estimated. Even then, the accuracy of the estimated parameters is not very high, but it was at least possible to obtain estimates that have no systematic bias with respect to the historical data. One way to achieve better-determined parameters is to use models and data on higher frequencies, e.g. monthly or even weekly. This requires, however, a very good understanding of the time aggregation properties of the involved models and the methods for their calibration. This approach is quite well understood for univariate models and on very high frequencies (intra-daily data), see Dacorogna et al. (2001). For multivariate models and long time horizons the situation is more difficult and usable results are sparse, see e.g. Kaufmann and Patie (2000). If models are kept

simple and parameters have clear practical meanings, then there is also the possibility to adjust the parameters “by hand”, e.g. based on insights or assumptions from other sources than statistical analysis of historical data. The approach of using simple but transparent models is often referred to as “Simplicity Postulate” or “Parsimony” in the statistical and econometric literature, see e.g. Clements and Hendry (1998).

Another problem is of more fundamental nature: When we have fitted a model to historical data, then the scenarios generated by this model have (roughly) the same stochastic behaviour as the historical data, which is not necessarily the best projection for the future. This means in practical terms:

- We implicitly assume that the future behaviour of the modelled economic variables is fundamentally the same as it used to be in that part of the past considered for calibration.
- The generated scenarios do, however, not account for fundamental changes or drifts in the regime governing the modelled variables, nor do they sufficiently account for hitherto unexperienced extremal events.
- The generated scenarios simulate the same risk as it emanated from the past behaviour of the respective variables. In particular: Scenarios account for “unusual” or “extreme” events to the same extent as such events occurred in the past. By using special models for the tails one can generate events with magnitudes and probabilities beyond what has been experienced. Refer to Embrechts et al. (1997) for the fundamentals and the discussion in Müller et al. (1998) about assessing extreme risks in the foreign exchange market.

See also Blumsohn (1999) for a presentation of this problem and possible solutions in a different context. The severity of this fundamental uncertainty depends on the time horizon of the study and also on non-statistical insights and expectations at the beginning of the period. Here are some possible ways to cope up with the problem of fundamental uncertainty. Their common characteristic is that they try to explore sensitivities rather than absolute levels of risk and return:

- Do several stochastic simulations with several different parameter sets for the scenario generator. These parameter sets can be based on statistical estimations, but also on manual adjustments reflecting non-statistical insights and assumptions.
- Identify those simulated scenarios that have led to extreme results (e.g. ruin) and explore their common characteristics, c.f. Dembo et al. (2000) for an application of this method in the context of finance.
- Complement the stochastic simulations with classical stress scenarios modelling particularly adverse courses of events.

4. Modelling FX risk exposure

4.1 FX risk in general

Whenever an amount of money A_t denominated in one currency YYY must be measured in terms of another currency XXX, this corresponds to multiplying the original amount by the currently prevailing FX rate between the two currencies: $S_t^{XXX/YYY} \cdot A_t$, i.e. the original volatility of A_t is superimposed by an additional volatility arising from the FX rate. In a typical DFA setup, there will usually be a large number of variables to be converted from one currency to another, and several currencies may be involved. There are correlations of different amount and sign between the FX rates and other stochastic variables: inflation and interest rates are correlated with the FX rate, inflation may influence loss experience, and interest rates influence investment returns. Hence, even in a relatively simple real case, it is already very difficult – if not impossible – to make a valid quantitative statement on the impact of FX volatility based only on analytical reasoning. See Loderer and Pichler (2000) for a survey of the difficulties that firms have in determining their actual FX risk exposure. Stochastic simulation, i.e. DFA, is one suitable means for resolving this issue. We can distinguish between three different types of FX risk exposure, i.e. *translation exposure*, *transaction exposure*, and *economic exposure*; see Luca (2000) for more details.

Translation exposure arises when foreign-denominated assets or liabilities are translated into the home currency for consolidation purposes at prevailing FX rates, e.g. for the annual financial statements of the company. If FX rates have changed since the last consolidation, this can lead to considerable changes in consolidated values of assets and liabilities. Translational changes are only nominal in the sense that no gains or losses are actually realized, since no assets or liabilities are liquidated. If translational fluctuations are high enough, they can nevertheless have an impact on the company's operations, e.g. through changes in taxation, loss of credit rating, reputational damage, or regulatory problems. Translational gains and losses can be identified by comparing financial statements with and without change in FX rates in the reporting period. They can also be a detector for possible future transaction exposures, in the case when the considered assets or liabilities must be liquidated. Accounting rules, however, offer means for mitigating the effect of translation exposure on the financial statements. The basic principle is that profits and losses from FX translations are assigned to an equity account and do not enter into income figures until the asset or liability is liquidated, but the details are rather complicated and vary between the different accounting standards; see FASB 52 (1981) for US Statutory and Chapter 22 of Bailey and Wild (2000) for IAS and GAAP. These rules make the implementation of translation exposure measurement in DFA highly complicated, therefore a treatment of translation exposure is beyond the scope of this paper.

Transaction exposure arises when funds are actually converted from one currency into another at prevailing FX rates. The respective gains and losses are no longer nominal, but they are realized

gains and losses and, therefore, fully affect the income of the company. Transactional gains and losses must be computed with respect to some reference FX rate. For a foreign-denominated investment, this would usually be the exchange rate at the time when the investment was made. Estimating transaction exposure and studying ways to anticipate it is the main aim of this paper.

Economic exposure is the impact that changes in FX rates can have on the competitive position of the company in the market. Economic exposure is very relevant for manufacturing companies, as the prices of their products are mainly made up by costs (labour, material, infrastructure). For (re)insurance companies, economic exposure is not so important, since expenses represent only a small part of the price of the policies, and the funds for covering risk can – at least partly – be moved from one country to another if necessary.

Besides the exposure to the FX rates themselves, there arises also additional exposure from the fact that there are different interest and inflation rates for each country. This may, for instance, lead to better investment returns in some country, but these returns may in turn be compensated by an adverse development of the respective FX rate. As we have seen previously, there exist correlations between interest rates, inflation and FX rates, but these correlations vary from case to case, and it is therefore not possible to make quantitative statements on the resulting overall exposure only based on analytical reasoning. Again, DFA is a good means to come to more insight in a given case.

4.2 Specific issues in insurance and reinsurance

Most of the concepts for FX risk management mainly apply to banks or other investment companies. An insurer or reinsurer is faced with some specific issues. On the asset side (re)insurers are often subject to regulatory constraints which prevent them from taking the positions that would best cover their needs (a Swiss insurer – for instance – is only allowed to hold a maximum of 20% of foreign-denominated bonds). Given the usually high proportions of bonds in their portfolios, (re)insurers are highly dependent on interest rates. Since (re)insurers are often only profitable because of investment returns, this issue must not be neglected.

The liability side of an insurance company is much more complicated than the one of a bank: There are often very large fluctuations in the loss development process, think e.g. of a Cat X/L that produces no claims in most years and large claims in a few years. More uncertainty arises from the fact that loss projections are often difficult due to insufficient data and due to the impossibility to foresee certain factors that affect loss development (e.g. changes in liability legislation). All these effects are increased by the fact that time horizons are often very long, e.g. up to 10 or 20 years. Hence, unlike banks, (re)insurers are faced with a high uncertainty as to amount and timing of their liabilities.

In a multi-currency setup, these issues concerning assets and liabilities become even more difficult to quantify since several inflation and interest rate regimes, as well as FX rates, and all the related correlations are involved.

An insurance or reinsurance company operating in a foreign market generally holds assets as well as liabilities in the respective currency, which gives the possibility of offsetting liabilities in one currency by assets in the same currency, thus reducing cross-currency cash flows and hence transaction exposure. This is one form of currency hedging and will be the subject of the next section. Notice however that the offsetting effect on the translation exposure is limited due to special accounting rules (see Bailey and Wild (2000) for details), even in case of foreign subsidiaries that are legal entities in the respective country. Moreover, in this case local regulations applying to the subsidiary may force a company to allocate its capital sub-optimally.

Let us now consider a truly international setup where a company runs business directly in a foreign country. Premiums are collected, claims are paid, and investments are made either in the company's home currency or in a foreign currency. Depending on the setup of international (re)insurance contracts, the risk emanating from the FX volatility is borne by different parties. We illustrate our reasoning by the following little example: Think of a UK cedant (calculating in GBP) that concludes reinsurance contracts with a US reinsurer (calculating in USD). At the present time $t = 0$ when the contracts are concluded, the exchange rate is at 1.5 USD per GBP. At the time $t = 1$ when the losses will occur, the exchange rate may be either at 1.2 or 1.5 or 1.8 USD per GBP, each with a certain probability. We consider the impact on the cash flows of cedant and reinsurer under different contract setups:

- i. Everything is quoted and settled in the currency of the cedant (here: GBP):

Rate:		1.2	1.5	1.8		1.2	1.5	1.8
Claim		Contract: 10 XS 10 denominated in GBP						
(GBP)		Cedant receives (GBP)				Reinsurer pays (USD)		
5		0	0	0		0	0	0
10		0	0	0		0	0	0
15		5	5	5		6	7.5	9
20		10	10	10		12	15	18

Whereas the cedant receives the same amounts irrespective of the exchange rate prevailing at $t = 1$, the liability of the reinsurer changes, i.e. the reinsurer bears the full currency risk. It is easily verified that the same applies also to a ground-up loss or a $x\%$ - quota share.

