LUISS

Dipartimento di
Economia e Finanza

Cattedra di Probability

Stochastic Skewness in the FX Market

Prof. Sara Biagini
SUPERVISOR

Prof. Valerio Marchisio
CO-SUPERVISOR

Emanuele Li Volsi - 742551
CANDIDATO

ANNO ACCADEMICO 2021/2022

Table of Contents

Abstract

In this paper I calibrate a jump diffusion model to dynamically extract log returns from exchange rate levels and one-month risk reversal quotes. The model is made up of a deterministic part representing the spread between overnight interest rates of the two countries and a stochastic part divided in two Levy processes to discern between up and down jumps. The estimation period covers 2021's daily observations and is mainly described by economic recovery, covid-19 spikes and risk sentiment swings reflected in both currency spot and options short-term trends. A fine connection is spotted between risk reversal jumps and spot levels, which corroborates the model's capability for both estimation procedures. The aim of this study is to point out the importance of skewness when accounting for currency returns and option pricing, more than in the equity counterpart. The work is divided in two big sections, namely a literature and an empirical analysis part. The former starts by making an introductory presentation of market returns studies and the importance of skewness for investment purposes; secondly the focus shifts to the main features of the currency market, economic drivers and market quotes used to spot skewness, also from a more practical perspective. The empirical analysis section introduces the model of choice, inspired by Carr et al. (2007), and points the attention on the peculiar features of such structure. Then the calibration procedure is thoroughly detailed, together with the chosen optimization procedures and obtained results. Lastly in the conclusions section it is summarized the whole study process and some suggestions are made for further developments.

Key words: jump diffusion, currency returns, stochastic skewness

1 Skewness and Market Returns

It is widely known among both practitioners and academics that currency returns have been historically described by a leptokurtic distribution (fat tails), together with some degree of asymmetry mostly connected with fundamental factors such as monetary policy or economic growth. Several studies have indeed incorporated skewness and other higher moments factors into their risk premia analysis, suggesting that currency, and not only, returns can be significantly explained by such distribution anomalies. This evidence is reflected into another finding which points to investors preference for positive skewness in their asset allocation (Brunnermeier and Pedersen (2009)). For instance, Sokol and Eguren-Martin (2020) have shown that countries with a low level of net foreign assets and high interest rates are relatively more exposed to currency depreciation. More interestingly for the purpose of this paper, Iseringhausen (2021) develops a model for time-varying skewness and finds out that idiosyncratic skewness related to currency returns varies significantly over time and is negatively related to the carry trade strategy, which entails going long high interest rates currencies and short lower-yielding currencies. Additionally, Brunnermeier et al. (2008) demonstrate that traders' funding constraints are a reasonable explanatory factor for the reduction of investment in high-yielding currencies, which explains why traders require a higher premium for holding such currencies, especially during periods of liquidity contraction. On the other hand, the two more used risk premia factors for currency returns, explained by fundamental theory, are the interest rate differential and the real exchange rate. Against the uncovered interest parity equation, literature based on Fama regressions have shown that higher yielding currencies tend to appreciate in the short run. Alina Steshkova (2021) with a study based on the above two fundamental factors, finds out that interest rate differential is negatively bounded to skewness, which again explains the traders' funding constraint concept. The real exchange rate has instead been found to have a somewhat monotonic relationship with the time horizon, i.e. it can be negligible to predict short-term risk premia and vice versa (Dahlquist and Penasse (2017)). All in all, skewness can be observed both in the short and long-run, thus one useful assumption is to see currency returns as described by a skewed t-distribution (Steshkova 2021). Accounting for skewness can further simplify not only empirical studies aimed at comparing investment strategies and asset allocation, but also the calibration of risk management metrics, namely Value-at-Risk and Expected Shortfall. Giot and Laurent (2003) provide evidence that including higher moments factors can improve the accuracy of downside risk analysis, as compared with computing such metrics using only mean and variance (gaussian distribution) .

2 Currency Market
2.1 Framework

FX market drivers can be mainly divided into three categories, namely the stochastic movement of the exchange rate, captured by the most used Black & Scholes model, then the stochastic behavior of volatility, reflected for example in Heston's and Bates' models through specific diffusion processes introduced above, and the observed skewness of the log returns' risk neutral distribution. The latter is priced by the market via the so called risk reversal, which is given by the difference in implied volatility between an out-the-money call option and an out-the-money put option contract. For consistency these contracts can be compared by looking at the same level of moneyness, which is practically linked to the delta, i.e. the probability that an out-the-money contract will expire in-the-money. The most conventional measure used by market participants is the 25 delta risk reversal which, given its high degree of tradability, can be perceived as a gauge of positioning in the currency. However, some currency pairs may be more resilient than others and react differently to market expectations, for instance because of a traded volume argument. Over 2021, to make an example, both the EUR and the GBP downward movements against the US dollar have reflected market worries with respect to stagflation signals, even though the EUR has sometimes proved to have a more solid support given its ample liquidity (and fiscal stimulus).

Over the decades, both academics and practitioners have studied and implemented several approaches to efficiently trade exchange rates, mainly from an hedging perspective. Both macro and technical factors have played a role in such task, even though the former resulted to be more reliable for longer-term estimates, whereas the latter has been widely used by traders and market makers to predict intra-day or daily exchange rate estimates. Fair value models, or more generally tools based on economic variables such as interest rate differentials, growth estimates or monetary policy, are widely used by institutional investors to gauge the market over/underreaction compared to the actual value of the exchange rate. It can be of great help from a positioning point of view, when accounting for currency allocation in a multi-diversified portfolio. Corporate flows, on the other hand, can indicate some specific interest with respect to one currency against the other, usually related with forward expectations related to data releases, macro events or even some seasonal portfolio adjustments, usually observed on a quarterly basis or towards the end of the month, in the form of a rebalancing market price action. Technical indicators can instead be very useful to identify shorter-term trends in the exchange rate level, being such indicators based on variables like traded volume, moving-averages and so forth. Moreover, they can be combined with economic variables to match a subjective belief about some specific trend or, more simply, for the best timing in terms of trade execution.

2.2 Skewness and OTC option quotes

Option prices have been widely used to describe the distribution of market returns, however it must be reminded that options are able to retrieve the risk-neutral probability density function (pdf) and not the actual density of such returns. Still, options are forward-looking instruments able to incorporate information that may only be relevant much later into the exchange rate's trend. Garch models, where volatility is assumed to follow a stochastic process, do not update their parameters often enough to reflect major regime changes while on the other side, Campa, Chang and Reider (1997) have proven that currency options are highly sensitive to market news and thus to sudden changes in the exchange rate level, which might affect the overall trend structure. Such big changes in the pdf can be attributed for example to unexpected uncertainty over an election programme or a monetary authority decision. The main advantage of using an option-based approach to retrieve market returns is that it does not depend on any specific functional form, which makes it a good fit to cope with different environments such as switching regimes, target zones and so forth. Empirically, it has been spotted a significant relation between the spot rate level and the degree of risk-neutral skewness, suggesting that exchange rate expectations (implied in options) are indeed "extrapolative", hence the connection between a strong currency and high probability of future appreciation.

In the over-the-counter market options are more liquid and thus more competitive relative to exchange-traded options, in particular the daily turnover and notional amount outstanding have historically been about 10% lower than its OTC counterpart. With regards to the latter, quotes are expressed in terms of implied volatility through which traders retrieve, via the Garman-Kohlhagen formula (i.e. the Black-Scholes adjusted for the foreign interest rate), the option premium. In this way option quotes do not require constant updating given that this approach implies a one-to-one mapping between implied volatility, representing traders' assessment of future changes in the underlying asset, and option prices. On the contrary, if prices were expressed in terms of strike price, it would fairly require greater connection between spot and option markets.

Dealing with over-the-counter option quotes, there are several approaches used to compute the degree of skewness. Firstly, skewness can be estimated via the probability density function (pdf), for which approach Campa et al. (1997) focused on the well-known trimmed binomial tree approach to derive the density and thus the skewness. Secondly, the level of asymmetry can be computed by looking at relative intensities around the tails. In particular, it can be compared the density function below and above a certain critical level as follows:

$$intensity(\overline{S}) = \int_{\overline{S}}^{\infty} (S_T - \overline{S}\ ) f(S_T) dS_T$$

Where $f(S_T)$ is the risk neutral density and the equation expresses the sum of any appreciation above a certain exchange rate multiplied by the probability that such realizations will take place, whilst on the other side:

$$intensity(S\_) = \int_{0}^{S\_} (S\_ - S_T\ ) f(S_T) dS_T$$

By the difference between the two intensities it can be computed the relative skewness of the distribution given that, by looking at only one of the tails the analysis could merely point to a leptokurtic distribution, i.e. described by fat tails. Accordingly, the two thresholds are chosen symmetrically around the forward rate, otherwise the relative intensity would always equal zero given that, as discussed above, the distribution's mean must always equal the forward rate. The magnitude of such appreciations is usually chosen as a function of volatility, so that changes in skewness are independent from changes in realized volatility. Similarly, even if not equally used, the total probability can be measured around the threshold levels and then obtain the difference. However, given that probability is centered around the median rather than the mean, for a positive skewed distribution the total probability would be greater below the forward rate, or mean, but the further the threshold is brought from the mean the more likely is the relative probability to switch sign.

Lastly, for OTC currency options, there is a specific market aimed at providing an estimate of skewness related to the price of two option contracts with symmetric features. More specifically, the so called 'risk reversal' denotes the difference in price between a call and a put option with the same level of moneyness, where the latter expresses the likelihood of such contracts of expiring in-the-money. The general idea can be reconnected to the relative intensity approach, even though for this case the threshold is chosen by the delta. In the market, 25 delta contracts are the most traded as opposed to 10 delta for instance, meaning that the mostly traded OTC options for risk reversal strategies are contracts with a risk-neutral probability of 25% of expiring in-the-money. This approach can be recognized as more reliable given that it reflects market sentiment and it is not subject to estimation errors, given by smoothing procedures used to compute the probability density function. More interestingly, given that over-the-counter options contracts are priced in terms of Black and

Scholes implied volatility, the difference in premium paid between call and put options simply reflects the spread in terms of volatility. It is important to remind that, by the Black and Scholes assumptions, the volatility surface across a series of strike prices should be flat, making risk reversal quotes approximately zero. Also, with stochastic volatility models like Heston's (1993), a random change in volatility is associated with an increase in the price of contracts far-from-the-money as opposed to contracts with a relatively higher intrinsic value, thus merely pointing towards a change in kurtosis.

2.3 Empirical findings

For this case study, attention will be put on risk reversal quotes, giving a direct market sense of one currency appreciation against one or more others. Risk reversals are expressed in terms of the 'quote currency', which in the case of GBPUSD determines the skewness towards an appreciation of the US dollar. Empirically, the correlation between skewness and exchange rate level is based on the assumptions regarding the exchange rate dynamics. For instance, in a target zone regime, when the level stays closer to the floor of the range, the skewness would be relatively more positive and vice versa when the level trades around the ceiling. Similarly, by looking at long-term fair value estimates, a depreciation of the spot level would be more likely when the latter sits above fair value, making the relative skewness 'more' negative. Conversely, when assuming a random walk process, the correlation between spot level and skewness should be zero since the relative skewness would be theoretically absent. When assuming volatility to have a stochastic process, it can be shown that the Black & Scholes implied volatility has its minimum at-the-money forward, i.e. when the strike price equals the forward rate, whereas it increases the further the option goes out-the-money. This pattern is most commonly known as the "volatility smile" but the shape can vary, quite frequently for currency options, depending on the strength of one currency against the other. The same pattern can be shown using either calls or puts given that, by assuming the put-call parity, the implied volatility is the same for any level of moneyness.

