POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

EXECUTIVE SUMMARY OF THE THESIS

Modelling correlation in foreign exchange markets

LAUREA MAGISTRALE IN MATHEMATICAL ENGINEERING - INGEGNERIA MATEMATICA

Author: FRANCESCO PASCUZZI

Advisor: PROF. MATTEO BRACHETTA

Co-advisors: MATTEO CERUTTI, GABRIELE MINEO

Academic year: 2022-2023

---

1. Introduction

Developing, and testing, a reliable model for correlation can be an extremely valuable project for any derivatives trading desk. It is a necessary task if one wants to trade any multi-asset exotic option, as the pricing model will need to account for the joint dynamics of the assets, but it can be difficult to find market instruments that can be used to calibrate correlations.
Even when basket instruments exist, and are liquidly traded, they often depend on the aggregate correlation between $N > 2$ assets and can not be used to calibrate the correlation between just two arbitrary assets. In this regard, foreign exchange (FX) markets luckily represent one of the few exceptions, since every FX rate can be seen as a derivative on two (different) FX rates.

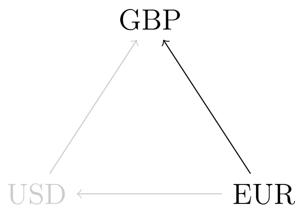

Figure 1: EURGBP represents the amount of Great British pounds (GBP) required to purchase one euro (EUR).

Definition 1 (FX rate). *The (spot) FX rate S_t = CCY1-CCY2 is the number of units of currency CCY2 required to buy one unit of currency CCY1 in a guaranteed currency exchange occurring on the spot date (1 or 2 business days after the trade date).*

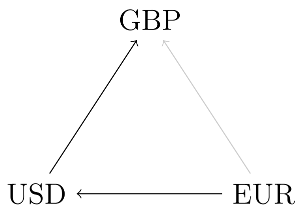

Figure 2: It is also "a derivative" on FX rates against a third currency like the U.S. dollar (USD).

The ability to look at FX rates through two different lenses, either as a single asset or as a derivative on two other assets, derives from the no-arbitrage relationships that prevent investors from getting a better FX rate by passing through a third currency.
When three currencies are considered together with their three FX rates, they are called an FX triangle and they are the key to modelling correlation in this market. In particular, when looking at an FX triangle it is common for two

FX rates to be traded more commonly than the third, these two rates are often called the *majors* (denoted with $S^{(1)}$ and $S^{(2)}$) while the third rate is called the *cross* (denoted with $S^{(12)}$). Using EURUSD and GBPUSD as the majors, and EURGBP as the cross, the value of the latter is given by:

$$S^{(12)} = \frac{S^{(1)}}{S^{(2)}} .$$

On its own, this relationship does not provide any insight into the correlation ρ between the majors, but this is solved by looking at the variances Σ of the log-returns of the three rates:

$$\Sigma^{(12)} = \Sigma^{(1)} + \Sigma^{(2)} - 2\rho\sqrt{\Sigma^{(1)}\Sigma^{(2)}} .$$

1.1. Objective

The objective of the thesis is to implement and test a correlation model for pricing multi-FX derivatives. From a review of the literature and numerical results obtained in the thesis, it is clear that a constant correlation model is not able to describe correlation in FX markets. The most complete theoretical framework for correlation would treat it as its own stochastic process, correctly reflecting the ever-changing perceptions of correlation from market participants. However, proceeding with the same line of reasoning that led to the creation of local volatility from stochastic volatility, we adopt the local correlation (LC) model proposed by Guyon [1] to obtain a non-stochastic model that, thanks to Gyöngy's theorem, replicates the option prices generated by a stochastic one.

The LC model is tested from the point of view of a trading desk, analysing the main points that must be considered before adopting a new market model in production. Among those points, the first question that must be considered is whether the model is able to replicate the calibration instruments. In the case at hand, this means simultaneously replicating any set of arbitrage-free plain vanilla options for all three FX rates in the triangle. A model that satisfies this requirement is called *admissible*.