- ii. Everything is quoted and settled in the currency of the reinsurer (here: USD):

Rate:		1.2	1.5	1.8		1.2	1.5	1.8
Claim		Contract: 15 XS 15 denominated in USD						
(GBP)		Cedant receives (GBP)				Reinsurer pays (USD)		
5		0	0	0		0	0	0
10		0	0	1.67		0	0	3
15		2.5	5	6.67		3	7.5	12
20		7.5	10	8.33		9	15	15

In this case the cedant as well as the reinsurer bear a currency risk. In particular, whether or not the contract triggers is now also dependent on the exchange rate. It is easily verified that for a quota share, only the reinsurer would bear currency risk.

- iii. Fixed rates can be agreed at which all losses will be valued or paid. As an example, loss amounts are translated at 1.5 USD per GBP, payments however are made at currently prevailing rates, in which case the full exchange rate risk is borne by the cedant:

Rate:		1.2	1.5	1.8		1.2	1.5	1.8
Claim		Contract: 15 XS 15 denominated in USD, valued at 1.5 USD/GBP						
(GBP)		Cedant receives (GBP)				Reinsurer pays (USD)		
5		0	0	0		0	0	0
10		0	0	0		0	0	0
15		6.25	5	4.17		7.5	7.5	7.5
20		12.50	10	8.33		15	15	15

In any case there exists an additional risk due to the additional volatility of the FX rate, and this risk must be borne by someone at some price. A detailed consideration of the pricing of FX risk or the implications of FX volatility on calculations of risk adjusted capital is, however, beyond the scope of this paper.

FX rates can also enter indirectly into an insurance contract. This is the case when a contract is concluded in the common home currency of cedant and reinsurer, but for a claims process that depends on FX rates, e.g. a reinsurance contract for a primary insurer that writes business in a foreign currency.

4.3 Bringing FX risk exposure into a DFA model

A thorough treatment of how exactly to model FX risk exposure in a DFA model is not possible at this point, as DFA models are usually very complex and the attachment points for the FX rates differ from case to case. As an example, one may bear in mind the Dynamo model described by D'Arcy et al. (1997). We restrict ourselves here to state some general rules and principles mainly for transaction exposure. The implementation of a translation exposure model is different, depending on the accounting standard in use, and highly complex due to the complicated rules for redirecting gains and losses to special equity accounts. The modelling of economic exposure, which would have to go in line with the modelling of insurance and reinsurance business cycles, is not treated here.

The principle for modelling transaction exposure is simple: Each cash flow crossing the currency border (and only the cash flows, but not foreign-denominated positions that are consolidated into some balance sheet) must be multiplied by the respective FX rate prevailing at the time the cash flow occurs. The concrete implementation is highly dependent on the DFA model used and on the structure of the modelled company.

Application of the above-stated rules corresponds to doing a DFA study that simply takes into account FX fluctuations. One might, however, also be interested in measuring the portions of risk and return emanating specifically from FX fluctuations in some given period. This can be

achieved by generating two sets of DFA results for that period: one with constant FX rates equal to the initial values but otherwise stochastic inputs, and one with stochastic FX rates. A word of warning for this approach: If FX rates are kept constant, then the variables correlated with them (i.e. interest rates and inflation) must be simulated according to their conditional probability law given the fixed value of the FX rate. Otherwise, the generated scenarios may become implausible.

By doing several simulations with FX rates kept constant on some hypothetic levels and inflation and interest rates simulated according to the conditional probability given the FX rates, one can explore the sensitivity of the results against FX rate levels under otherwise equal circumstances. This is a dynamic version of the well-known scenario testing approach. Classical scenario testing would also be a possibility, but given that we have five dependent variables per currency pair, the number of scenarios may quickly become rather high. Finally, one might also do statistical analysis of the output values of the simulation against the respective values of the input scenarios in order to make inference about sensitivity to FX rates. This approach is largely dependent on the drill-down capabilities of the DFA software used.

4.4 Simulation results

The following results from simulations with our example introduced in Section 2 show the impact that FX rates and their volatility have on the claims as experienced by the Home Office in its reporting currency USD. The first result shows the historical impact of the FX rates on the consolidated loss development:

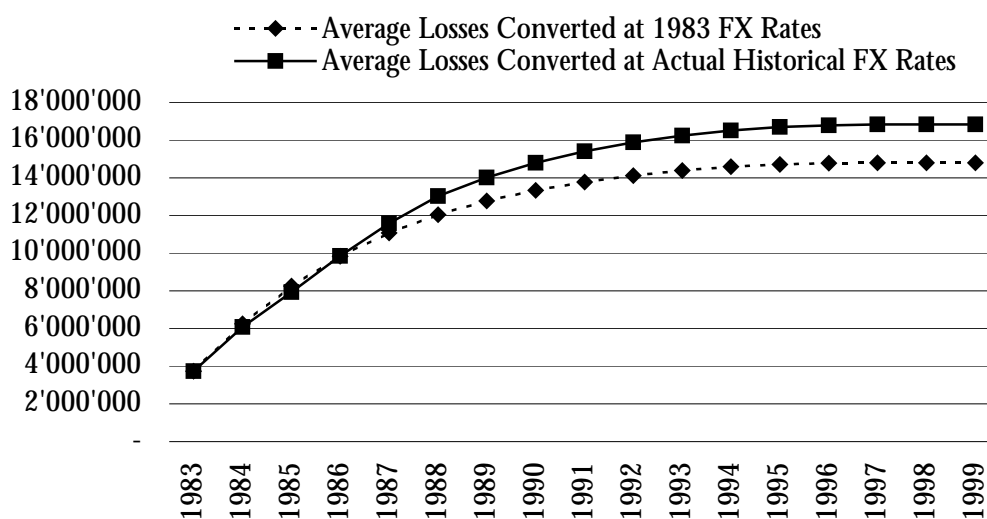


Figure 4: Historical impact of FX rate fluctuations (cumulated)

The (stochastic) loss amounts of the different lines of business were converted into USD according to two different deterministic exchange rate regimes:

- In the first case (dotted line) the losses of each year were converted at the FX rates of 1983, thus excluding any kind of FX rate fluctuation.

- In the second case (solid line) the same loss amounts of each year were converted at the historical FX rates that prevailed in the respective years.

Hence, the difference between the two lines shows the actual historical impact that the changing FX rates had on the cumulated consolidated losses as experienced by the head office in its reporting currency USD. It is easily seen that the difference in ultimate loss amounts is quite considerable: USD 16.82M under the actual FX rate development against USD 14.79M under constant rate. In the present case, the results from the historical FX rates are higher than the ones under the constant rate, since the USD underwent a considerable depreciation with respect to GBP, JPY, and CHF during the considered period (e.g.: 1 CHF = 0.48 USD in 1983, but 1 CHF = 0.67 USD in 1999). On the contrary, an appreciating USD with respect to the other currencies would have had a favourable effect on the results of the company.

Next, we compare projections of future losses (same stochastic simulation for the losses as before) converted at deterministic FX rates and converted at stochastically simulated FX rates.

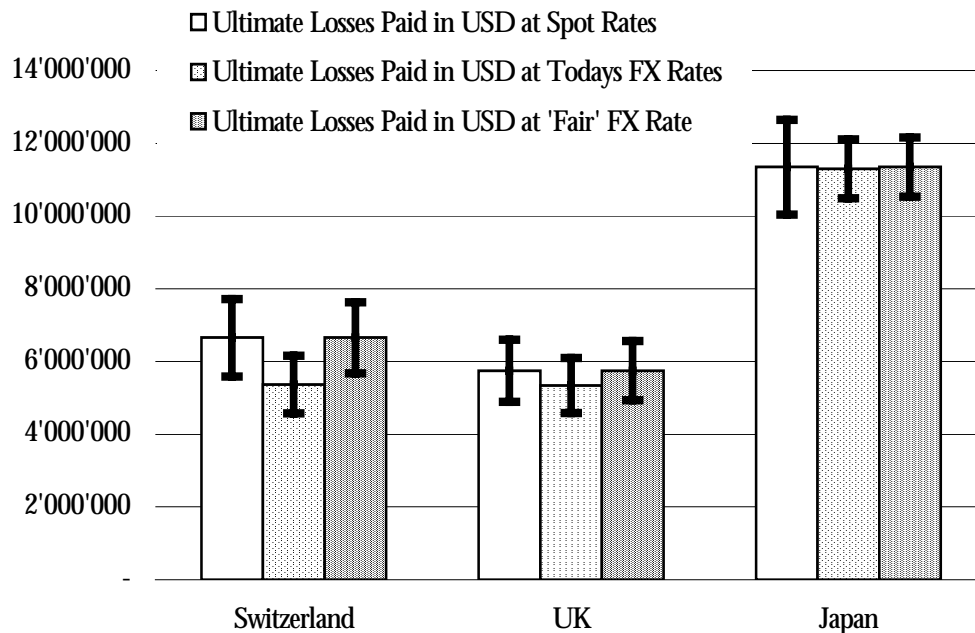


Figure 5: Ultimate losses under deterministic and stochastic FX rates

The figure shows the ultimate losses per line of business in USD, where the conversion of the yearly payments into USD was made according to three different FX rate regimes (the solid black bars indicating a confidence interval of plus/minus one standard deviation):

- FX rates stochastically simulated according to the bottom-up modular model introduced in Section 3 (referred to as “Ultimate Losses Paid in USD at Spot Rates” in the legend), with the values of year 2000 as initial values.
- FX rates deterministic and constant over time with the actual values prevailing in the year 2000 (“... Today’s FX Rates”).