Furthermore, in relation with the observed volatility smile, skewness has empirically resulted from a negative correlation between price level and volatility shocks, as discussed by Hull and White (1987) and Heston (1993). Unlike in equity markets, as shown by Jackwerth and Rubinstein (1996), skewness in the currency market has been relatively less systematic over time. Studies on dollar crosses during mid-1980's have found negative skewness on the back of an unusually overbought greenback, which would coincide with mean-reversion. More specifically, positive skewness would indicate that the right tail of the distribution is fatter, meaning that large positive realizations are more probable than a depreciation. However, it is important to remember that the mean of the risk-neutral distribution must always be the forward rate, thus broadly speaking a positive risk-neutral skewness does not imply that an appreciation is more likely than a depreciation. An interesting interpretation of such risk-neutral based concepts provided by Campa, Chand and Reider (1997) is that "markets do not necessarily view large appreciation of a strong currency as statistically more likely; they may simply attach a higher valuation to states of large appreciation".

Going back to risk reversal quotes, by the Black and Scholes formula, calls and puts with the same level of moneyness, or delta, should have the same sensitivity to changes in volatility. Hence option traders would not need to hedge vega when buying or selling risk reversals. However, to compute the actual level of volatility for out-the-money options it is useful to look at strangles quotes as well, i.e. contracts made of an out-the-money call plus an out-the-money put with the same level of moneyness. By looking at both contracts it is possible to retrieve a more accurate price for each individual option contract. More interestingly, Bertola and Caballero (1992) showed how the positive relation between spot and skewness reflects the economic interpretation that exchange rates follow target zones. On this respect, when the spot rate gets closer to the ceiling, the probability of an upward realignment should be greater and vice versa, as well as support and resistance levels within the technical analysis framework. Moreover, traditional currency models based on interest rate parity have been empirically found less reliable over the short term. On that respect, the so called 'carry trade' strategy has been widely implemented to anticipate an appreciation of the currency caused by a spike in interest rate (opposite idea of the interest rate parity). Because of the success of such strategy over the years, currencies have started to be categorized according the appeal of their respective yield curve. For instance, currencies like the US dollar, the Canadian dollar or the New Zealand dollar have been classified as high-yielding, whilst the Japanese Yen or the Swiss Franc as low-yielding or funding currencies. The latter name has been reasonably attached to these currencies given that, according to the carry trade strategy, the investor would borrow funds in the low-yielding currency and invest in some higher-yielding currency to make a profit between lending and borrowing costs. The chosen time period for this case study (2021) has seen a rise in those strategies, given that a global economic recovery has pushed central banks and governments to reconsider their monetary policy and fiscal stimulus. Accordingly, countries like the United States or New Zealand have perceived a more hawkish change of policy, to be reflected in a rise in interest rates, unlike the Japan whose central bank has always been known for its low interest rates policy. However, such strategies have been empirically found to last for a certain amount of time followed by a so called 'policy normalization' period, when all countries tend to realign in terms of monetary policy. All in all, a plurality of factors have pushed towards the acceptance of a more stochastic trend in terms of currency pairs, to be backed by either macro-economic or technical factors which may indeed affect the overall skewness of currency returns.

Another interesting study by Brunnermeier et al. (2008) explains how carry trades might be related to currency crash risk, and thus to the degree of skewness. In particular, they state that the carry trade strategy delivers negatively skewed returns given that sudden abandonment of these strategies can be positively linked to currency crashes. Moreover, these kind of market events can be related to shocks in overall volatility, which might force traders to set higher margins to overcome the illiquidity issue, hence the unwinding of carry trades. These findings support the view that macroeconomic factors can severely affect currencies' interest rate differential, while funding constraints and volatility spikes can instead be linked to the likelihood of a currency crash. Following the same idea of currency crashes, Farhi and Gabaix (2008) worked on a model that studies the link between currency forward premium and the sensitivity to global economic crunch periods. Unlike the study of Brunneimeier, Farhi et al. point the attention on more exogenous productivity shocks, while still accounting for skewness patterns observed in currency option quotes. Galati et al. (2007) focus instead on net banking flows between countries, which can be fairly useful to gauge the degree of carry trade activity. All in all, currencies have been found to have similar patterns when accounting for similar yield curves, and the relation between spot level and skewness or between carry trade and liquidity constraints can be further useful to anticipate spikes in the forward premium.

Finally, Sunjin Park (2016) carries on a case study aimed at predicting carry trade returns based on changes in conditional expected growth and global conditional skewness. The model is built on agents with heterogeneous beliefs about growth rate expectations, according to which model the degree of skewness can vary quite frequently. The latter is seen by Park as a global risk factor that affects the discounting factor of each country and thus the exchange rate level. Unlike disaster risk, skewness has the benefit of accounting for both positive and negative directions of asymmetry, which in a currency trade environment can result in a more parsimonious model able to measure not only unexpected shifts related to bad news but also extreme positive events which could weigh on one currency strength.

3 Modelling

3.1 Stochastic Volatility

Empirical studies on volatility patterns have shown that, unlike traditional option pricing models that assume constant volatility over the life time of a contract, the implied volatility increases as the underlying asset moves more in-the-money or out-the-money. This simply contradicts the assumed flat volatility surface plotted against a series of strike prices (moneyness), and suggests a more realistic u-shaped function of moneyness with regards to volatility implied into option prices. The latter would thus suggest that the risk-neutral conditional distribution of market returns is fat tailed. Accordingly, second generation models have successfully considered such behavior when it comes to pricing option contracts. Namely the Heston stochastic volatility model (1993) is able to capture not only the volatility smile but also the stochastic variation in the implied volatility level for a given degree of moneyness, or the jump-diffusion stochastic volatility model of Bates (1996) which captures the same features through the presence of a jump parameter. These kind of models have unfortunately been unable to generate relevant time variations in the risk-neutral skewness attached to returns. The latter is a feature which can be practically recognized as deterministic in some kinds of option markets, like equity. In particular, equity options returns have been historically described by negatively skewed distributions, on the back of the so called "crunch premium" which reflects higher prices related to bottom-side hedges as opposed to top-side bets.

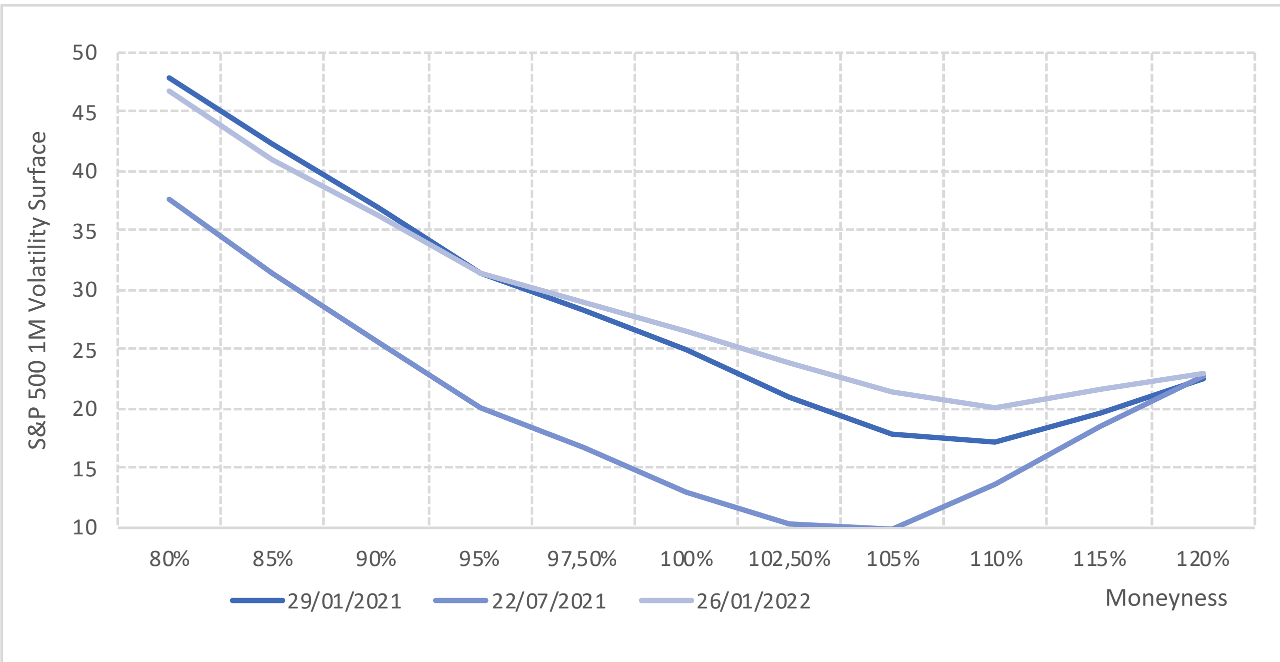

Figure 1: S&P 500 one-month call options implied volatility surface. *The three lines represent three different dates, to depict the surface change on a six-month basis. The shape of the surface slightly*

changed over the year, maintaining the so called "smirk" form. Interestingly the minimum of the curves stays systematically closer to the in-the-money zone.

The above chart depicts the latter concept indeed, despite a change in the relative price of put and call options over six-month period, the S&P 500's volatility surface remains of the same smirk shape. The slope of the curve is merely changing between time periods, which indicates a fairly preference from equity investors for put options, as a form of hedging strategy. Currency market returns are instead described by strong time variation in risk-neutral skewness, which has been found not only to vary stochastically but also to change sign depending on option markets expectations about future exchange rates movements. In particular, some exchange rates have historically shown some relevant trend moves right after such "psychological" changes in skewness, to go hand in hand with significant market events expected within the life time of the option contract, such as expectations of rate hikes or the start of a tapering cycle. For instance an interesting currency pair, analyzed by Carr and Wu (2007) with respect to skewness behavior affecting changes in price, would be the USDJPY. The latter has mainly shown negative skewness historically, to be found in a preference for JPY OTM calls, given its safe-haven flavor even relative to the highly liquid greenback during risk-off periods of time. However a few changes in the skewness sign have been followed by significant spikes in the exchange rate, via a relative increase in the demand for USD OTM calls.