Unfortunately, to the best of our knowledge, there is no theoretical guarantee that the LC model will always be admissible (provided that the input data is arbitrage-free). As an alternative to theoretical guarantees, however, the model can be tested on empirical data that spans a wide range of market conditions to infer the likelihood of the LC model being inadmissible.

After having confirmed that the model is admissible often enough to be considered useful, the following points pertain to the quality of the model. A model can be considered satisfactory if it offers fast calibration and pricing, produces prices for exotic options that are in line with the rest of the market, and does not display excessive numerical instability in the Greeks. These requirements are necessary for a trader in order to always have a perspective, on both prices and exposures to market risks, that is aligned with the rest of the market.

1.2. Methodology

To observe the performance of the model in a wide range of market conditions, we chose to apply it to three different FX triangles. The analysis can not be carried out on a single FX triangle, because different FX markets can display radically different volatility surfaces and the results on one triangle would only apply to other triangles with similar volatility dynamics. In particular, to encompass the widest possible range of volatility dynamics, the chosen FX triangles are the following:

1. euro, U.S. dollar, and Great British pound (EUR-USD-GBP);
2. euro, U.S. dollar, and Japanese yen (EUR-USD-JPY);
3. euro, U.S. dollar, and Mexican peso (EUR-USD-MXN).

This allows the tests to consider three different scenarios for the cross: a scenario with low volatility and low skew in the EUR-USD-GBP triangle, a scenario with higher volatility and high negative skew in the EUR-USD-JPY triangle, and a scenario with higher volatility and high positive skew in the EUR-USD-MXN triangle. While there are structural reasons for these volatility shapes[1], it should be pointed out that these reasons are susceptible to changes and the descriptions of these triangles only reflect the

[1] EURGBP has low volatility because of the tight economic link between the United Kingdom and the European Union, EURJPY has negative skew because of the unconventional interest rate policy pursued by Japan, and EURMXN has positive skew because of the perceived risk of a sudden depreciation of the peso.

market as of 2023.
For each of the three triangles, options data was manually collected on every active trading day from March 29, 2023, to May 5, 2023. Choosing to perform the analysis on multiple days allowed us to examine how the model performs even in stressed market conditions, for example following large moves in the spot market (e.g. USD-MXN on April 5) or before economic events (e.g. Bank of Japan meeting on April 28, or Federal Reserve rate decision on May 3).
Then, the quality of the model is examined on a case study where it is used to price and measure the exposure of a dual range accrual. The (non-dual) range accrual is an exotic option that is quite common with investors, it usually consists of a note that accrues interest when the underlying FX rate stays inside a predetermined range. The dual version adds a second FX rate and accrues interest only when both FX rates stay inside their respective ranges, this is usually done to reduce the price of this option as it lowers the probability of collecting interest.

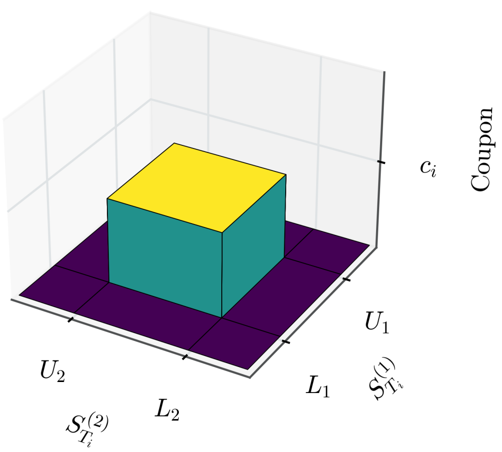

Figure 3: A dual range accrual with ranges $[L_1, U_1]$ and $[L_2, U_2]$ can be seen as a strip of late-delivery European options with the pictured payoff.