- FX rates deterministic and constant over time with artificial values such that the mean ultimate losses become equal to the ones under the stochastic simulation (“... ‘Fair’ FX rates”). This approach is unrealistic, but allows best to compare volatilities with the case of stochastic FX rates.

For UK and Switzerland, the estimations for the means of the ultimate losses provided by the simulation with stochastic FX rates are quite different from the ones provided by the simulation with constant FX rates equal to the values of the year 2000. This is due to the fact that the actual rates in 2000 for CHF and GBP are below the long-term mean of the rates generated by the simulator. For JPY, the year 2000 rate is close to this mean, and hence there is no big difference.

The confidence intervals (indicating the estimated volatility of the ultimate losses) are larger for the case of stochastic FX rates than for the case of deterministic ones, but the difference is not always very big, i.e. in some cases (UK, see below) the inherent variance of the claims process is considerably higher than the extra variance added by the FX rate fluctuations. To get a clearer view of these differences, we give here a diagram of the 99th percentiles of the losses minus their expectation belonging to each of the above simulation results:

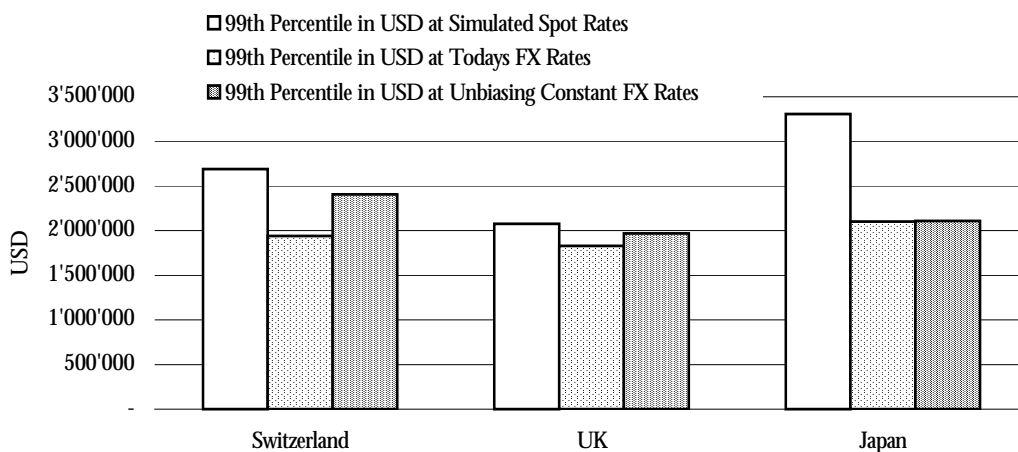


Figure 6: 99th Percentiles of (Losses - Expected Losses)

The picture here is also not uniform: In certain cases, the added volatility is quite considerable (Japan), whereas in other cases (UK) it is lower (but still above 5%).

From these investigations one can conclude that FX fluctuations definitely have an impact on amounts and volatilities of the losses, but the impact can be quite different in size. Comparing simulations with deterministic and stochastic FX rates provides in any case a good feeling for the order of magnitude of the impact.

5. FX risk management strategies

5.1 General strategies

Underwriting liabilities denominated in a foreign currency is almost daily business for reinsurers. Yet, they need to measure their performance in terms of their home currency and are thus exposed to FX risk. We start here by introducing general methods for the management of FX risk that apply to any company doing foreign business (see Luca (2000) and Loderer and Pichler (2000) for more details), and put them into the context of (re)insurance in the next section. There are many ways to deal with this type of risks, here are brief descriptions of some of them:

- ***Avoid*** the occurrence of FX risk by not doing business abroad. This may be an option if FX risks turn out to be uncontrollable, but not in general.
- ***Accept*** the FX risk and do nothing. This may be an option if exposure is sufficiently low or if the treaty is written in a country with a weaker currency than the reinsurer's.
- ***Diversify*** risk by doing business in different countries such that the respective FX risks are uncorrelated or even negatively correlated. Implementation of this approach will usually be difficult as business volume in different countries cannot be fully controlled but depends on demand and other collateral factors.
- ***Transfer*** risk to the customer by denominating the contract in home currency or by fixing a FX rate for all transactions related with the contract. The former is a full transfer of the FX risks to the customer, the latter a partial transfer; check the examples in section 4.2 to get a feeling. Notice however that this kind of risk transfer is already an implicit forward contract that will come at a price.
- ***Transfer*** risk to third parties: buy currency derivatives (futures, forwards, swaps, or options), or buy insurance of FX risks. These methods are called "off-balance sheet hedging". The effectiveness of this approach for large players starts, however, to be questioned, see Lyon (2001) for the example of a large international commodity trader.
- ***Reduce*** likelihood and severity of losses due to FX fluctuations by doing offsetting transactions. I.e. take sufficient asset positions in the foreign currency at the time (and FX rate) the contract is incepted and pay the liabilities from these positions, thus avoiding or reducing cash flows across the currency border that are exposed to FX risk. This approach is also called "on-balance sheet hedging", see Dacorogna et al. (2001) for more details.

The widely held wisdom in the (re)insurance industry is that foreign exchange exposure is not a problem as long as the reinsurer maintains assets in the currency in which the liabilities are denominated, i.e. does on-balance sheet hedging. In principle, this is true, although such a strategy might be difficult to follow in practice and not optimal as it is equivalent to fully hedging the position. For investments, it has been shown by Froot (1993) that at horizons of several years, complete hedging not only does not lower return variance, it actually increases it for many portfolios. In another context, Levich and Thomas (1993) analyze the impact of active currency

management on internationally diversified portfolios and show that it can significantly improve returns without overly affecting the risk. Dacorogna et al. (2001) show how it is possible to separate the different aspects of the investment strategies and to treat the exchange exposure by itself. In the financial industry the foreign exchange risk is not only an object of study but also a practical problem leading to financial products. There are firms like Pareto Partners that offer specially designed currency overlay programs to protect foreign investment in an internationally diversified portfolio.

5.2 Special issues in insurance and reinsurance

Unlike in international banking and investment or commodity trading, insurers and reinsurers are faced with liabilities (i.e. claims) that are stochastic in amount and timing, i.e. there may be considerable uncertainty on when to pay and how much (see also Section 4.4 to get a feel). Moreover, (re)insurance contracts may imply very long time horizons for paying, e.g. ten up to twenty years. Finally, (re)insurance companies are subject to regulations that may prevent them from placing their assets in the way they want. These specific factors have an impact on some of the general FX risk management strategies set forth in the previous section:

Taking out currency derivatives (futures, forwards, options, swaps) becomes difficult due to the timing uncertainty as one does not know exactly for what maturity and for what amount to buy them, which can result in considerable mismatches. Exchange-traded FX derivatives are only available for time horizons up to one or two years, and it will also be difficult to take out over-the-counter (OTC) products for longer times to maturity. Hence, the time horizon of many (re)insurance liabilities is too long for protecting them with FX derivatives.

On-balance sheet currency hedging also becomes more difficult: One does not know how much to invest and at what duration in order to match the liabilities. One may invest too much (at possibly sub-optimal returns), or too little (such that funds must be nevertheless transferred across the currency border). If one strictly matches all liabilities in foreign currencies by sufficient amounts of assets in the same currency, this may also lead to a silo effect and to the use of more risk based capital (which also has a cost) than would be necessary under a more flexible currency hedging regime. Moreover, regulatory constraints may prevent the (re)insurer from taking certain hedge positions. Finally, an insurer is also an investor that must invest its assets at an optimal rate of return, and investments in certain currencies (dictated by hedging requirements) may lead to sub-optimal returns in certain time periods.