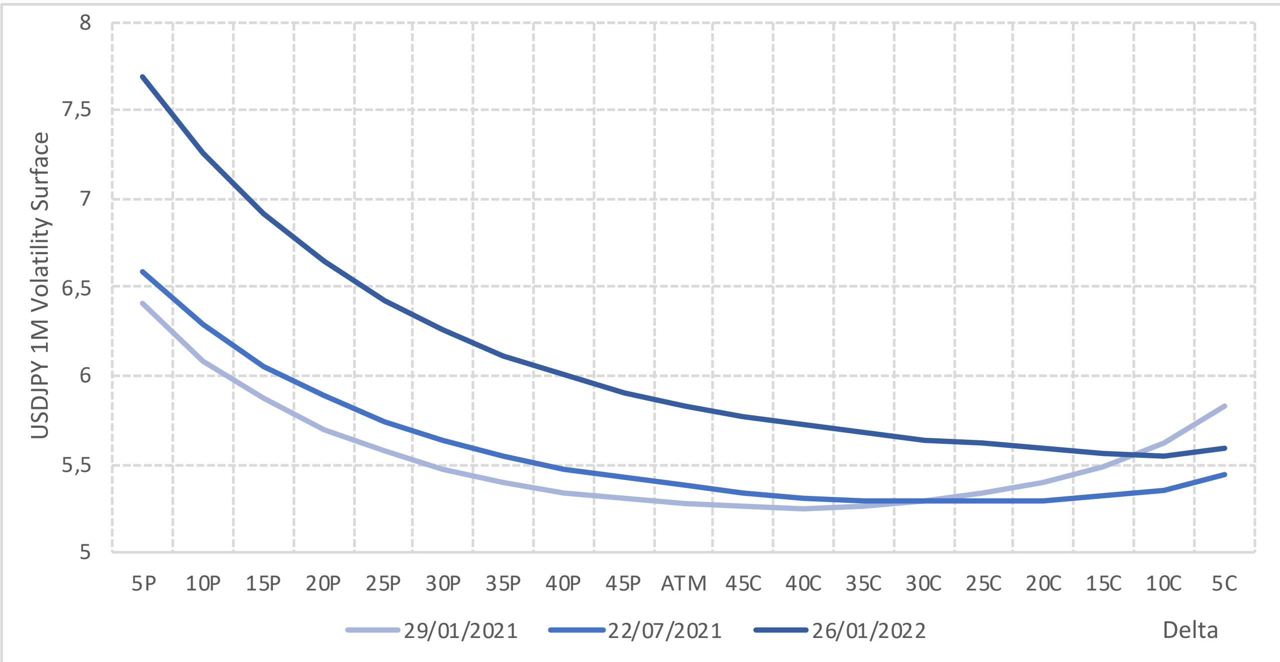

Figure 2: USDJPY one-month implied volatility surface. The three lines represent three different dates, to depict the surface change on a six-month basis. It is noticeable how in January 2021 there

was a somewhat symmetry between out-the-money call and put options, whereas the asymmetry increased in negative sign over the year, favoring out-the-money put options' implied volatility.

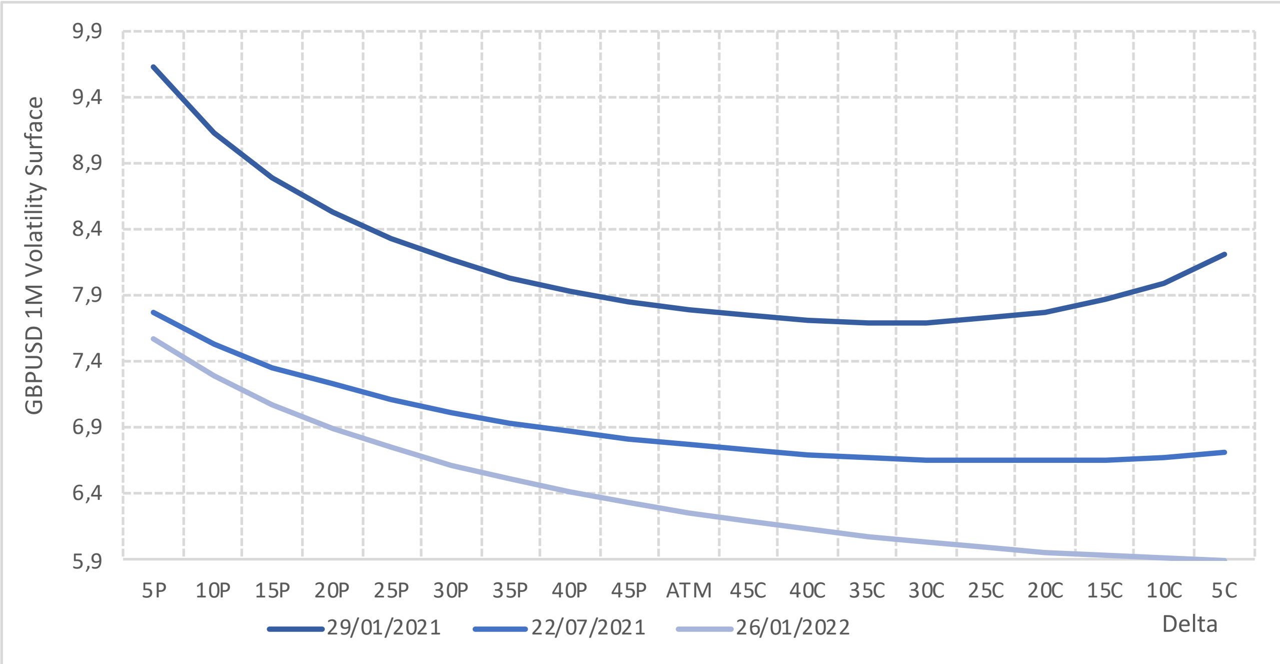

Figure 3: GBPUSD one-month implied volatility surface. *The three lines represent three different dates, to depict the surface change on a six-month basis. It can be seen how the smile shape flattens over the year and the overall volatility level decreases, depicting a shift in risk appetite.*

The shown volatility surfaces describe the more frequent change in the relative price of puts versus calls with regards to exchange rates, as opposed to equity options. In particular, the charts shows how the implied volatility has been almost symmetric from the start of 2021 and has then slowly moved towards a more negatively skewed shape in favor of JPY-OTM puts. Furthermore, for OTM currency options, the absolute increase in implied volatility can further express a shift in market risk sentiment, which is sometimes correlated with the degree of skewness. The second chart, depicting GBPUSD options volatility surface, can further simplify the concept. The flattening of the surface happens simultaneously with an overall decrease in implied volatility, which could suggest future lower volatility for the exchange rate. It can be sometimes cumbersome to find the right causality behind a change in skewness, especially when accompanied by a solid change in volatility. The aim of this paper is indeed to define a model able to capture these features by calibrating both the spot level and risk reversal quotes of the exchange rate, in order to identify some degree of correlation and at the same time optimize the performance of the model via root mean squared error and maximum likelihood. The calibration of a skewed process is indeed not the same as the one used for a process described more generally by fat tails. The degree of kurtosis, though essential for tail risk analysis and more, does not assume any difference between the two tails of the distribution. Even if volatility has been empirically found to be positively related with skewness, the sign of this relation may often differ in the currency market, especially when accounting for a shift in the likelihood of appreciation and thus strength of one currency against the other. An interesting approach related to the latter would be to consider a set of volatility regimes and observe the forward rate process, or mean of the distribution, across each different regime. The obtained span system would enable to discern between environments with a strong correlation where spikes in skewness cannot be discerned from spikes in volatility and with a higher degree of independency, where the level of asymmetry is not entirely affected by the deviation around the distribution mean. As will be thoroughly explained in the empirical analysis section, to obtain a stochastic skew process I will set the variance term equal to a constant, similar to setting the optimal volatility regime for a specific time period, and then, by constraining the stochastic time change and the jump process' diffusion terms to be related by a specific correlation term, I will obtain a stochastic skew process both in the long and short-run. It is indeed crucial to account for some degree of correlation between skewness and volatility, which in this environment enables to identify which currency has a higher likelihood of depreciation.

3.2 Jump Diffusion Models and Alternatives to Capture Stochastic Volatility and Skewness

Jump diffusion models have been introduced and widely used in computational finance after spotting some biases in more traditional models like the Black-Scholes (1973). These findings (Jarrow & Rosenfeld (1984), Jorion (1988)) have brought academics to shape the underlying stock price function as a mixture between a continuous diffusion part and a discontinuous jump part, the latter being able to account for extreme events. To name the first introductions of such models, Cox & Ross (1976) developed a pure jump model where the underlying is modeled by a jump process of deterministic size without a diffusion part, whilst Merton (1976) defined a model described by a so called "normal" vibration part, through a geometric Brownian motion, and an "abnormal" vibration part modeled by a jump process. Interestingly for option pricing matters, these models can generate implied volatility curves which are more similar to the smile curves observed in the market, as same as stochastic volatility models (Hull & White (1987)). More generally, the process can be defined by the following stochastic differential equation (SDE):

$$\frac{dS_t}{S_{t\_}} = \alpha dt + \sigma dW_t + dI_t$$

where the equation focuses on the nearest point in time preceding t and alpha is the drift term, sigma expresses the diffusion part, W is the standard Weiner process and I_t indicates the jump component. In particular the dynamics of I_t is described by the sum of J Poisson processes N and jump amplitudes Y, the latter to be greater than -1 to ensure non negative stock prices:

$$dI_t = \sum_{j=1}^{J} Y_{j,t} dN_{j,t}$$

where it is important to assume independency between the jump amplitudes and the related Poisson processes. As an example, in the Merton Model the jump amplitude is log-normally distributed. These models performed surprisingly well for derivatives pricing and risk management purposes, especially in the short-run more than in the long-run. Indeed, for longer time periods the probability of jump within a time series may become redundant, hence the relative underperformance (see, for example, Honore (1998)). Additionally, more recent studies (i.e. Lau et al. (2019)) have found out that by treating the jump event asymmetrically it can be further improved the overall accuracy of the model. The bottom idea is that jumps should be treated differently for upward and downward trends. As for the aim of this paper, the discontinuous jump part will be modelled in order to differentiate between positive and negative jumps in the exchange rate trend, to reflect the stochastic behavior of skewness in the forex market.

Another interesting approach, thoroughly detailed by La Bua and Marazzina (2021), to model volatility for pricing purposes is the so-called stochastic local volatility (hybrid) pricing paradigm. Local volatility models are well-known for derivatives pricing and risk management purposes, having been introduced by two main papers from Derman & Kani (1994) and Dupire (1994) through the introduction of a state-dependent diffusion process able to reproduce observed European options volatility surfaces. In particular, contributions from more recent papers such as Adreasen and Huge (2011), have managed to obtain a complete surface of arbitrage free option values, despite the existence of a discrete set of combinations of moneyness and maturity observed in the market. Albeit such results, local volatility models struggle with more exotic structures, in fact the deterministic nature of volatility does not perfectly match with the volatility dynamics of the underlying assets, thus delivering a much flatter surface than what observed in the market (Gatheral 2011). The latter is partially offset by adding a stochastic diffusion process into the volatility dynamics. Despite the more realistic parametrization, such alternative approach fails at matching implied volatility quotes for bottom levels of moneyness, for which a solutions has been attempted by the above mentioned jump diffusion model proposed by Bates (1996). The mixture of the two approaches has been in high demand among academics during the last decade, which have found the opportunity to separately calibrate those and subsequently merge them through the so-called leverage function. Bua et al (2021) introduce the Wishart Stochastic Local Covariance model in the attempt to provide a more comprehensive multidimensional variance dynamics within the stochastic local volatility framework. As for the proposed leverage function, it acts as a compensator when the stochastic volatility expected level is considerably different from the local volatility.

4 Empirical Analysis
4.1 The Model

The model used for this analysis aims at calibrating the log currency return and has been built on the back of the model proposed by Carr et al. (2007), which differentiates from traditional jump diffusion models, i.e. Bates (1996), because of the presence of two Levy processes described by left and right skewness respectively. The bottom idea is that the log currency return is described by a time-changed Levy process, within a probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, \mathcal{Q})$ defined through the risk-neutral probability measure Q:

$$s_t \equiv \ln \frac{S_t}{S_0} = \left(r_d - r_f\right)t + \left(L^R_{T^R_t} - T^R_t\right) + \left(L^L_{T^L_t} - T^L_t\right) \quad (1)$$

where the domestic and foreign interest rates are assumed to be continuously compounded and deterministic, $L^R \; and \; L^l$ are the two Levy processes, and $T^R_t \; and \; T^L_t$ denote the two stochastic time changes. The correlation coefficient is constrained to be different from zero only for the right Levy process and its related stochastic time change as well as for the left Levy process and the left stochastic time change.