2. LV-LC model

The local correlation (LC) model [1] assumes that the instantaneous correlation ρ is a function of the spot FX rates:

$$\mathrm{d}\langle W_t^{(1)}, W_t^{(2)}\rangle = \rho\left(t, S_t^{(1)}, S_t^{(2)}\right)\,\mathrm{d}t \qquad (1)$$

On its own, the LC model does not make any claim regarding the single-asset dynamics of the majors but is only concerned with the dependence between their driving Brownian motions. In any case, the most natural choice for single-asset dynamics is the well-known local volatility (LV) model because of its simplicity and its ability to replicate arbitrary volatility smiles without the introduction of non-traded sources of risk (e.g. jumps, stochastic volatility, ...).
While it is possible to use more complex models (e.g. stochastic volatility) with local correlation, the LV-LC combination is the simplest one and also represents the industry standard for most use cases that do not involve path-dependent payoffs. Thus, each major follows the dynamics below:

$$\mathrm{d}S_t = \mu_t S_t\,\mathrm{d}t + L(t, S_t) S_t\,\mathrm{d}W_t\,. \qquad (2)$$

On the other hand, once the dynamics of the cross are determined by (2) and the correlation is determined by (1), the dynamics of the cross do not have any residual degree of freedom as the cross is a deterministic function of the majors. In particular, since it is not possible to know the value of $\left(S^{(1)}, S^{(2)}\right)$ from $S^{(12)}$, the volatility V of the cross will not be measurable with respect to $S^{(12)}$ making it a stochastic, rather than local, volatility.
This means that the LV-LC model is guaranteed to replicate the volatility surfaces of the majors (thanks to the LV model) but it must be calibrated in order to match the volatility surface of the cross. This calibration involves finding the solution of a McKean's stochastic differential equation, and it is possible to approximate this solution with the very efficient particle method developed by Guyon and Henry-Labordère [2]. The efficiency of this method derives from the fact that it is able to calibrate the local correlation using the same Monte Carlo paths that will then be used for pricing.
During the calibration procedure, inadmissible LC models can be identified from the fact that they produce correlations that assume values outside $[-1, 1]$. It is possible to still use those models in practice (e.g. by capping the correlation to $[-1, 1]$) but they will not perfectly replicate the volatility surface of the cross.

2.1. Local in cross representation

The most important contribution to local correlation made by Guyon is, arguably, the introduction of the *local in cross* representation of $\rho\left(t, S_t^{(1)}, S_t^{(2)}\right)$. In fact, the volatility surfaces of all three FX rates of a triangle are not sufficient to identify a single local correlation function ρ, there could be an infinite number of possible functions that exactly replicate all three volatility surfaces. Guyon solved this problem by introducing two arbitrary functions, a and b, to obtain the following representation:

$$\rho_{(a,b)}\left(t, S_t^{(1)}, S_t^{(2)}\right) =$$
$$a\left(t, S_t^{(1)}, S_t^{(2)}\right) + b\left(t, S_t^{(1)}, S_t^{(2)}\right) f\left(t, \frac{S_t^{(1)}}{S_t^{(2)}}\right)$$

After choosing a and b, the local in cross function f is the only free parameter that controls the correlation, greatly simplifying the calibration. However, the choice of (a, b) is up to the researcher and different choices will lead to different prices for exotic options.
Except for two special cases, the model implemented in the thesis always chooses (a, b) to be equal to $(0, 1)$, the so-called Reghaï local correlation, as it is a smooth function that can be plotted as a surface and inspected for potentially critical points.

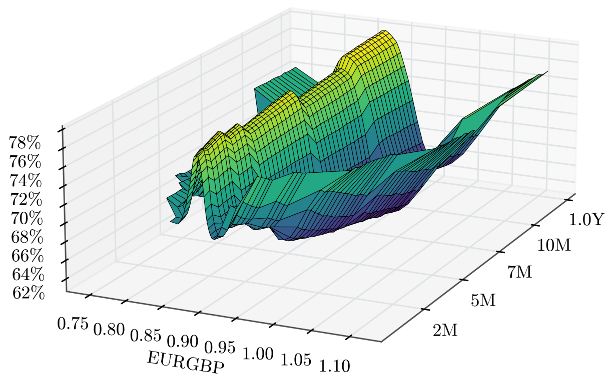

Figure 4: Reghaï local correlation for EUR-USD-GBP (March 29, 2023).