To get a feeling for the working of on-balance sheet hedging under stochastic liabilities, we revisit our simple example from Section 4.2 and consider the case of a 10 XS 10 contract denominated in GBP where the reinsurer bears the full currency risk. In order to (partly) hedge its possible liability, the reinsurer opens a position of 5 GBP at $t = 0$ and at 1.5 USD per GBP, hence the

value at $t = 0$ of the hedge position is 7.5 USD. At $t = 1$ when the loss occurs, the USD/GBP rate is at 1.2, 1.5, or 1.8. The balance of the reinsurer after settling the claim looks as follows:

Rate:				1.2	1.5	1.8
Claim (GBP)	Payment (GBP)	Hedge Account Balance (GBP)	Reinsurer's Balance (USD)			
5	0	5	6	7.5	9	
10	0	5	6	7.5	9	
15	5	0	0	0	0	
20	10	-5	-6	-7.5	-9	

If the contract does not trigger (claims of 5 and 10 GBP), then the reinsurer has a gain or loss on its unused deposit depending on the FX rate prevailing at $t = 1$. If the claim amount is 15 GBP (i.e. the payment is 5 GBP), the hedge account exactly covers the loss and no funds transaction over the currency border takes place. If the full loss occurs (i.e. 10 GBP payment), then the money on the hedge account is not sufficient and another 5 GBP must be transferred at prevailing FX rates, leading to transaction gains or losses with respect to the original rate of 1.5. The example can easily be adapted for other hedge amounts (e.g. full hedge or no hedge). In every case, due to the uncertainty on the actual amount of the claims payment, there is a potential for gains or losses due to FX rate changes.

5.3 On-balance sheet hedging concepts

On-balance sheet currency hedging means that a (re)insurer who has to cover liabilities of value s in a foreign currency invests a hedge amount of s_h in the respective currency in order to cover the liability, i.e. in order to be able to pay (at least a part of) the claims without funds transactions over the currency border. The concepts presented in this section come from Dacorogna et al. (2001). The *hedging ratio* is defined as $h = s_h / s$ and denotes the portion of the total liability covered by the hedge amount. In a multi-period setup, s denotes the total liabilities from “now” up to some fixed time horizon.

The determination of the liability amount s can already be a problem in the context of (re)insurance: If the maximum liability amount is known (e.g. in case of an XL b contract), there is no problem and s would be set to this maximum liability amount (e.g. $s = a$). In many cases, however, the maximum liability amount is unknown, as in our simulation example. Possibilities to deal with this situation include:

- Set s equal to the $x\%$ (e.g. 99%) quantile of the loss distribution and express s_h in terms of this s . The advantage here is that the hedge amount s_h can be expressed as a percentage of a fixed amount s . The disadvantage is that there exists a residual risk that is not covered by the hedging considerations, i.e. the liability amounts above the $x\%$ quantile.
- Directly take s_h as the $y\%$ quantile of the loss distribution and do not treat s explicitly. The disadvantage here is that the hedge amount is not expressed with respect to some fixed maximum, but the advantage is that there is no residual to be treated separately.

Other concepts, e.g. based on expected loss plus some security loading, or on PML would also be possible. We find however that the concepts based on the quantiles of the loss distribution are the most objective ones.

The simplest strategy is *static hedging*. This strategy consists of taking positions in the foreign currency at inception of the contract and keeping them until the end without adding or withdrawing funds (except for actually paying liabilities and unless the liabilities exceed the funds and the negative balance must be covered by transfer payments). In practice this could be to place the necessary reserves in the currency in which the claims have to be paid. There are two ways of implementing such a hedge: either by covering fully the liabilities as they are expected to occur (i.e. $s_h = s$), this is called *full hedging*, or by covering only a part of the liabilities with an investment in the respective currency, this is called *partial hedging* (i.e. $s_h < s$). A portion of $h \cdot 100\%$ of each claim would then be paid from the hedge amount, and the rest from funds transferred across the currency border.

In a multi-period setup, it is also possible to readjust the hedge amount s_h for each time interval depending on the actual course of events and changed expectations of the future. This approach is called *dynamic hedging*. Funds can be added to or removed from the foreign investment, or the portion of the liabilities of a certain time interval that is paid from local funds can be adjusted dynamically depending on the actual development of the FX rate. As can be seen from the simple example in the previous section, it can be worthwhile for our US reinsurer to pay its GBP-denominated liabilities with USD funds (although a sufficient GBP amount would be in place), if the GBP has depreciated with respect to the USD.

Besides these forms of hedging, it is possible to combine a partial static hedging with a dynamic hedging strategy designed to take profit of the FX movements. The ideal situation would be to only hedge the FX risk when the rate moves in an unfavorable direction and to fully profit from the movement in the other direction by not hedging. Of course, the problem remains to determine in advance when the movement will occur in the right direction.

In all cases, it is essential to find methods to determine an optimal value for the hedge amount: once and for all times in the case of static hedging, and for each period afresh in the case of dynamic hedging. These methods must take into account the possible development of the liabilities and of the drivers of the exposure: inflation and interest rates in both countries. This can be done by using a DFA model as will be shown in this paper. The hedge ratio will be determined by considering the return and the risk of the entire portfolio and choosing the solution which is the closest to the efficient frontier. The optimization of the hedge ratio can be made for each currency separately (this is *single-currency hedging*), or for all currencies simultaneously (*multi-currency hedging*), thus trying to achieve an overall optimum. Optimality can be defined in different ways: In the most basic case, the only goal is to service the liabilities

reliably, in other cases, this may be superimposed by profitability requirements for the invested funds. When optimizing hedge amounts (static and dynamic case) constraints must also be taken into account. This can be position limits for foreign investments, or – in more advanced applications – duration matching issues between liabilities and invested assets. Particularly when designing dynamic hedging strategies, one should also be aware of transaction costs. These are usually low for FX transactions, but may nevertheless consume a part of possible benefits when transactions occur very frequently.

5.4 Modelling issues

As a general prerequisite for modelling currency hedging, the DFA model must be enabled for international modelling, i.e. a scenario generator for FX rates and related variables (as explained in Section 3), a model for investments in different currencies, and a model for at least FX transaction exposure (as explained in Section 4) must be in place.

Modelling of static (full or partial) hedging is simple. The liability and hedge amounts can be determined prior to the simulation. A preliminary DFA simulation with only the loss model can be used to obtain the distributions of the liability amount and other variables that may be used to determine the hedge amount. The determination of the hedge amount can be done outside the DFA model (along the lines presented in the previous section), hence no specific extensions of the DFA model are necessary. Modelling of cash flows from assets positions to pay claims is basic DFA functionality.

Modelling of dynamic hedging is simple as long as the hedge amounts are not constant over time, but determined prior to the DFA simulation (and hence independent from the actual development of the respective variables in some given scenario). This approach makes sense if one models clear trends in some of the underlying input variables. Think e.g. of a US reinsurer who expects the GBP to depreciate in the subsequent years. Then it would make sense to decrease the hedge ratio accordingly over time, not based on the concrete scenario but only on the drift in the expectations of the scenario generator.

Modelling of dynamic hedging becomes difficult if the hedge amounts (or hedge ratios) are adaptively adjusted depending on the actual development of the losses or FX rates. In this case, the basic international DFA model must be extended with respective dynamic adjustment functions. A simple example would be a hedge ratio that is made proportional to the changes in the underlying FX rate. In more complex cases, this can also include the computation of conditional expectations and embedded optimization functions, in particular in the case of multi-currency hedging. In this context, if one considers foreign currencies also as a mean for obtaining trading gains, then a model for the revenues of a trading system must be in place (see Dacorogna et al. (2001) for more information).

5.5 Simulations

In this section we investigate static partial on-balance sheet hedging strategies with different hedge ratios for our example. To this end we consider our simulation setup as a full loss portfolio transfer that must be funded appropriately. The distributions of the loss amounts in local currency (L^{CHF} , L^{GBP} , L^{JPY}) and of the consolidated total loss amount in USD (L^{USD}) were obtained by simulations as shown in Section 4.4. Conversion of losses from local currency into USD was done at stochastically simulated FX rates. The reference liability amounts s for hedging (as introduced in Section 5.3) were defined to be the expected loss plus standard deviation for each LOB (in local currency), i.e.

$$s^{CHF} = E[L^{CHF}] + \sigma[L^{CHF}], s^{GBP} = E[L^{GBP}] + \sigma[L^{GBP}], s^{JPY} = E[L^{JPY}] + \sigma[L^{JPY}].$$

The total funding amount in USD that is available at time 0 for covering the loss portfolio was determined in the same way:

$$s^{USD} = E[L^{USD}] + \sigma[L^{USD}].$$

At time 0, the amounts $h \cdot s^{CHF}$, $h \cdot s^{GBP}$, $h \cdot s^{JPY}$ were invested in CHF, GBP, and JPY respectively (converted at the initial FX rate). The remainder of the total capital s^{USD} was kept in USD. h denotes the hedge ratio as introduced in Section 5.3. All investments were simulated as portfolios of a cash account and bonds with different times to maturity. The cash flows of the investments were matched with the expected cash flows of the liabilities. During the simulation, all transfer payments from USD into the other currencies were converted at current (simulated) FX rates, and at the end all foreign-denominated investments were converted back into USD at the final FX rates. The following figure shows the net present value in USD of the loss portfolio hedged in the described way for various hedge ratios:

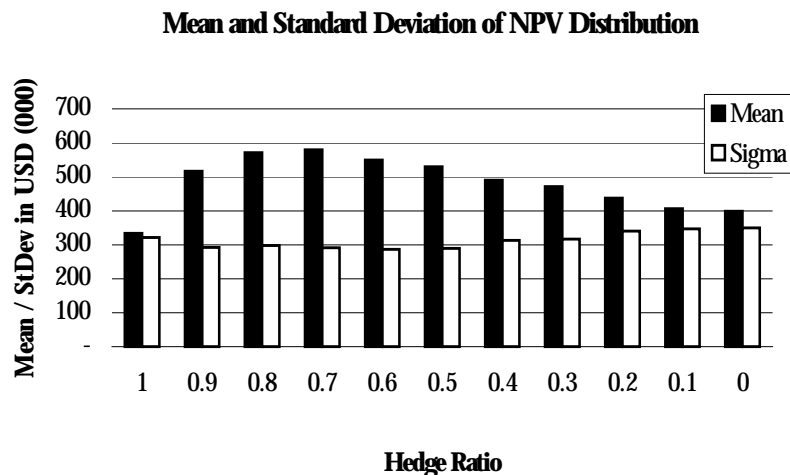


Figure 7: NPV of hedged loss portfolio

The figure shows very clearly that neither full hedging nor no hedging is the optimal strategy for the given setup. A hedge ratio of about 0.7 yields the highest return with a relatively low risk. There are different factors that may have contributed to this result:

- US interest rates are considerably higher than the ones in Japan and Switzerland and about equal to the British ones. Hence, investments kept in USD generate higher returns than investments in CHF and JPY, and investments in GBP do not yield much more than US ones.
- None of the three foreign currencies have a clear upward or downward trend with respect to the USD. Therefore, in some cases it is favourable to keep the money in USD, in other cases it is favourable to have the money in foreign currencies. On the average, a mixed portfolio with investments in USD as well as in foreign currencies (i.e. partial hedging) seems the most reasonable.

This is, of course, only one example, and results may look different for other loss portfolios and other involved currencies. But the results here show that it is definitely worthwhile to consider partial currency hedging instead of full currency hedging. Returns from investments in foreign currencies could be further optimized by using dynamic currency overlay strategies as described in Dacorogna et al. (2001). The implementation of such strategies in the simulation model was, however, well beyond the scope of this study.

6. Reinsurance cover

Three reinsurance strategies for our example have already been treated implicitly by the simulations in Sections 4.4 and 5.5: Retention of the whole portfolio by the cedant, loss portfolio transfer (an ART transaction), and a quota share where the cedant retains only a part of the whole loss portfolio (and where all considerations apply to the cedant and to the reinsurer for their respective share of the loss portfolio). We have seen for these cases that, indeed, FX rate fluctuations bring additional volatility that has to be dealt with, and that careful selection of the hedge ratio can have a considerable impact on risk and return of the loss portfolio.

We complement these cases here by the consideration of a non-proportional treaty. We introduce a USD 20M XS 12.5M Stop Loss treaty on the consolidated losses of the portfolio in USD, cumulated over time. I.e. cumulated losses above USD 12.5 Million and up to USD 32.5 Million would be paid by the reinsurer. For this case we have a look at expected claims and their volatility in the layer of the reinsurer:

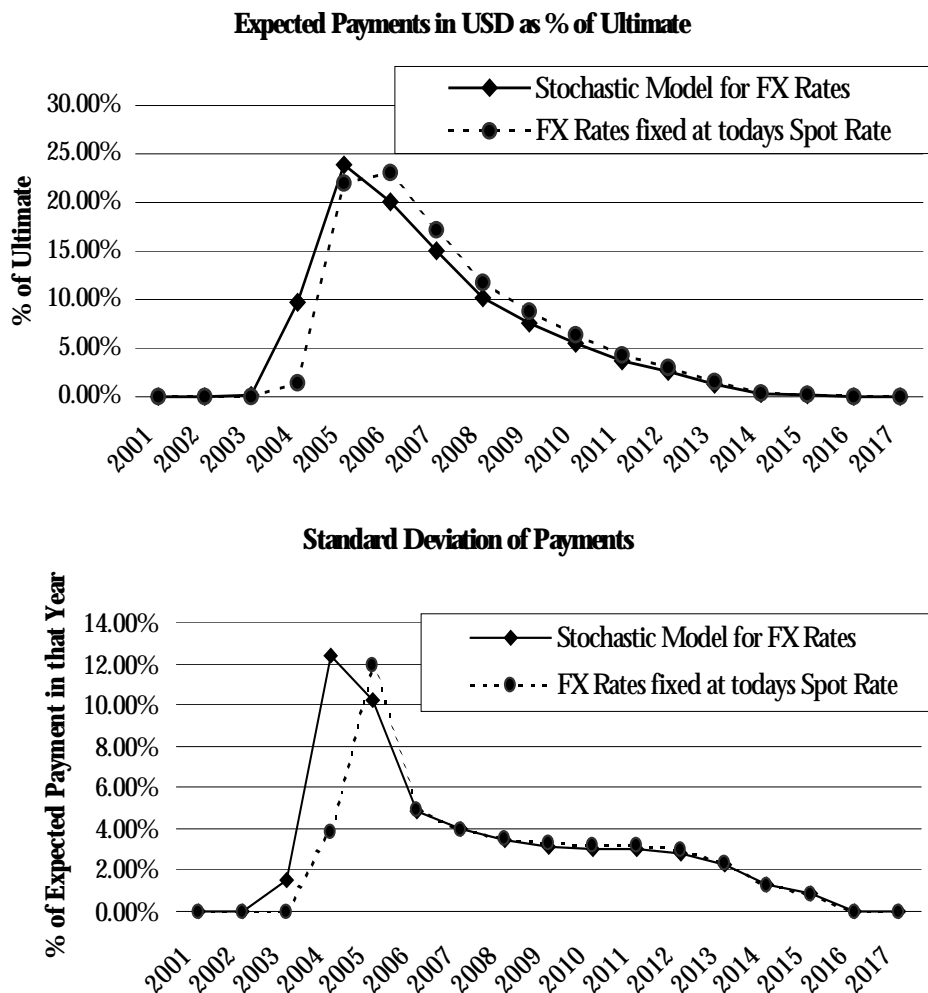


Figure 8: Mean and standard deviation for payments in R/I layer

The payment figures here are given as percentages of the ultimate loss of the portfolio, knowing that the latter is different for fixed and stochastic FX rates. In the first years, the contract does not

trigger at all under both FX rate regimes. The differences in the levels and in the volatility of the payments are very considerable in the third up to the sixth year, because there is not only the added volatility of the FX rate itself, but FX fluctuations have also an impact on whether or not the contract triggers at all. Moreover, in this time span the absolute size of claims is still quite high; see also Section 4.4. For the rest of the time interval the differences between fixed and stochastic FX rates are not very high, since there is no longer uncertainty as to whether or not the contract triggers, and since the absolute size of additional claims arising in these years is lower. Looking at the ultimate cumulated payment amount of the reinsurer, the difference in volatility is quite low: the risk ratio (i.e. standard deviation divided by mean) of the reinsurer's payments is 26.2% at fixed FX rate as opposed to 27.8% under stochastic FX rates. Hence, because of the structure of the claims and because of the long duration of the contract, the difference in volatility for the reinsurer is not very high, but the changes in the FX rate add considerable uncertainty to the timing of the payments in the time period when the FX rate has an impact on whether or not the contract triggers at all (a phenomenon that is well known for X/L treaties in the case of claims inflation). This fact may be important for the reinsurer when structuring the hedge account from which to pay the claims. The added volatility due to uncertainty whether or not the contract triggers would be higher in the case of a series of individual X/L treaties for each year, because in that case contract triggers could arise each year and not only once for the whole period. On the absolute scale, the mean ultimate cumulated payments of the reinsurer amount to USD 4.46 Million if the FX rates are kept constant at initial values. Under stochastic scenarios for the FX rates this increases to USD 5.46 Million.

We now apply exactly the same currency hedging approach as in Section 5.5 to the claims of the reinsurance layer. Be aware that the situation for the reinsurance company is now slightly different. The present reinsurance cover is fully denominated and payable in USD (which we assume to be the reinsurer's home currency); we only have a dependence of the claims process on the FX rates. The same approach could, however, also be worked out for a foreign-denominated reinsurance contract where the FX rates enter more directly. The figure below shows that a low hedge ratio (20%), or even no hedging, is clearly better than a high hedge ratio or full hedging. Recall that for the whole loss portfolio the optimal hedge ratio was found to be at 70%.

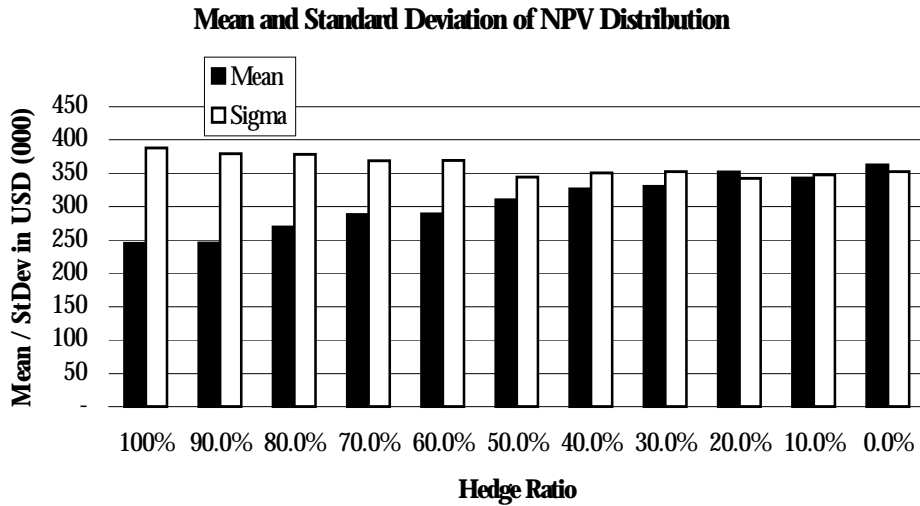


Figure 9: Hedge results of reinsurer

The reason for this lower hedge ratio may be that the Stop Loss contract does not cause any claims at all during the first two or three years (see above), i.e. the full hedge amount can be invested during this time period. As we have mentioned earlier, the US interest rates are considerably higher on average than the Japanese and Swiss ones, and not much lower than the British ones. Therefore, high US investment returns over the first years may outweigh the FX risks for the later years.