The Levy components are made up of a pure jump component and a standard Brownian motion, according to the following equations:

$$L^R_t = J^R_t + \sigma W^R_t; \;\; L^L_t = J^L_t + \sigma W^L_t \quad (2)$$

where the two Brownian motions W as well as the pure jump components J are independent, whilst the latter can generate right and left skewness according to the following Levy densities:

$$f^R(x) = \begin{cases} \lambda e^{-x/v} x^{-\alpha-1}, & x > 0 \\ 0, & x < 0 \end{cases} ,$$

$$f^L(x) = \begin{cases} \lambda e^{-|x|/v} |x|^{-\alpha-1}, & x < 0 \\ 0, & x > 0 \end{cases} \quad (3)$$

Such densities, taken from the CGMY model of Carr et al (2002), enables the right skewed jump component to generate only up jumps and vice versa. For simplicity, λ and v are the same for both jumps. The former can be interpreted as an approximation of the aggregate activity level, whilst the latter controls the rate of exponential decay attached to the two densities. The alpha coefficient is instead affecting the sample path of the jump process and needs to remain less or equal than 2 to maintain finite quadratic variation. In particular the jump process can show finite activity ($\alpha<0$), infinite activity with finite variation ($0<\alpha<1$), or infinite variation ($1<\alpha<2$) (Carr (2007)).

4.2 Stochastic time change and Levy processes

By applying a stochastic time change to the Levy component, the model is able to generate stochastic volatility and thus it can generate implied volatility smiles, when it comes to converting the above returns into OTC option quotes. Furthermore, given two different time changes into the model, the relative weight of the Levy components can generate a skewed conditional return distribution and thus a different from zero risk reversal. All in all, such features brought the author to label it as a stochastic skew model. The concept of time change has firstly been applied in finance by Clark (1976), from the basis that security prices volatility tends to be time-varying and show some clustering along the time series. The simplest case is represented by a standard Brownian motion X and an independent continuous time change T. The computed time-changed process $Y_s = X_{T_s}$ has a conditional normal distribution with mean zero and variance equal to the time change factor T. The latter relies on the Dambis-Dubins-Schwarz theorem (1965), according to which:

"Every continuous local martingale $M = (M_s)_{s\geq 0}$ can be written as a time–changed Brownian motion $(B_{[M]s})_{s\geq 0}$, where $[M] = ([M]_s)$ $s\geq 0$ is the quadratic variation of M."

Thus, a quadratic variation $[M]_s = \int_0^s \sigma_r dW_r = \int_0^s \sigma_r^2 dr$ entails the independence between the standard Brownian motion W and the volatility parameter σ, which is equivalent to the independence between the standard Brownian motion and the continuous time change (Clark (1976)). Furthermore, it should be reminded the scaling property of a Brownian motion, which states that:

"Given a standard Brownian motion W_t and a positive constant c, the stochastic process $X_t = \frac{1}{\sqrt{c}} W_{ct}, t > 0$ is also a Brownian motion."

In this environment, the theorem proves that it is possible to switch from "spatial" scaling σX_t into "temporal" scaling $X_{\sigma^2 t}$. From a practical perspective, we can note that the independence between sigma and the Brownian motion excludes the leverage effect, according to which there is an observable negative correlation between asset returns and volatility.

For our financial modelling purpose, the time change can be defined as subordinator or absolutely continuous time change. Subordinators are stationary and not autocorrelated non-decreasing Levy processes. The absolutely continuous time changes are instead of the form $T_s = \int_0^s \tau_u \, du$ for a positive and integrable process theta. It is important to note that T is always continuous, whilst theta can have jumps. Indeed the latter is commonly known as the "instantaneous activity rate". The advantage of such structure is that it enables to obtain an affine model, which is then highly tractable. Consequently, I have used the second type of factor, to be specified below.

The activity rate follows a mean-reverting square root process of the following form:

$$d\tau_t^j = k(1 - \tau_t^j)dt + \sigma_v \sqrt{\tau_t^j} dZ_t^j \quad j = R, L \tag{4}$$

Where k expresses the speed of reversion, v is the so called 'vol of vol' parameter and Z are the Brownian motions related to each activity rate. It can be seen that, for normalization matters, I imposed the long run mean to be equal to 1 and the two remaining parameters, mean reversion k and vol of vol coefficient, to be the same for both left and right processes. From equation (4), the first two moments can be found through the decomposition described in the calibration section. The activity rate can be indeed seen as a random variable which follows a gaussian distribution with specific mean and variance (see the proof below). Furthermore I let the Brownian motions in the jump process and in the activity rate process to be correlated by p^L and p^R respectively. The latter is only assumed within left or right processes and not between the two. By setting p^L negative and p^R positive , it is possible to generate positive skewness in the short term via the jump component J, and in the long term via the positive correlation p^R, and viceversa on the other side.

4.3.1 Calibration

The above model has been calibrated to different exchange rates and risk reversal quote series, to gauge the accuracy of such structure to capture the stochastic behavior of skewness in the FX market. Firstly, to define the pure jump component J in equation (2) it is needed to define the initial jump mean, which will be constrained to be positive for the right jump and negative for the left jump. Unlike Carr (2007) I decided to make such means different between right and left jump so that the optimization procedure can weigh more on the magnitude of one of the two directional jumps if needed, clearly in relation with the real market observations. As for the stochastic time change factor, equation (4) refers to the so called Cox-Ingersoll-Ross model (CIR). The latter is widely used in mathematical finance to price interest rates derivatives, given that the peculiar diffusion term allows the interest rate, which in this case is instead a time change, not to turn negative. The exact solution, though not explicit, can be defined via the following proof.

Proof. Equation (4) can be rearranged as:

$$d\tau(t) + k\tau(t)dt = kdt + \sigma\sqrt{\tau(t)}dZ(t)$$

then by multiplying both sides of the equation by e^{kt}:

$$e^{kt}d\tau(t) + ke^{kt}\tau(t)dt = ke^{kt}dt + \sigma e^{kt}\sqrt{\tau(t)}dZ(t)$$

and by integrating both sides on the interval 0-t:

$$e^{kt}\tau(t) - \tau_0 = k\int_0^t e^{ks}ds + \sigma\int_0^t e^{ks}\sqrt{\tau(s)}dZ(s)$$

Finally, I obtain the solution of the following form:

$$\tau(t) = e^{-kt}\tau_0 + (1 - e^{-kt}) + \sigma e^{-kt}\int_0^t e^{ks}\sqrt{\tau(s)}dZ(s) \tag{5}$$

The latter can be then used to compute explicitly the mean and variance, to be used for the calibration procedure.

Taking expectation from both sides of (5):

$$\begin{aligned} E[\tau(t)] &= e^{-kt}\tau_0 + (1 - e^{-kt}) + \sigma e^{-kt}\int_0^t e^{ks}\sqrt{\tau(s)}E[dZ(s)] \\ &= e^{-kt}\tau_0 + (1 - e^{-kt}) \end{aligned} \tag{6}$$

where the last step is due to the property of the Brownian motion, according to which for all points in time s and t, t greater than s, the Brownian motion increments follow a gaussian distribution with mean zero and variance given by t – s. Additionally, from the obtained first moment of the process, it can be noted that the limit for t that tends to infinity of v(t) is equal to 1, which is the imposed long-run mean of the activity rate. Last but not least, the variance is computed as:

$$Var[\tau(t)] = E[\tau^2(t)] - (E[\tau(t)])^2$$

$$= E[(e^{-kt}\tau_0 + (1 - e^{-kt}) + \sigma e^{-kt}\int_0^t e^{ks}\sqrt{\tau(s)}dZ(s))^2]$$

$$-[e^{-kt}\tau_0 + (1 - e^{-kt})]^2$$

$$= 2\left(e^{-kt}v_0 + (1 - e^{-kt})\right)\sigma e^{-kt}E[\int_0^t e^{ks}\sqrt{v(s)}dZ(s)]$$

$$+\sigma^2 e^{-2kt}E[(\int_0^t e^{ks}\sqrt{\tau(s)}dZ(s))^2]$$

$$= 2\left(e^{-kt}\tau_0 + (1 - e^{-kt})\right)\sigma e^{-kt}[\int_0^t e^{ks}\sqrt{\tau(s)}E[dZ(s)]]$$

$$+\sigma^2 e^{-2kt}E[(\int_0^t e^{2ks}\tau(s)ds)]$$

$$= \sigma^2 e^{-2kt}\int_0^t e^{2ks}E[\tau(s)]ds$$

$$= \sigma^2 e^{-2kt}\int_0^t e^{2ks}(e^{-kt}\tau_0 + (1 - e^{-kt}))ds$$

$$= \frac{\sigma^2}{k}\tau_0(e^{-kt} - e^{-2kt}) + \frac{\sigma^2}{2k}(1 - 2e^{-kt} + e^{-2kt})$$

$$= \frac{\sigma^2}{k}\tau_0(e^{-kt} - e^{-2kt}) + \frac{\sigma^2}{2k}(1 - e^{-kt})^2 \tag{7}$$

where it has been applied both the zero mean property and the so called isometry property, which states that:

$$E\left[(\int_0^t X(s)dZ(s))^2\right] = \int_0^t E[X^2(s)]ds$$

Additionally, an important feature to generate skewness in the long-run is the correlation coefficient between the jump processes and their respective stochastic time changes. To generate two correlated Brownian motions, both with gaussian distribution and same variance, it can be used the well-known Kaiser-Dickman algorithm (Kaiser & Dickman, 1962). The latter is a simplified bivariate version of the Cholesky decomposition and has the following form:

$$Y = \rho X_1 + \sqrt{1-\rho^2} X_2$$

where X_1 and X_2 are two not correlated gaussian distributions with same variance and Y is a third gaussian distribution correlated by the coefficient ρ with X_1. However, for this case study I have used a slightly modified version of the above to account for two different diffusion terms. The equation can be proved by fixing the correlation coefficient between X_1 and a linear combination between X_1 and X_2 and imposing the variance preservation as follows:

$$corr(aX_1 + BX_2, X_1) = \rho \,;$$
$$var(aX_1 + BX_2) \equiv a^2 var(X_1) + B^2 var(X_2) = var(X_2)$$

which can be further developed into:

$$\rho = \frac{cov(aX_1 + BX_2, X_1)}{\sqrt{var(aX_1 + BX_2) var(X_1)}} = \frac{acov(X_1, X_1) + Bcov(X_1, X_2)}{\sqrt{var(aX_1 + BX_2) var(X_1)}}$$

$$= a\sqrt{\frac{var(X_1)}{a^2 var(X_1) + B^2 var(X_2)}}$$

$$a = \rho\sqrt{\frac{Var(X_2)}{Var(X_1)}} \,; B = \sqrt{Var(X_2)(1-\rho^2)}$$

Thus, thanks to the independence assumption between X_1 and X_2 and the variance preservation between X_2 and the linear combination Y, we obtain:

$$Y = \rho \frac{\sigma_{X_2}}{\sigma_{X_1}} X_1 + \sigma_{X_2}\sqrt{1-\rho^2} X_2$$

For the sake of completeness even if not used in this paper, which directly focuses on log currency returns and risk reversal quotes, to price European options from the computed log returns it is used the fast Fourier inversion method. Indeed, Carr and Wu (2004) explain how the problem of computing the generalized Fourier transform of a time-changed Levy process under the risk neutral measure Q can be simplified by finding the Laplace transform of the random time change under a complex valued measure. The Laplace transform depends on the characteristic exponent, which is provided by the Levy-Khintchine theorem. The latter specifies a drift term, a constant variance and the so called Levy density, which determines the arrival rate of jump. Accordingly, to change measure , it is applied the Girsanov's Theorem in order to obtain the activity rate process under the new complex measure. Having obtained in closed form the generalized Fourier transforms, Carr et al. (2007) compute the option values by applying the fast Fourier inversion on the transform.