3. Calibration results

Throughout the entire observation period, the LV-LC model has been able to (almost perfectly) replicate the volatility surfaces of all three crosses. As shown in Figure 5, the additional complexity introduced by LC over a constant correlation is an acceptable cost for the vastly superior fit obtained on the cross. Simpler models, like constant correlation, are often unable to replicate the skew and convexity of pairs like EURJPY and EURMXN.
Although LV-LC still has some replication errors, these are caused by the necessary numerical approximations made during both calibration and pricing. In any case, the magnitude of these errors (almost always below 20 basis points) is almost insignificant compared to the bid-ask spread of the options market.

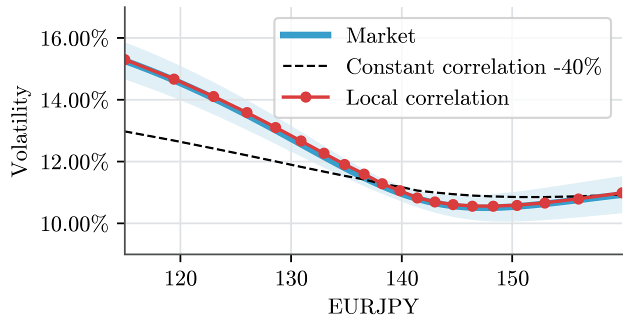

Figure 5: Replication of the one-year smile of EURJPY (March 29, 2023).

The accurate replication of the cross smile for all three triangles on all observation days already suggests that the correlation function should be admissible in most market conditions.

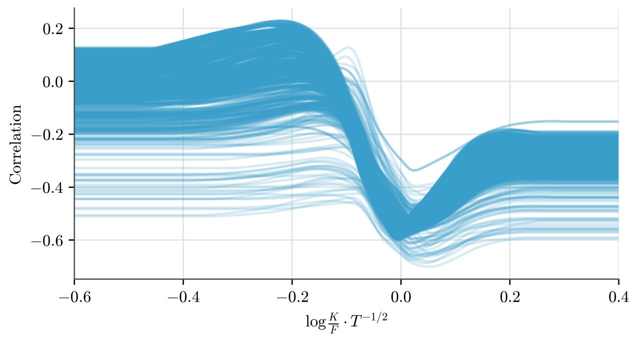

Figure 6: Values assumed by the Reghaï local correlation for EUR-USD-JPY on all maturities and all observation dates.

Nevertheless, it can be useful to examine the range of values assumed by the correlation to check if it ever gets dangerously close to 100% in absolute value to obtain some intuition concern-

ing how close the local correlation is to becoming inadmissible (recalling that inadmissibility $\iff$ correlation assumes values outside $[-1, 1]$). Figure 6 shows two important features of the local correlation model. The first feature is that the local correlation function never exceeded values of 70% in absolute value, meaning that it always remained reasonably far from inadmissibility. The second feature is that the shape of the local correlation function for a fixed t remains stable across maturities with sticky-delta dynamics, as all lines show a high degree of overlapping when plotted against the maturity-adjusted log-moneyness $\log \frac{K}{F} \cdot T^{-1/2}$.
These observations also hold for the other two triangles (not pictured in this summary), except for the EUR-USD-GBP correlation that exhibits correlations as high as 90%. However, the extreme values of EUR-USD-GBP do not raise any major concerns, as the triangle is known for having high (albeit stable) correlations.

4. Case study results

The case study in the thesis considers a (slightly simplified) FX dual range accrual with the following details.

- **Underlying:** EURUSD as $S^{(1)}$ and GBPUSD as $S^{(2)}$.
- **EURUSD range:** $[L_1, U_2] = [1.10, 1.20]$.
- **GBPUSD range:** $[L_1, U_2] = [1.15, 1.30]$.
- **Payment dates:** coupons are fixed and paid every quarter for one year (i.e. $\{T_i\} = \{3M, 6M, 9M, 1Y\}$).
- **Coupon:** 5.9% of the notional is paid in USD at each fixing date (contingent on the accrual condition being satisfied).