7. Assessment and outlook

Measurement, modelling, and management of FX risks is a vast and difficult field, for internationally operating companies in general and for insurers and reinsurers in particular. This paper can only scratch the surface of the issue and, hopefully, give some input for further research and practical application. The style of this paper was kept informal, and many topics had to be simplified or omitted in order not to overburden the presentation. Here is a survey of the most important findings:

In scenario generation, the modelling of dependences between the FX rates and other economic variables, that could be significantly measured in data analysis, is very important. In yearly data there are simple and also economically defensible equilibrium relations that can be used for scenario generation. On finer time scales (quarterly and monthly) the situation is more complicated, as dependences on past values at different lags become highly significant. There is no silver bullet for scenario generation. On the yearly level, the presented bottom-up modular approach provides a simple and easy-to-use solution, which comes, however, at the price of decreased flexibility if one wants to incorporate deterministic trends or simulate according to conditional probabilities. The alternatives are purely statistical multivariate time series models, e.g. the VAR model as presented in this paper. Finding an international economic model that captures the behaviour and dependences of the relevant economic variables on the one hand, but that is also understandable and fully tractable by the actuarial and financial practitioner on the other hand is clearly an important issue for future applied research. The approach to be chosen is likely to be rather statistical than economical, as the economic relations between the involved economic variables mainly apply on very long time horizons. After all, even the best scenario generator leaves us with some fundamental uncertainty on the adequacy of the simulated scenarios for the future. It is therefore important to complement all simulations with sensitivity analyses that give at least a hint for the exposure of the company to events or developments not covered by the generated scenarios.

FX fluctuations affect the business of (re)insurers in different ways: through transaction, translation and economic exposure. Translation exposure is difficult to model in DFA as it largely depends on the accounting standard in use and was therefore not substantially treated in this paper. Exploring the details of modelling translation exposure in DFA under some important accounting standards (in particular IAS and GAAP) would be an interesting subject for more accounting-oriented researchers. Economic exposure was also only mentioned here. A thorough treatment of economic exposure due to FX rate changes would have to go in line with research on economic cycles in insurance. Modelling of transaction exposure in DFA was found to be relatively straightforward, although dependent on the concrete model in use. Our investigations show that, indeed, FX rate changes can have significant effects on the business of an insurance or reinsurance company, particularly over long time horizons and there is also an additional

volatility due to the fluctuations of the FX rates. Whether or not these additional uncertainties must be taken into account will usually depend on their quantitative relation to the genuine volatility in the claims process. The message here is that, whenever FX rates influence the business, one should not neglect them a priori, because they potentially have a quite dramatic impact on the business results of the (re)insurer. DFA-style stochastic simulation is one way of determining the impact of FX rate fluctuations on the business results. An alternative would be scenario testing with different fixed FX scenarios but otherwise stochastic input variables. This can very well show the sensitivity of the results against FX rate changes, but it does not yield information about additional volatility due to FX rate uncertainty.

There exists a number of approaches for FX risk management. The respective concepts mainly come from the area of banking and investment. In (re)insurance, some important specialties must be taken into account: liabilities are stochastic in amount and timing, time horizons may be very long, and regulatory constraints may prevent (re)insurers from implementing certain strategies. On-balance sheet currency hedging in its various forms was identified as the most important universally applicable strategy for FX risk management. Simulations showed that there exist situations where full currency hedging is clearly dominated by partial hedging strategies. Currency hedging in an insurance context is more complex than in a banking context because there is additional uncertainty as to amount and timing of the liabilities to be hedged. Hedging considerations can be based on different approaches, e.g. mean plus standard deviations or also percentiles. Only the most basic static hedging strategies were investigated here. Dynamic multi-currency hedging strategies (so-called currency overlay, see Dacorogna et al. (2001)) would bear a potential for even more optimization, but adaptation of these methods for the specific needs of insurance liabilities is an open point for further research.

Our findings for primary insurance apply to reinsurance as well. For the cases of loss portfolio transfer and proportional reinsurance, everything is equivalent. In the case of non-proportional treaties (X/L and Stop Loss) FX rate fluctuations create an additional uncertainty as to whether or not a layer triggers at all, with the respective impact on expectation and volatility of the claims and, hence, on the price for the cover.

This paper has shown very informally that FX fluctuations can bear dangers as well as opportunities for an internationally operating insurer or reinsurer, and it is undoubtedly worthwhile to explore the downside as well as the upside of this risk. Further, probably more analytical, research on the following topics could add considerable value:

- impact of FX rate volatility on the pricing of international reinsurance contracts and on the methodology for risk based capital, in particular: determination of possible extra amounts of risk based capital to cover FX risks (in particular: determination of the sign of this amount),
- general methodology for the management of a portfolio of assets and (insurance) liabilities in the presence of volatile FX rates.

Bibliography

- BAILEY G.T., K. WILD (2000): *International Accounting Standards: A Guide To Preparing Accounts*. Second Edition, abg Professional Information.
- BLUMSOHN G. (1999): "Levels of determinism in workers compensation reinsurance commutations", *CAS Proceedings*, Vol. LXXXVI, Part 1.
- CHAN K.C., G.A. KAROLYI, F.A. LONGSTAFF, A.B. SANDERS (1992): "An empirical comparison of alternative models for the short-term interest rate", *Journal of Finance*, 47, 1209 – 1227.
- CLEMENTS M., D. HENDRY (1998): *Forecasting economic time series*. Cambridge University Press.
- CLOSTERMANN J., B. SCHNATZ (2000): "The determinants of the euro-dollar exchange rate", Discussion paper 2/00, Economic Research Group of the Deutsche Bundesbank (www.bundesbank.de).
- DACOROGNA M. M., R. GENÇAY, U. A. MÜLLER, R. B. OLSEN, O. V. PICTET (2001): *Introduction to High Frequency Finance*. Academic Press.
- D'ARCY S.P., GORVETT R.W., HERBERS J.A., HETTINGER T.E., LEHMANN S.G., MILLER M.J. (1997): "Building a public access PC-based DFA model", Casualty Actuarial Society.
- DAYKIN C.D., T. PENTIKÄINEN, M. PESONEN (1994): *Practical Risk Theory for Actuaries*, Chapman and Hall, 1994.
- Dembo R., A. Aziz, D. Rosen, M. Zerbs (2000): *mark-to-future, A Framework for Measuring Risk and Reward*. Algorithmics Publications, (www.mark-to-future.com).
- EMBRECHTS P., C. KLÜPPELBERG, T. MIKOSCH (1997): *Modelling Extremal Events for Insurance and Finance*. Springer, New York.
- EMBRECHTS P., A. MCNEIL, D. STRAUMANN (1999): "Correlation: pitfalls and alternatives", *RISK Magazine*, May, 69 – 71.
- FASB 52 (1981): "Statement No. 52: Foreign currency translation", Financial Accounting Standards Board (www.fasb.org).
- FROOT K. A. (1993), "Currency hedging over long horizons", NBER Working Paper No. 4355.
- KAUFMANN R. AND PATIE P. (2000), "Strategic long-term financial risks", Intermediate research report, Risklab / ETH Zürich (www.risklab.ch).
- KIM J., MALZ A.M., MINA J. (1999): *LongRun Technical Document*. RiskMetrics Group (www.riskmetrics.com).
- LEVICH, R. M. AND L.R. THOMAS, L. R. (1993): "The merits of active currency risk management: evidence from international bond portfolios", NBER Working Paper No. 4340.
- LODERER C., K. PICHLER (2000): "Firms, do you know your currency risk exposure? Survey results", *Journal of Empirical Finance*, 7, 317 – 344
- LUCA C. (2000): *Trading in the Global Currency Markets*. New York Institute of Finance (Prentice Hall).
- LYON (2001): "BHP drops derivatives", *RISK Magazine*, February, 13.
- MACDONALD R. (2000): "Concepts to calculate equilibrium exchange rates: an overview", Discussion paper 3/00, Economic Research Group of the Deutsche Bundesbank (www.bundesbank.de).
- MÜLLER U. A, DACOROGNA M. M. AND PICTET O. V. (1998): "Heavy tails in high-frequency financial data" in *A practical Guide to Heavy Tails*, ed. by R. J. Adler, R. E. Feldman and M. S. Taqqu, Birkhäuser, Boston.