4.3.2 Optimization procedures

Having set the log return formula, I converted each observation to the respective price level and compared it to the observed market quote. To optimize such comparison, I implemented the Generalized Reduced Gradient non-linear optimization method. It receives the objective function and releases the local optimum solution by setting the first order partial derivative equal to zero. One drawback of convergence algorithms is that they are highly dependent on initial conditions, thus the solution might not be a global optimum. However it is one of the fastest approaches and by setting the right initial framework, made up of the unknown parameters, it can provide better solutions than more robust methods like the evolutionary, which is based on natural selection theory and is more likely to find a global optimum. Using appropriate initial guesses for the unknown parameters, the GRG algorithm searches for values that minimize the sum of squared residuals relative to market bid prices, while setting the cross-sectional average error to zero. Throughout the loop the initial guesses are set according to the observed trend in market quotes, thus there could be some degree of autocorrelation in the final outcome.

Additionally, I measured the performance of the model through the maximum likelihood estimation. The latter concerns the maximization of a likelihood function by changing the unknown parameters of the model. Being in an environment with stochastic skewness I decided to base the procedure on the estimation of the conditional variance through the EWMA scheme. The latter belongs to the generalized autoregressive conditional heteroskedastic models (GARCH), being a conditional varying volatility model. Indeed, one-period log-returns are assumed to be Gaussian with zero mean and time-varying conditional variance:

$$r(t, t+\Delta)|F_t \sim N(0, \sigma^2(t, t+\Delta))$$

Moreover, the variance term is assumed to be a weighted average of past squared returns, given a smoothing parameter lambda comprised between 0 and 1. All said allows for: changing volatility by assigning larger weights to more recent observations; autocorrelation in squared returns and thus volatility clustering and excess kurtosis through the following form:

$$\sigma^2(t, t+\Delta) = (1-\lambda)\sum_{j=0}^{\infty} \lambda^j r^2(t-(j+1)\Delta, t-j\Delta)$$

Thanks to this structure, the model is also able to capture sudden market shocks, as opposed to the sample variance procedure which weighs equally each past observation. The latter equation can be decomposed to obtain the daily EWMA estimate. In order to do that, it is needed to subtract from both sides the past period variance term multiplied by the smoothing parameter:

$$\sigma^2(t, t+\Delta) - \lambda\sigma^2(t-\Delta, t) =$$
$$(1-\lambda)(\sum_{j=0}^{\infty}\lambda^i r^2(t-(j+1)\Delta, t-j\Delta) - \lambda\sum_{j=0}^{\infty}\lambda^j r^2(t-(j+2)\Delta, t-(j+1)\Delta)$$
$$= (1-\lambda)(\lambda^0 r^2(t-\Delta, t) + \sum_{j=1}^{\infty}\lambda^i r^2(t-(j+1)\Delta, t-j\Delta) - \lambda\sum_{j=1}^{\infty}\lambda^j r^2(t-(j+1)\Delta, t-j\Delta))$$
$$= (1-\lambda)r^2(t-\Delta, t)$$

It has been empirically proved that, for daily observations, the smoothing parameter lambda should be approximately equal to 0.94, whilst the initial variance term can be inserted into the set of unknown parameters. Finally, the sample likelihood of $\{r_0, \dots, r_{(T-1)\Delta}\}$ is given by the following productory:

$$\mathcal{L}\{\} = \prod_{j=0}^{T-1} f(r_{j\Delta}|F_j; \lambda, \sigma_0^2)$$

Where the term inside the productory is the conditional density of the log returns based on the information set at each point in time and the two unknown parameters. Finally the log likelihood has the following form and can be quickly maximized through an optimization algorithm:

$$log\mathcal{L}\{\} = \sum_{j=0}^{T-1} \ln\left(\frac{1}{\sqrt{2\pi\sigma_{j\Delta}^2}} e^{-\frac{1}{2}\left(\frac{r_{j\Delta}}{\sigma_j}\right)^2}\right)$$
$$= \sum_{j=0}^{T-1}\left(-\frac{1}{2}\ln(2\pi) - \frac{1}{2}\ln(\sigma_{j\Delta}^2) - \frac{1}{2}\left(\frac{r_{j\Delta}}{\sigma_{j\Delta}}\right)^2\right)$$

4.4 Results and Interpretations

The above calibration procedure has been applied to spot exchange rate levels and 1M risk reversal quotes to reflect the performance of such alternative jump diffusion model in reflecting price levels characterized by a non-deterministic skew process. To improve the reliability of results I focused the attention on 2021 daily observations, being this year characterized by a global recovery from the pandemic and thus the focus of investors towards the development of the yield curve, to go hand in hand with policy divergence between central banks. The latter has seen many exchange rates being strongly correlated with interest rates differential, from which the resurgence of carry trade strategies. In particular I have chosen to calibrate the model to USDJPY and GBPUSD exchange rates and risk reversal quotes, having the former been highly dependent on yield differential and box spread (difference between the steepness of each country's yield curve) and the latter been shaped by stagflation worries and risk aversion shifts.

Given the features of the chosen optimization algorithms, detailed in the calibration section, the starting values of the unknown set of variables can affect the overall optimization result. Because of the latter, many trials for each exchange rate and risk reversal have led to the conclusion that the most important parameters to be set are: the correlation coefficient between the Brownian motion attached to the jump component's diffusion term and the Brownian motion related to the square root process of the activity rate, or stochastic time change; and the alpha coefficient related to the jump process's frequency. In particular the two correlation coefficients enable the model to generate positive and negative skewness across the whole time series, and are found to be necessary despite the presence of the two directional jumps. The alpha coefficient instead determines the likelihood of a jump throughout the time series, and it has been constrained within the 1-2 range, in order to favor a high frequency jump specification, as found in Carr (2007).

As shown in the performance table the parameters can be quite different between crosses, but more interestingly between the exchange rate level and its related risk reversal quote. The former can be seen by the difference in the left jump mean and related correlation coefficient with the activity rate. Indeed for the chosen time period, but more generally for the Japanese yen against the US dollar, there has been a solid appreciation of the greenback, sustained by a positive market sentiment, an hawkish Fed and a rebound of the Japanese market, which has a negative relationship with the currency. For all the said, the model turns towards a more positive skew process, even if a lower number of negative jumps is still significant during risk-off periods, such as the resurgence of hospitalizations and covid-19 cases. On the other side, the British pound against the US dollar has shown a more mixed performance during the year, being the cross not as much related to fixed income movements as the previous exchange rate. Indeed the model does not seem able to capture the initial spikes in the level, despite the increase of the initial correlation coefficients or jump mean. However the overall performance of the model is significantly positive and, taking into consideration the fact that this exchange rate does not perfectly relate with the overnight rates spread, it is evident how the two directional jump components can play a role in explaining the more choppy movement of the pound against the dollar. Moving towards option quotes, the calibration of the model turns out to deliver a higher jump frequency factor, standard deviation and jump mean. The latter has been fully expected given that risk reversal quotes, and computed log returns, tend to have a more stochastic trend as opposed to the underlying spot levels. In particular, through the optimization algorithm, the alpha coefficient tends to the upper limit of 2, in order to maximize the frequency of each directional jump. The jump means are also greater than their spot level's counterparts to reflect the magnitude of each jump.

Overall, the two optimization methods deliver very similar set of parameters, which proves the reliability of one approach against the other and, more interestingly, the importance of a conditional variance model in order to maximize the likelihood of forex log returns. The average error term approximates to -0.49/0.02 and -0.08/0.10 volatility points for the USDJPY/GBPUSD levels and risk reversal quotes respectively. The latter are expressed in relative terms, thus it explains the difference between the two exchange rate levels' errors, being the first quoted around 110 and the second around 1.35.

Such findings demonstrate the reliability of the used jump diffusion model defined by two distinctive Levy processes in modelling short-term spot and option markets quotes when characterized by a non-deterministic skew process. Moreover, similarities between the two markets in terms of skewness can be found through the connection between significant risk reversal swings and spot level spikes, to be more interestingly found in the yen against the dollar price quotes. Indeed, the latter has been historically characterized by investors preference towards downside protection (negative risk reversal) in view of an appreciation of the yen, backed by a loss in risk appetite. However the subsequent reversal moves have been firstly characterized by an increase in the implied volatility of USD call options, sometimes reflected in a change of sign in the implied volatilities spread, followed by a less steep yet more durable dollar appreciation.

5 Conclusions

In this paper I calibrated a jump diffusion model described by two separated Levy processes to discern between up and down jumps in order to estimate log returns from currency spot levels and one-month risk reversal quotes. Literature on similar models and studies focused on skewness have pointed out the importance of skewness and kurtosis for both risk management and trading purposes. In particular in the currency option market, unlike the equity counterpart, quotes and related returns tend to have a more dynamic degree of skewness, to be reflected in investors preference for one currency against the other. Conversely, with regards to equity options, the so-called "crunch premium" makes the price of put option contracts almost systematically more expensive than call options' with the same level of moneyness, making the volatility smile more similar to a "smirk". Accordingly, whereas the level of kurtosis can be useful to control the magnitude of tail risk and so forth, skewness can be assumed negative for most equity options analysis. The main features of Carr's model are indeed tailored to the currency market and enable the process to generate stochastic skewness both in the short and long-run. Such development has been obtained with the presence of a stochastic time change factor, or activity rate, described by a square root process. The long-run skewness comes from a fixed correlation coefficient between the activity rate's diffusion term and the jump process' so that a positive coefficient for the up jump can generate positive skewness and vice versa for the opposite directional jump. Another important factor is the frequency parameter attached to the jump's density, which is constrained to express the highest possible frequency of jumps throughout the time series. However, the latter is found to be relatively higher for more volatile trends, to be observed in risk reversal quotes compared to spot levels, but also in GBPUSD against USDJPY spot levels. One added feature, to differentiate from the starting model, is the difference in absolute terms between up and down jump means. Being the case study focused on only one year of daily observations, I decided to optimize the two jump means independently between each other, so that the USDJPY price estimation process favored a more positive skew process, reflected by USD-OTM calls approaching the extremes versus JPY-OTM calls. Indeed, throughout the year, one-month and three-month risk reversals even changed sign on a couple of occasions, leading the greenback towards a strong appreciation against the yen, which has been one of the worst performers of the year within the G10 FX space. The model can be further improved through optimization procedures more tailored to the specific securities and could be used to model longer time series or even shorter, intra-day price trends. It is however of paramount importance the understanding of stochastic skewness within currency option prices, to be economically linked to currency risk premia through fundamental factors such as monetary policy, geo-politics and other macro differences between countries.

6 Appendix

Calibration results

The four charts depict the differences in terms of observed market quotes trends and the related performance of the model. The dotted line describes the performance of the model, whilst the straight line indicates the observed market quotes. Finally, the vertical lines indicate the position of the right and left jumps. Moreover, the two performance tables show the optimization parameters for the two optimization approaches, namely the GRG algorithm and the Maximum Likelihood estimation.