The main price is obtained with Reghaï local correlation, while two other prices are obtained with bespoke choices of (a, b) that were created with the sole purpose of maximising and minimising the value of the dual range accrual. The comparison between the three prices in Table 1 offers some insight into the degree of uncertainty introduced by the incompleteness of the multi-FX market.

Reghaï	Best case	Worst case
6.23%	6.35%	6.19%

Table 1: Price of the dual range accrual (as percentage of notional) under different local correlation functions.

The small impact on price is consistent with the results obtained by Guyon in [1] and confirms that the choice of (a, b) is not particularly critical (at least for this specific derivative on this specific triangle).
After having obtained a price, the following step for a trader is computing the Greeks in order to manage the risks generated from trading this exotic product. In particular, we focus on spot Greeks (i.e. deltas and gammas) as they can often be unstable when prices are computed with Monte Carlo simulations.
As shown by Figure 7, a naïve finite differences approach on the LV-LC model generates unacceptable numerical instability.

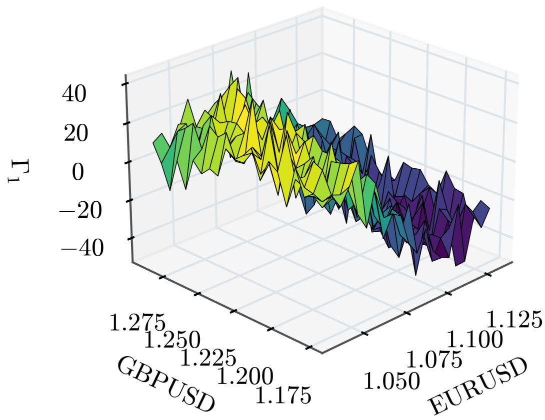

Figure 7: Dual range accrual EURUSD gamma computed with finite differences.

To use the model in practice, Greeks can be smoothed with the Chebyshev's approach proposed in 2021 by Maran, Pallavicini, and Scoleri [3].

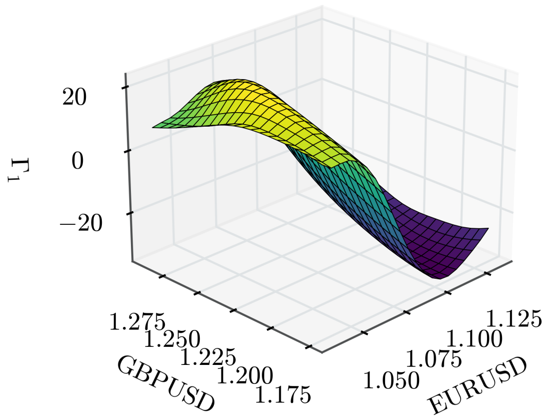

Figure 8: Dual range accrual EURUSD gamma computed with Chebyshev's approach.

This method computes the price of the derivative on Chebyshev's knots and uses a Lagrange interpolant to obtain a smooth Greek without introducing bias.

5. Conclusions

Empirical evidence on three extremely different scenarios shows the robustness of the LV-LC model on a wide range of market conditions, while also showing that calibration errors deriving from the numerical implementation are well within market tolerance. In addition, despite the LV-LC model naturally introducing additional complexity compared to a constant correlation approach, the computational costs are more than justified by the inability of constant correlation to replicate the calibration instruments.
The case study shows that, thanks to the particle method and Chebyshev's approach, it is also possible to obtain fast and accurate Greeks from the model.
These results strongly support the viability of the LC-LV model for multi-FX exotic derivatives.

6. Acknowledgements

This thesis was developed with the FX Derivatives Trading team at Intesa Sanpaolo. I would like to thank Matteo Cerutti, Gabriele Mineo and Samuele Gatti for the support and thought-provoking conversations that made this thesis possible.

References

[1] Julien Guyon. A new class of local correlation models. *Available at SSRN 2283419*, 2013.

[2] Julien Guyon and Pierre Henry-Labordère. The smile calibration problem solved. *Available at SSRN 1885032*, 2011.

[3] Andrea Maran, Andrea Pallavicini, and Stefano Scoleri. Chebyshev greeks: Smoothing gamma without bias. *arXiv preprint arXiv:2106.12431*, 2021.