Appendix A: Economic data

The following table shows the yearly data for the economic variables that were used throughout our study. All data were obtained from *Datastream*.

- **R** denotes the short-term interest rate of the respective country in percent (annualized). The actual figures are interbank money market rates.
- **I** denotes the inflation rate of the respective country in percent (annualized). The actual figures are consumer price inflation rates.
- **S** denotes the spot exchange rate of the respective currency, i.e. the spot price in USD for one unit of foreign currency (except for Japan: USD price of 100 Yen).

Year	USA		United Kingdom			Switzerland			Japan		
	R [%]	I [%] ⁶	R [%]	I [%]	S [\$]	R [%]	I [%]	S [\$]	R [%]	I [%]	S [0.01\$]
1981	16.955	9.076	13.993	12.302	2.027	9.246	6.782	0.511	7.370	4.500	0.455
1982	13.260	3.883	12.351	5.607	1.750	5.168	5.619	0.494	6.777	2.372	0.403
1983	9.725	3.931	10.186	5.485	1.516	4.210	2.167	0.477	6.358	1.980	0.421
1984	10.871	3.939	10.003	4.744	1.336	4.435	2.963	0.427	6.147	3.024	0.421
1985	8.392	3.773	12.297	5.848	1.295	5.039	3.353	0.411	6.415	1.513	0.422
1986	6.829	1.168	10.998	3.799	1.467	4.332	0.069	0.559	4.971	-0.271	0.597
1987	7.176	4.447	9.765	3.834	1.639	3.888	1.956	0.672	4.123	0.933	0.694
1988	7.973	4.510	10.380	6.960	1.781	3.202	1.986	0.685	4.352	1.124	0.781
1989	9.273	4.695	13.942	7.850	1.639	7.068	5.104	0.612	5.259	2.944	0.726
1990	8.267	6.107	14.814	9.817	1.784	8.955	5.346	0.723	7.429	3.978	0.694
1991	5.961	3.090	11.550	4.577	1.769	8.244	5.370	0.700	7.058	2.790	0.744
1992	3.822	2.890	9.675	2.807	1.764	7.913	3.496	0.713	4.310	1.260	0.790
1993	3.306	2.811	5.983	2.092	1.502	4.943	2.532	0.677	2.922	1.075	0.902
1994	4.739	2.625	5.559	2.969	1.532	4.173	0.463	0.733	2.226	0.782	0.980
1995	6.022	2.570	6.749	3.419	1.578	3.074	2.018	0.847	1.176	-0.338	1.070
1996	5.484	3.401	6.089	2.551	1.562	2.055	0.814	0.810	0.558	0.654	0.920
1997	5.715	1.648	6.930	3.603	1.638	1.723	0.406	0.690	0.478	2.110	0.827
1998	5.542	1.706	7.381	2.910	1.657	1.561	-0.160	0.691	0.511	0.709	0.768
1999	5.414	2.637	5.514	1.796	1.618	1.399	1.693	0.666	0.166	-0.985	0.883
2000	6.529	4.115	6.187	3.231	1.516	3.078	1.539	0.593	0.342	0.029	0.928

Figures for interest rates and spot exchange rates were obtained by aggregating daily data: In a first step, a monthly aggregate was generated by taking the median of the daily quotes, then yearly aggregates were obtained by taking the average of the monthly aggregates. Inflation was aggregated by averaging annualized monthly inflation figures.

Below we give a series of charts that present the above-stated time series in different combinations so as to give a feeling not only for the behaviour of the single time series, but also for their relations with one another. The solid lines always denote the FX rate, the dashed lines denote the interest rates, and the dotted lines denote the FX rates. The difference between the dashed and the dotted line is the real rate of return, which turned out to be the most significantly explanatory factor for the FX rates in the bottom-up modular model, see also Appendix B. These charts also allow to check more systematically the findings of the statistical analysis presented in Section 3.2.

⁶ Notice: The inflation figures given here do not correspond to the ones usually used by US practitioners. We used here a CPI All Items series from Datastream and averaged annualized monthly inflations. Usual practice is the US is to use the CPI Urban series from the Bureau of Labor Statistics and to compute the inflation for a given year based only on the CPI values at the end of that year and the previous year. Given the purely illustrative nature of our example, this deviation from common practice is, however, irrelevant.

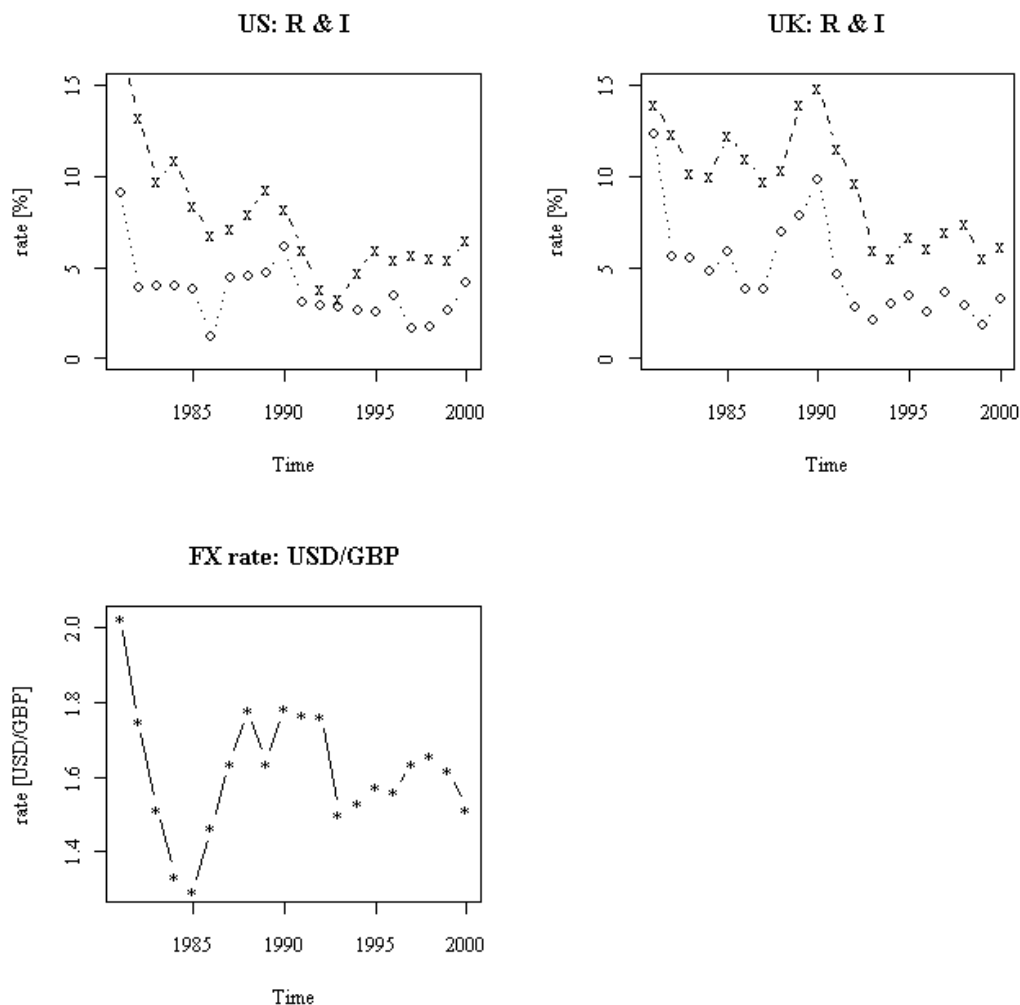


Figure A 1: USD/GBP rate and its constituents

Remarks:

- The strong correlation between interest rates and inflation can be clearly observed for all countries. The real rate of return (interest minus inflation) is almost always positive, but by far not constant over time, and there are also considerable differences between the countries.
- Also, it can be observed that the interest rates of the different countries are quite strongly correlated with each other. Periods of high and low interest rates occur quite simultaneously in all countries, although the absolute levels of interest are rather different from country to country, and also the relative increases and decreases differ.
- The interest rates of the US and the UK are generally higher than the ones of Switzerland and Japan, which may have an impact on investment decisions when investment returns matter (see the example in Section 6).
- All three FX rates against the USD show a similar pattern in terms of increases and decreases, which shows the role of the USD as a world reference currency. Changes in FX rates from one year to the next can be quite drastical: 20% and more.
- All series show very clear signs of autocorrelation.