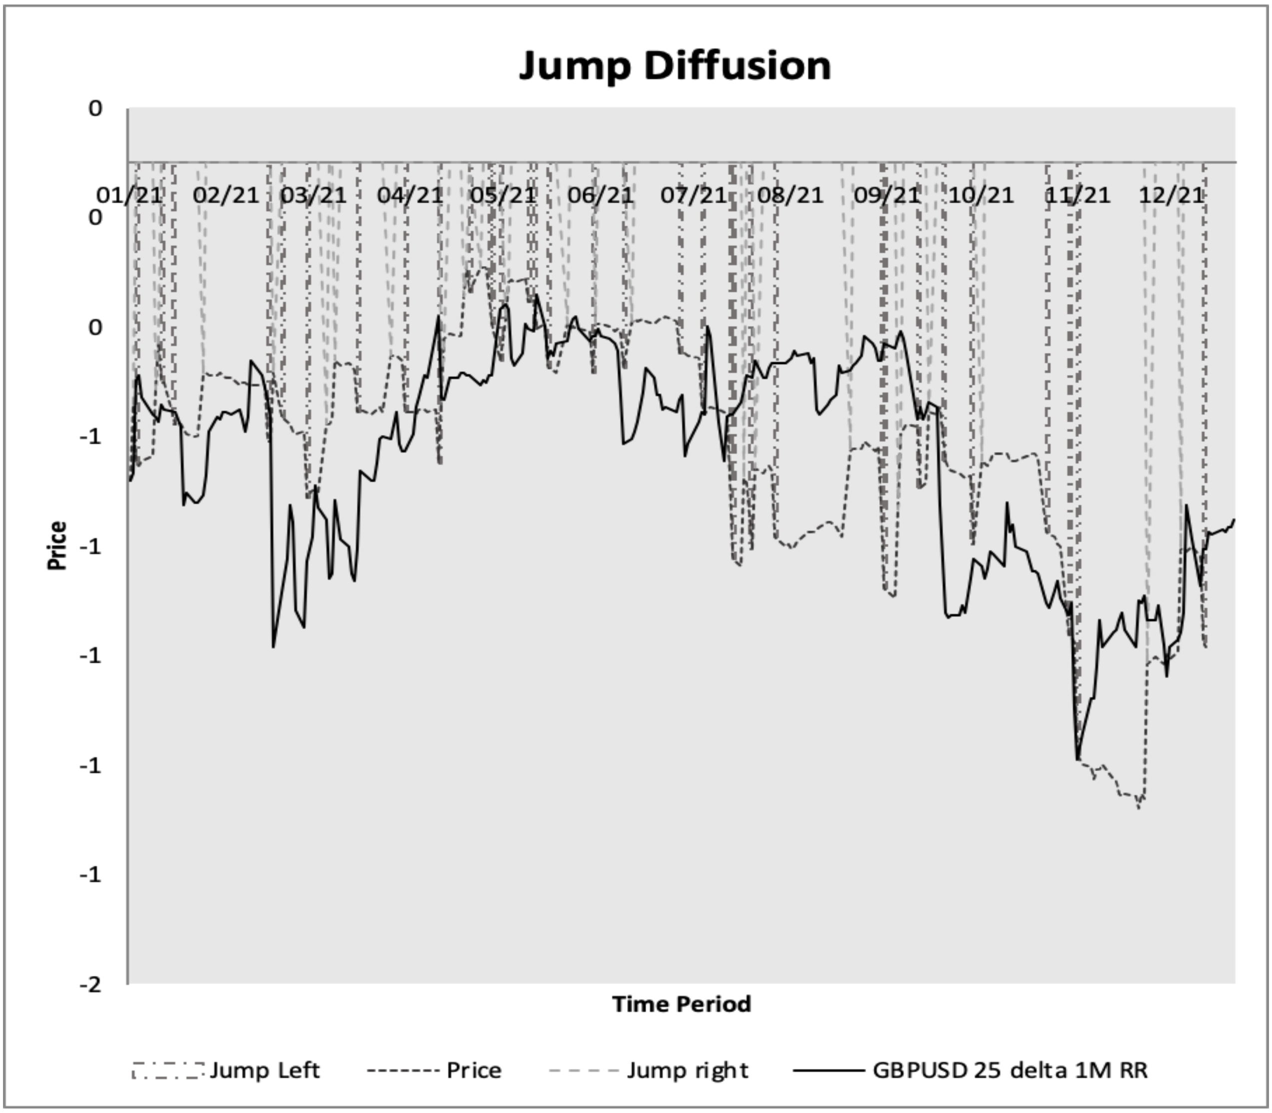

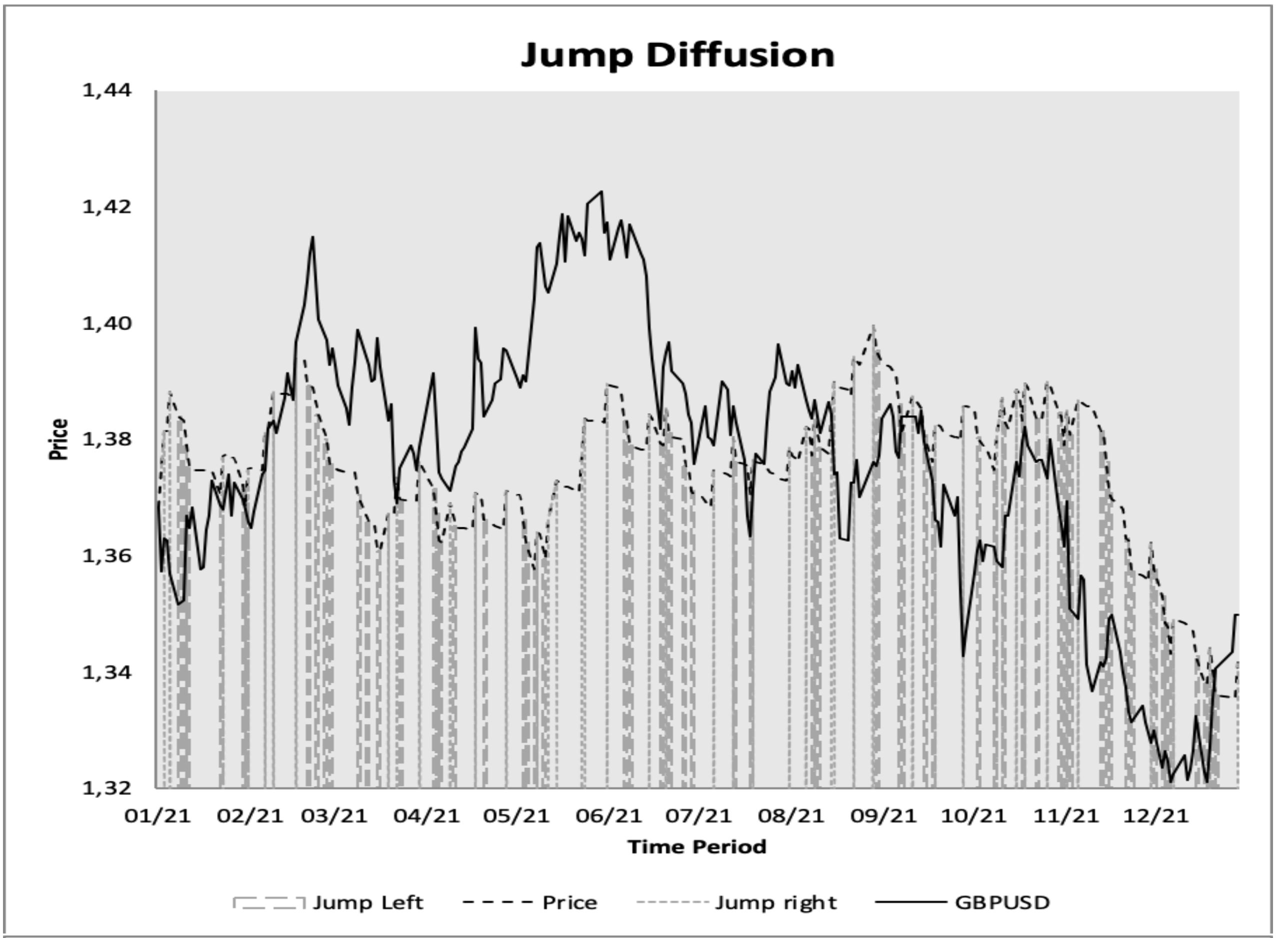

Jump Diffusion
1,44
1,42
1,40
1,38
1,36
1,34
1,32
Price
01/21
02/21
03/21
04/21
05/21
06/21
07/21
08/21
09/21
10/21
11/21
12/21
Time Period
Jump Left
Price
Jump right
GBPUSD

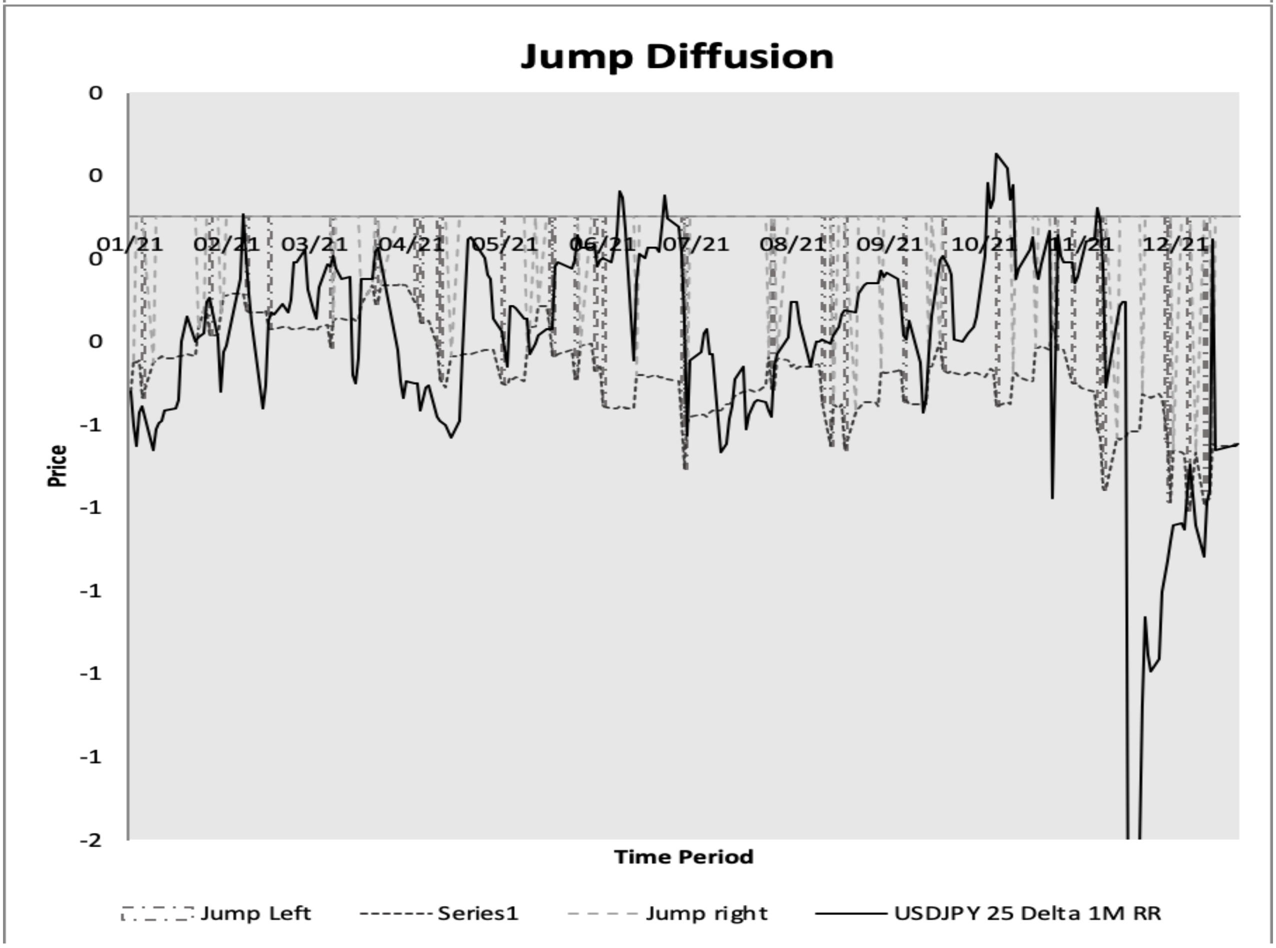

Jump Diffusion
0
0
0
0
-1
-1
-1
-1
-1
-2
Price
01/21
02/21
03/21
04/21
05/21
06/21
07/21
08/21
09/21
10/21
11/21
12/21
Time Period
Jump Left
Series1
Jump right
USDJPY 25 Delta 1M RR

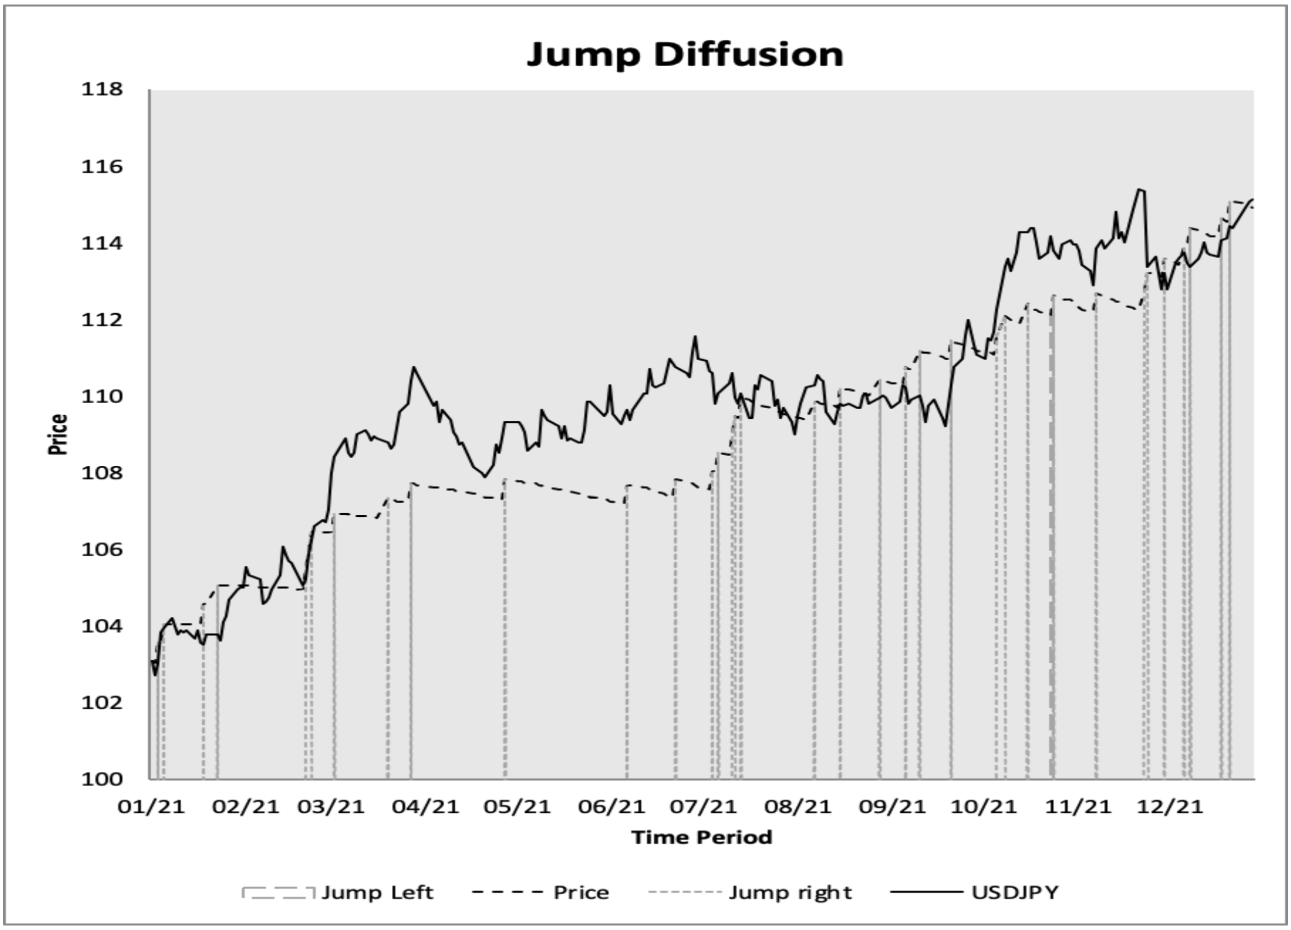

		USDJPY		GBPUSD	
	symbol	**spot**	**25 1M RR**	**spot**	**25 1M RR**
Optimization		GRG	GRG	GRG	GRG
Arrival rate	lambda	1.37	1.52	0.67	1.52
speed of reversion	k	0.45%	0.43%	0.48%	0.43%
Std Dev Of Time change	v_i	0.15%	1.83%	0.17%	1.73%
Std Dev of Levy density	v	9.8%	27.3%	9.0%	19.7%
alpha	α	1.15	1.82	1.57	1.60
Jump Mean Left	J_L	-0.11	-21.16	-0.30	-23.42
Jump Mean Right	J_R	0.49%	2.23%	0.47%	2.27%
Std Dev Of Jump Value	σ	0.93%	10.38%	0.96%	10.47%
correlation coefficient R	p_R	0.67	0.69	0.20	0.69
correlation coefficient L	p_L	-0.21	-0.32	-0.73	-0.40
average error	ε	-0.15	0.01	-0.22	0.01

	symbol	*USDJPY* spot	*USDJPY* 25 1M RR	*GBPUSD* spot	*GBPUSD* 25 1M RR
Optimization		ML	ML	ML	ML
Arrival rate	lambda	1.35	1.53	0.69	1.55
speed of reversion	k	0.48%	0.44%	0.46%	0.38%
Std Dev Of Time change	v_i	0.17%	1.85%	0.21%	1.66%
Std Dev of Levy density	v	8.7%	25.4%	8.9%	18.7%
alpha	a	1.21	1.93	1.59	1.71
Jump Mean Left	J_L	0.00	-0.02	0.00	-0.03
Jump Mean Right	J_R	0.45%	2.59%	0.45%	2.57%
Std Dev Of Jump Value	σ	0.89%	11.43%	0.99%	10.62%
correlation coefficient R	p_R	0.56	0.72	0.34	0.55
correlation coefficient L	p_L	-0.24	-0.38	-0.81	-0.28
initial volatility (EWMA)		0.015	0.0167	0.0123	0.0155

7 References

1. Albanese Claudio, Aleksandar Mijatovic (2009) “A stochastic Volatility Model for Risk-Reversals in Foreign Exchange” – International Journal of Theoretical and Applied Finance | Vol.12 pp. 887-899

2. Andreasen Jesper, Brian Norsk Huge (03/2010) “Volatility Interpolation” – available at SSRN: https://ssrn.com/abstract=1694972

3. Bates David S. (1996) “ Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options” – The Review of Financial Studies 0893-9454

4. Bertola Giuseppe, Ricardo Caballero (1992) “Target Zones and Realignments” – American Economic Review, vol. 82, 520-36

5. Brunnermeier Markus k., Stefan Nagel, Lasse H. Pedersen (11/2008) “Carry Trades and Currency Crashes” – National Bureau of Economic Research, Working Paper No. 14473

6. Brunnermeier Markus K., Lasse Heje Pedersen (11/2008) “Market Liquidity and Funding Liquidity” – The Review of Financial Studies, vol. 22, pp. 2201-2238

7. Campa Josè M., P.H. Keving Chang, Robert L. Reider (09/1997) “Implied Exchange Rate Distrbutions: Evidence from OTC Option Markets” – National Bureau of Economic Research Working Paper No. 6179

8. Carr Peter, Helyette Geman, Dilip B. Madan, Marc Yor (2002) “The Fine Structure of Asset Returns: an Empirical Investigation” – Journal of Business, vol.75, no. 2

9. Carr Peter, Liuren Wu (2004) “Time-changed Levy processes and options pricing” – Journal of Financial Economics 71 113-141

10. Carr Peter, Liuren Wu (2007) “Stochastic skew in currency options” – Journal of Financial Economics 86 213-247

11. Carr Peter, Liuren Wu (2007) “Theory and evidence on the dynamic interactions between sovereign credit default swaps and currency options” – Journal of Banking & Finance, vol. 31 2383-2403

12. Chang Young Bo, Peter Christoffersen, Kris Jacobs (01/2013) “Market skewness risk and the cross section of stock returns” – Journal of Financial Economics, vol. 107, pp. 46-68

13. Christensen B.J., N M Kiefer (1999) "Simulated moment of methods for empirical equivalent martingale measures" – Centre for Analytical Finance, University of Aarhus, 31

14. Daehwan Kim, Fabio Moneta (2021) "Long-term foreign exchange risk premia and inflation risk" – International Review of Financial Analysis 78 (2021) 101901

15. Dahlquist Magnus, Julien Penasse (05/2017) "The Missing Risk Premium in Exchange Rates" – Econometric Modeling: International Economics eJournal No. 157124987

16. Derman Emanuel, Iraj Kani (01/1994) "The Volatility Smile and Its Implied Tree" – Quantitative Strategies Research Notes, Goldman Sachs

17. Dubins Lester E., Schwarz Gideon (1965) "On Continuous Martingales" – National Academy of Sciences of the United States of America, Vol. 53, No. 5, pp. 913-916

18. Dupire, B. (1994) "Pricing with a Smile" – Risk, vol. 7 18-20

19. Gatheral Jim, Antoine Jacquier (2011) "Convergence of Heston to SVI" – Quantitative Finance, vol.11 1129-1132

20. Giot Pierre, Sebastien Laurent (05/2003) "Value-at-risk for long and short trading positions" – Journal of Applied Econometrics pp. 641-663

21. Heston Steven L.(1993) "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options" – The Review of Financial Studies, vol. 6, no. 2

22. Honore' P.(11/1998) "Pitfalls in Estimating Jump-Diffusion Models" – Centre for Analytical Finance . Aarhus School of Business Working Paper Series No. 18

23. Hull John, Alan White (06/1987) "The Pricing of Options on Assets with Stochastic Volatilities" – The Journal of Finance 281-300

24. Iseringhausen Martin (2021) "A time-varying skewness model for Growth-at-Risk" – European Stability Mechanism, Working Papers 49

25. Jackwerth Jens Carsten, Mark Rubinstein (12/1996) "Recovering Probability Distributions from Option Prices" – The Journal of Finance 1611-1631

26. Kaiser, H. F., & Dickman, K. (1992) "Sample and population score matrices and sample correlation matrices from an arbitrary population correlation matrix" – Psychometrika 27 (2), 179-182

27. La Bua Gaetano, Daniele Marazzina (09/2021) "A new class of multidimensional Wisart-based hybrid models" – Decisions in Economics and Finance: https://doi.org/10.1007/s10203-021-00357-4

28. Lau Joe Kein, Goh Kheng Yong, Lai Chow An (05/2019) “An empirical study on asymmetric jump diffusion and annuity pricing” – Bucharest University of Economic Studies

29. Park Sunjin (2017) “Global Macroeconomic Conditional Skewness and the Carry Risk Premium” – Essays in Financial Economics, Chapter 1, University of North Carolina Press

30. Sokol Andrej, Fernando Eguren-Martin (04/2020) “Attention to the tail(s): global financial conditions and exchange rate risks” – European Central Bank, Working Paper Series No 2387

31. Steshkova Alina (02/2021) “Conditional Skewness in Currency Markets” – Vienna University of Economics and Business Press

32. Veraart Almut E.D., Matthias Winkel (2010) “Time Change” – Encyclopedia of Quantitative Finance. Vol. 4, pp. 1812-1816

Summary

In this paper I calibrate a jump diffusion model to dynamically extract log returns from exchange rate levels and one-month risk reversal quotes. The model is made up of a deterministic part representing the spread between overnight interest rates of the two countries and a stochastic part divided in two Levy processes to discern between up and down jumps. The estimation period covers 2021's daily observations and is mainly described by economic recovery, covid-19 spikes and risk sentiment swings reflected in both currency spot and options short-term trends.