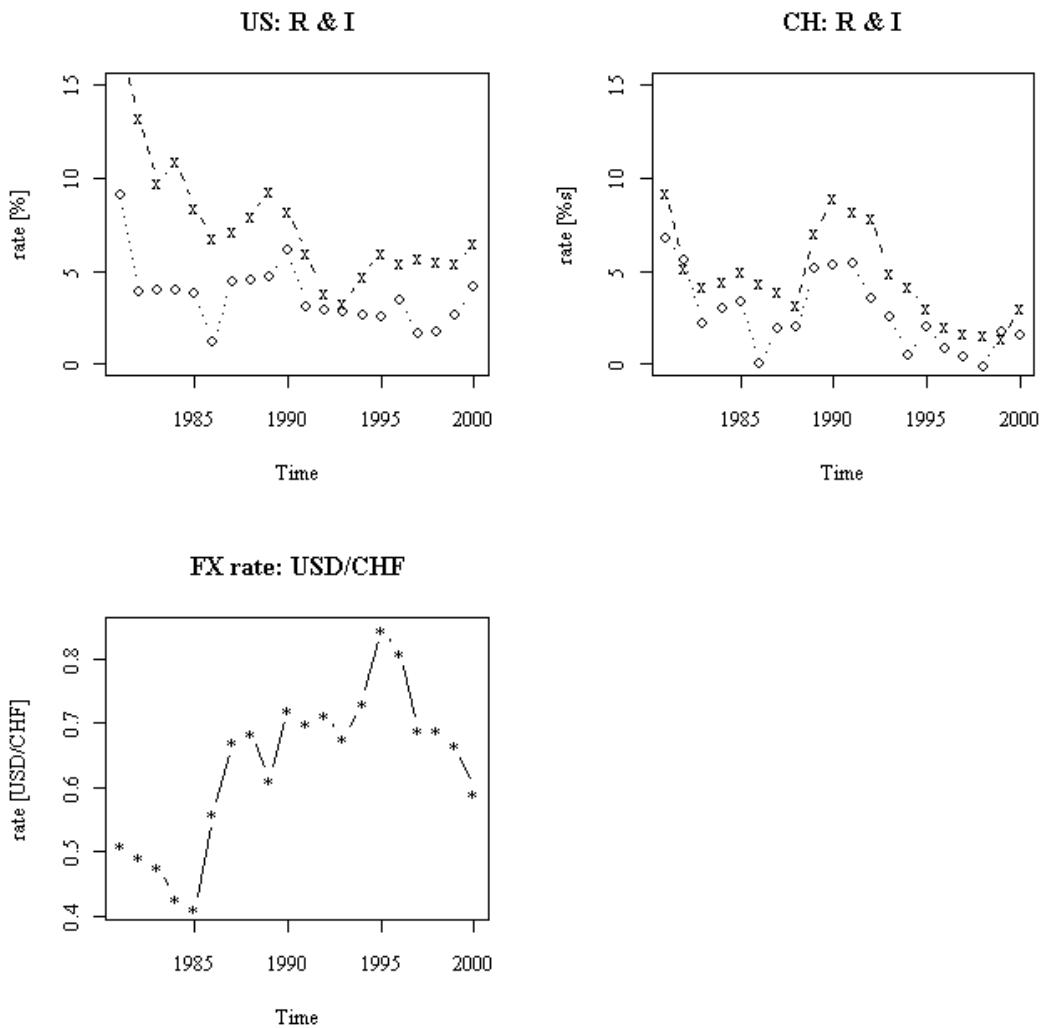


Figure A 2: USD/CHF rate and its constituents

Remarks (cont'd):

- In spite of its reputation as a safe haven currency, also the CHF makes rather strong moves with respect to the USD, see also Luca (2000).
- The high degree of simultaneity in all series suggests that the relevant economic variables of the involved countries are driven by some common economic cycle.

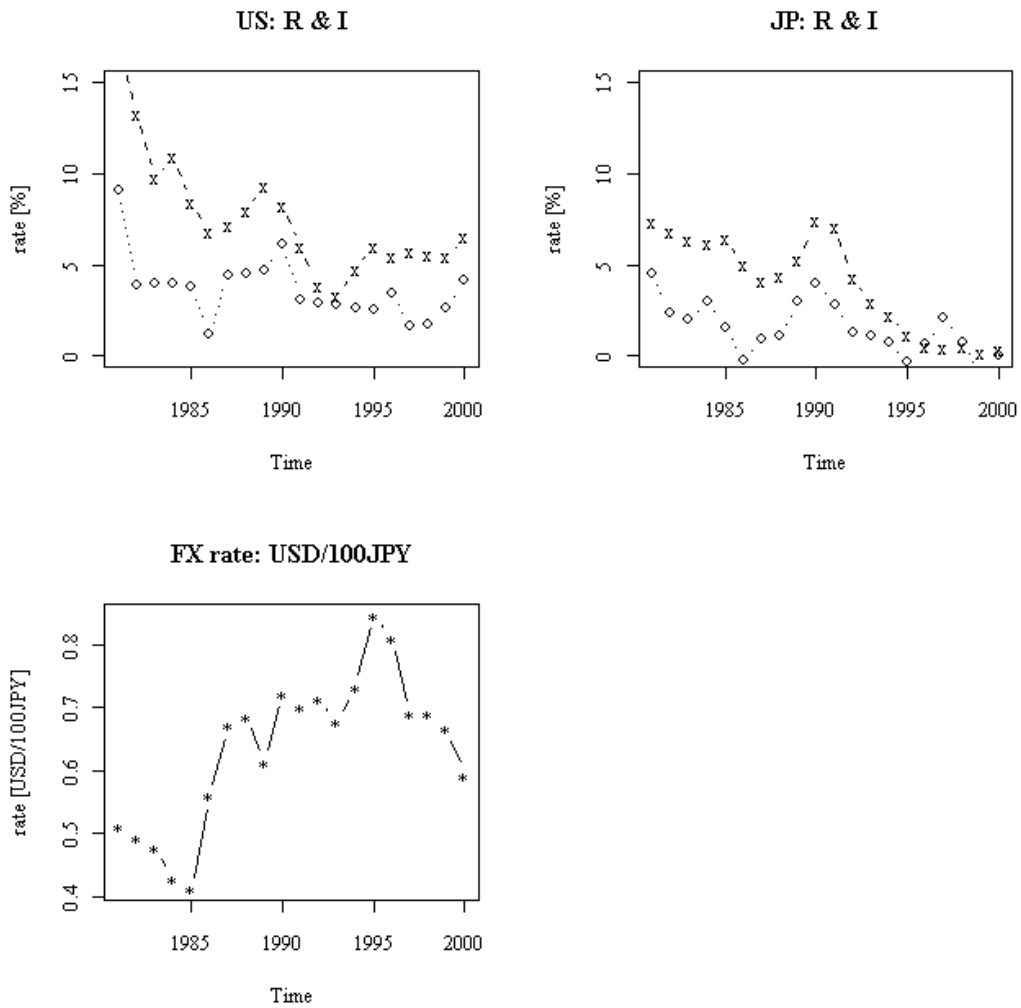


Figure A 3: USD/JPY rate and its constituents

Remarks (cont'd):

- Interest rates in Japan and also inflation can attain remarkably low levels (close to zero), which can have a strong impact when hedge ratios are computed and investment returns play a role.

Appendix B: Model parameters

We now present the estimated parameters for the bottom-up modular model as we introduced it in Section 3.3. Refer the latter section for a detailed explanation of the formulae and parameters.

i.) Interest rate generator

$$r_t^{XXX} = r_{t-1}^{XXX} + a^{XXX} (b^{XXX} - r_{t-1}^{XXX}) + s^{XXX} \sqrt{r_{t-1}^{XXX}} \varepsilon_t^{r,XXX}$$

	USD	CHF	GBP	JPY
a	0.29157	0.24343	0.18029	0.02769
b	0.05	0.035	0.06	0.03
s	0.04202	0.07029	0.05072	0.04492
r₀	0.06325	0.06004	0.06004	0.00341

Correlation matrix for residual terms ε :

	USD	CHF	GBP	JPY
USD	1	0.31022	0.55942	0.43474
CHF	0.31022	1	0.64407	0.63582
GBP	0.55942	0.64407	1	0.66107
JPY	0.43474	0.63582	0.66107	1

ii.) Inflation generator

$$i_t^{XXX} = \beta_0^{XXX} + \beta_1^{XXX} r_t^{XXX} + \varepsilon_t^{i,XXX}$$

	USD	CHF	GBP	JPY
β_0	0.00584	-0.006662	-0.02255	-0.0008
β_1	0.41237	0.722173	0.76457	0.40863
$\sigma(\varepsilon)$	0.011053	0.010601	0.013627	0.0096097

iii.) FX rate generator

$$s_t^{USD/XXX} = \alpha_0^{XXX} + \alpha_1^{XXX} r_t^{USD} + \alpha_2^{XXX} i_t^{USD} + \alpha_3^{XXX} r_t^{XXX} + \alpha_4^{XXX} i_t^{XXX} + \varepsilon_t^{s,XXX}$$

	CHF	GBP	JPY
α_0	-0.1666	0.52786	0.17831
α_1	-7.3806	0.28439	-8.71995
α_2	7.3806	-0.28439	8.71995
α_3	1.8608	-1.41399	-5.97009
α_4	-1.8608	1.41399	0
$\sigma(\varepsilon)$	0.14102	0.10203	0.13337

Notice that for CHF and GBP we have applied the simplifications $\alpha_1 = -\alpha_2$ and $\alpha_3 = -\alpha_4$ as they led to significantly better fits, i.e. the FX rate is modelled by the real rates of return of the involved countries, and the effective number of parameters to be estimated is reduced. For JPY, this simplification turned out to be applicable only for US inflation and interest rate.