Currency returns have been historically described by a leptokurtic distribution (fat tails), together with some degree of asymmetry mostly connected with fundamental factors such as monetary policy or economic growth. Several studies have indeed incorporated skewness and other higher moments factors into their risk premia analysis, suggesting that currency, and not only, returns can be significantly explained by such distribution anomalies. For instance, Iseringhausen (2021) develops a model for time-varying skewness and finds out that idiosyncratic skewness related to currency returns varies significantly over time and is negatively related to the carry trade strategy, which entails going long high interest rates currencies and short lower-yielding currencies. Additionally, Brunnermeier et al. (2008) demonstrate that traders' funding constraints are a reasonable explanatory factor for the reduction of investment in high-yielding currencies, which explains why traders require a higher premium for holding such currencies, especially during periods of liquidity contraction. With regards to fundamental factors analysis, Alina Steshkova (2021) with a study based on interest rate differential and real exchange rate, i.e. the two more used risk premia factors for currency returns, finds out that interest rate differential is negatively bounded to skewness, which again explains the traders' funding constraint concept.

FX market drivers can be mainly divided into three categories, namely the stochastic movement of the exchange rate, captured by the most used Black & Scholes model, then the stochastic behavior of volatility, reflected for example in Heston's and Bates' models through specific diffusion processes, and the observed skewness of the log returns' risk neutral distribution. The latter is priced by the market via the so called risk reversal, which is given by the difference in implied volatility between an out-the-money call option and an out-the-money put option contract. Option prices have indeed been widely used to describe the distribution of market returns, even if these contracts are able to retrieve the risk-neutral probability density function (pdf) and not the actual density of such returns. Still, options are forward-looking instruments able to incorporate information that may only be relevant much later into the exchange rate's trend. Garch models, where volatility is assumed to

follow a stochastic process, do not update their parameters often enough to reflect major regime changes while on the other side, Campa, Chang and Reider (1997) have proven that currency options are highly sensitive to market news and thus to sudden changes in the exchange rate level, which might affect the overall trend structure. The main advantage of using an option-based approach to retrieve market returns is indeed that it does not depend on any specific functional form, which makes it a good fit to cope with different environments such as switching regimes, target zones and so forth. Empirically, it has been spotted a significant relation between the spot rate level and the degree of risk-neutral skewness, suggesting that exchange rate expectations (implied in options) are indeed "extrapolative", hence the connection between a strong currency and high probability of future appreciation.

For OTC currency options, there is a specific market aimed at providing an estimate of skewness related to the price of two option contracts with symmetric features. More specifically, the so called 'risk reversal' denotes the difference in price between a call and a put option with the same level of moneyness, where the latter expresses the likelihood of such contracts of expiring in-the-money. This approach, aimed at retrieving skewness, is fairly reliable given that it reflects market sentiment and it is not subject to estimation errors, given by smoothing procedures used to compute the probability density function. More interestingly, given that over-the-counter options contracts are priced in terms of Black and Scholes implied volatility, the difference in premium paid between call and put options simply reflects the spread in terms of volatility. It is important to remind that, by the Black and Scholes assumptions, the volatility surface across a series of strike prices should be flat, making risk reversal quotes approximately zero. Also, with stochastic volatility models like Heston's (1993), a random change in volatility is associated with an increase in the price of contracts far-from-the-money as opposed to contracts with a relatively higher intrinsic value, thus merely pointing towards a change in kurtosis.

These kind of models have unfortunately been unable to generate relevant time variations in the risk-neutral skewness attached to returns. The latter is a feature which can be practically recognized as deterministic in some kinds of option markets, like equity. In particular, equity options returns have been historically described by negatively skewed distributions, on the back of the so called "crunch premium" which reflects higher prices related to bottom-side hedges as opposed to top-side bets. On the contrary currency option prices, and thus implied volatilities, have historically changed in a more asymmetric way, reflecting both fundamentals differences between countries but also swings in risk appetite (e.g USD vs JPY). The calibration of a skewed process is indeed not the same as the one used for a process described more generally by fat tails. The degree of kurtosis, though essential for tail risk analysis and more, does not assume any difference between the two tails of the distribution. Even if volatility has been empirically found to be positively related with skewness, the sign of this relation may often differ in the currency market, especially when accounting for a shift in the likelihood of appreciation and thus strength of one currency against the other.

The model used for this analysis aims at calibrating the log currency return and has been built on the back of the model proposed by Carr et al. (2007), which differentiates from traditional jump diffusion models, i.e. Bates (1996), because of the presence of two Levy processes described by left and right skewness respectively. The bottom idea is that the log currency return is described by a time-changed Levy process, within a probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, \mathcal{Q})$ defined through the risk-neutral probability measure Q:

$$s_t \equiv \ln\frac{S_t}{S_0} = \left(r_d - r_f\right)t + \left(L^R_{T^R_t} - T^R_t\right) + \left(L^L_{T^L_t} - T^L_t\right) \quad (1)$$

where the domestic and foreign interest rates are assumed to be continuously compounded and deterministic, L^R *and* L^l are the two Levy processes, and T^R_t *and* T^L_t denote the two stochastic time changes. The correlation coefficient is constrained to be different from zero only for the right Levy process and its related stochastic time change as well as for the left Levy process and the left stochastic time change.

The Levy components are made up of a pure jump component and a standard Brownian motion, according to the following equations:

$$L^R_t = J^R_t + \sigma W^R_t; \; L^L_t = J^L_t + \sigma W^L_t \quad (2)$$

where the two Brownian motions W as well as the pure jump components J are independent, whilst the latter can generate right and left skewness according to the following Levy densities:

$$f^R(x) = \begin{cases} \lambda e^{-x/v} x^{-\alpha-1}, & x > 0 \\ 0, & x < 0 \end{cases} ,$$

$$f^L(x) = \begin{cases} \lambda e^{-|x|/v} |x|^{-\alpha-1}, & x < 0 \\ 0, & x > 0 \end{cases} \quad (3)$$

Such densities, taken from the CGMY model of Carr et al (2002), enables the right skewed jump component to generate only up jumps and vice versa. For simplicity, λ and v are the same for both jumps. The former can be interpreted as an approximation of the aggregate activity level, whilst the latter controls the rate of exponential decay attached to the two densities. The alpha coefficient is instead affecting the sample path of the jump process and needs to remain less or equal than 2 to maintain finite quadratic variation. In particular the jump process can show finite activity ($\alpha<0$), infinite activity with finite variation ($0<\alpha<1$), or infinite variation ($1<\alpha<2$) (Carr (2007)).

The stochastic time change is of the form $T_s = \int_0^s \tau_u \, du$ for a positive and integrable process theta. It is important to note that T is always continuous, whilst theta can have jumps. Indeed the latter is commonly known as the "instantaneous activity rate". The advantage of such structure is that it enables to obtain an affine model, which is then highly tractable.

The activity rate follows a mean-reverting square root process of the following form:

$$d\tau_t^j = k\left(1 - \tau_t^j\right)dt + \sigma_v\sqrt{\tau_t^j}dZ_t^j \quad j = R, L \tag{4}$$

Where k expresses the speed of reversion, v is the so called 'vol of vol' parameter and Z are the Brownian motions related to each activity rate. It can be seen that, for normalization matters, I imposed the long run mean to be equal to 1 and the two remaining parameters, mean reversion k and vol of vol coefficient, to be the same for both left and right processes. From equation (4), the first two moments can be found through the decomposition described in the calibration section. The activity rate can be indeed seen as a random variable which follows a gaussian distribution with specific mean and variance (see the proof below). Furthermore I let the Brownian motions in the jump process and in the activity rate process to be correlated by p^L *and* p^R respectively. The latter is only assumed within left or right processes and not between the two. By setting p^L negative *and* p^R positive , it is possible to generate positive skewness in the short term via the jump component J, and in the long term via the positive correlation p^R, and vice versa on the other side.

The above model has been calibrated to different exchange rates and risk reversal quote series, to gauge the accuracy of such structure to capture the stochastic behavior of skewness in the FX market. Having set the log return formula, I converted each observation to the respective price level and compared it to the observed market quote. To optimize such comparison, I implemented the

Generalized Reduced Gradient non-linear optimization method. It receives the objective function and releases the local optimum solution by setting the first order partial derivative equal to zero. One drawback of convergence algorithms is that they are highly dependent on initial conditions, thus the solution might not be a global optimum. However it is one of the fastest approaches and by setting the right initial framework, made up of the unknown parameters, it can provide better solutions than more robust methods like the evolutionary, which is based on natural selection theory and is more likely to find a global optimum. Additionally, I measured the performance of the model through the maximum likelihood estimation. The latter concerns the maximization of a likelihood function by changing the unknown parameters of the model. Being in an environment with stochastic skewness I decided to base the procedure on the estimation of the conditional variance through the EWMA scheme. The latter belongs to the generalized autoregressive conditional heteroskedastic models (GARCH), being a conditional varying volatility model. Indeed, one-period log-returns are assumed to be Gaussian with zero mean and time-varying conditional variance.

Given the features of the chosen optimization algorithms, the starting values of the unknown set of variables can affect the overall optimization result. Because of the latter, many trials for each exchange rate and risk reversal have led to the conclusion that the most important parameters to be set are: the correlation coefficient between the Brownian motion attached to the jump component's diffusion term and the Brownian motion related to the square root process of the activity rate, or stochastic time change; and the alpha coefficient related to the jump process's frequency. In particular the two correlation coefficients enable the model to generate positive and negative skewness across the whole time series, and are found to be necessary despite the presence of the two directional jumps. The alpha coefficient instead determines the likelihood of a jump throughout the time series, and it has been constrained within the 1-2 range, in order to favor a high frequency jump specification, as found in Carr (2007). Furthermore, being the case study focused on only one year of daily observations, I decided to optimize the two jump means independently between each other, so that the USDJPY price estimation process favored a more positive skew process, reflected by USD-OTM calls approaching the extremes versus JPY-OTM calls.

Overall, the two optimization methods deliver very similar set of parameters, which proves the reliability of one approach against the other and, more interestingly, the importance of a conditional variance model in order to maximize the likelihood of forex log returns. The average error term approximates to -0.49/0.02 and -0.08/0.10 volatility points for the USDJPY/GBPUSD levels and risk reversal quotes respectively. The latter are expressed in relative terms, thus it explains the difference between the two exchange rate levels' errors, being the first quoted around 110 and the second around 1.35.

Such findings demonstrate the reliability of the used jump diffusion model defined by two distinctive Levy processes in modelling short-term spot and option markets quotes when characterized by a non-deterministic skew process. Moreover, similarities between the two markets in terms of skewness can be found through the connection between significant risk reversal swings and spot level spikes, to be more interestingly found in the yen against the dollar price quotes. Indeed, the latter has been historically characterized by investors preference towards downside protection (negative risk reversal) in view of an appreciation of the yen, backed by a loss in risk appetite. However the subsequent reversal moves have been firstly characterized by an increase in the implied volatility of USD call options, sometimes reflected in a change of sign in the implied volatilities spread, followed by a less steep yet more durable dollar appreciation.

The model can be further improved through optimization procedures more tailored to the specific securities and could be used to model longer time series or even shorter, intra-day price trends. It is however of paramount importance the understanding of stochastic skewness within currency option prices, to be economically linked to currency risk premia through fundamental factors such as monetary policy, geo-politics and other macro differences between countries.