Multivariate Additive Subordination with Applications in Finance

Multivariate Additive Subordination with Applications in Finance

Giovanni Amici[a] Laura Ballotta[b] Patrizia Semeraro[c]

October 10, 2024

Abstract

We introduce a tractable multivariate pure jump process in which the trading time is described by an additive subordinator. The multivariate process retains the additivity property, and therefore is time inhomogeneous, i.e., its increments are independent but non stationary. We provide the theoretical framework of our process, perform a sensitivity analysis with respect to the time inhomogeneity parameters, and design a Monte Carlo scheme to simulate the trajectories of the process. We then employ the model in the context of option pricing in the FX market. We take advantage of the specific features of currency triangles to extract the joint dynamics of FX log-rates. Extensive tests based on observed market data show that our model outperforms well established pure jump benchmarks.

Keywords: Finance, Additive subordination, Multivariate stochastic processes, Option pricing, Currency triangles.

1 Introduction

The aim of this paper is to develop a multidimensional model for time inhomogeneous processes within the additive class (i.e., processes with independent and non-stationary increments), and to examine both its effectiveness in accurately fitting option implied volatility surfaces, and its ability to capture time variations in the dependence structure of financial assets. The proposed construction is obtained by additive subordination of Brownian motions (in the spirit of Mendoza-Arriaga and Linetsky, 2016 and Semeraro, 2022), in particular the multivariate process is obtained by component-wise Sato subordination.

The main contribution of this paper lies in the fact that our model allows for additivity both at marginal and joint level, and a flexible time structure for the correlation. Therefore the proposed construction circumvents simultaneously the issues of stationarity of the increments and constant correlation which affect multivariate models based on Lévy processes.

In traditional financial models based on Lévy processes, such as Variance Gamma (Madan et al., 1998), Normal inverse Gaussian (Barndorff-Nielsen, 1997) and CGMY (Carr et al., 2002), the log-return process displays independent and stationary increments which limits the model's ability to reproduce the variations of the implied volatility smile across maturities. Parsimonious generalizations

[a]Corresponding Author: giovanni.amici@polito.it, Department of Mathematical Sciences, Politecnico di Torino
[b]Faculty of Finance, Bayes Business School (formerly Cass), City St George's, University of London, UK
[c]Department of Mathematical Sciences, Politecnico di Torino

are represented by Sato processes (see Eberlein and Madan, 2009, Sato, 1999) which are additive processes constructed from time-1 self-decomposable laws. Compared to Lévy processes, Sato processes only carry one additional parameter which controls time inhomogeneity.

Financial univariate models based on Sato processes have been extensively studied in Carr et al. (2007), whose empirical results show that the best performing constructions stem from the time-1 law of the Variance Gamma (VG) and the Normal inverse Gaussian processes (NIG). The multivariate extensions of these distributions though are not self-decomposable. Marena et al. (2018a) and Guillaume (2018) bypassed this issue by building processes with Sato margins; however, these constructions are not jointly additive, and their correlation is constant over time.

Given the results of Carr et al. (2007), and the fact that both VG and NIG are subordinated Brownian motions, we focus on the construction of additive subordinators. Indeed, modelling log-returns by subordination of Brownian motions is appealing due to both their analytical tractability and their financial interpretation. This class of processes is tractable because they present known characteristic functions for all the relevant quantities, and therefore option pricing can be performed efficiently using Fourier transform techniques (see Eberlein et al., 2010, Fang and Oosterlee, 2009, for example). Furthermore, as subordinated Brownian motions can be interpreted as Brownian motions that evolve on a stochastic time scale, these constructions explicitly model business time, i.e., the time at which trades occur. Advocated by Clark (1973), models which explicitly capture the flow of trades are also empirically supported by Engle and Russell (1998) and Engle (2000).

Univariate nonnegative-valued Sato processes, i.e., Sato subordinators, have been developed by Sun et al. (2017), whilst multivariate Sato subordinators based on tempered stable distributions have been introduced by Semeraro (2022) to develop time inhomogeneous multivariate subordinated Brownian motions for the dynamics of the log-returns. However, the self-similarity of the multivariate subordinator allows for only one parameter devoted to model time inhomogeneity, which can be a significant restriction in a multi-asset context.

Our strategy to overcome this limitation is based on a factor construction, so that the multivariate process capturing the risk of n assets is given by the linear combination of a n-dimensional Brownian motion subordinated by independent univariate Sato subordinators, and a multivariate Brownian motion subordinated by a common Sato subordinator. The first component of the linear combination captures the idiosyncratic risk of each asset, the second component instead originates the systematic risk affecting all assets. The proposed approach uses tempered stable time-1 laws to build the Sato subordinators; as such, our model can be considered as a generalization of the constructions of Luciano and Semeraro (2010), Semeraro (2022) and also the one of Ballotta and Bonfiglioli (2016) for subordinated Brownian motions.

Our multivariate model presents the advantage of parsimony in terms of parameters compared to other multivariate constructions incorporating stochastic volatility features such as the ones proposed in De Col et al. (2013), Brigo et al. (2019), and Fontana et al. (2022) amongst others. Nevertheless, it is sufficiently rich to offer control on the variations in level, slope and curvature of the implied volatility smile (i.e., the so called volatility skew), as well as the variations of these smiles across maturity, and the variations over time of the correlation between assets. This is achieved also by means of two parameters of time inhomogeneity, also referred to as 'time scalers', for each asset: one of these is common across all assets, whilst the other characterizes the unique inhomogeneity nature of each asset. In particular, the role of these time scalers is studied by means of an extensive sensitivity analysis.

An additional contribution of this paper is the development of an algorithm for Monte Carlo simulation of these processes; the approach is based on sampling from a numerical cumulative distribution function recovered from the characteristic function by Fourier inversion in the tradition of Broadie and Kaya (2006), Ballotta and Kyriakou (2014) and Kang et al. (2017) amongst others.

The performance of our model in reproducing observed market implied volatilities is analysed in the context of the FX market. We show, in fact, that the multivariate model based on Sato subordinators proposed in this paper is closed under triangulation, and consequently ideal for this type of application (see the discussion in De Col et al., 2013, Ballotta et al., 2017, Fontana et al., 2022, for example). We also focus on the issue of recovering market consistent information on the correlation among assets. The fact that our FX model permits an arbitrage free pricing of the cross-rate facilitates the extraction of an 'implied' correlation matrix from the joint calibration of traded quotes of FX triangles.

Alternative approaches to recover plausible values for the correlation amongst assets rely either on historical estimates (see e.g., Luciano and Schoutens, 2006), or implied correlation extracted from option prices with a fixed expiry, which are set up by traders irrespective of liquidity (see Amici et al., 2023). Such solutions though cannot guarantee the absence of arbitrage, as shown by Ballotta et al. (2017) for the case of historical correlation.

The results from our empirical study based on observed market quotes confirm the ability of the model to reproduce the FX implied volatility surface of the margin currencies, and to consistently re-price options on the cross-rate. Furthermore, we analyze the validity of the above mentioned practice of using correlation values extracted from options with a given maturity, and show that arbitrage can arise when these values are used for the re-pricing of options on the cross-rate with different expiries. Therefore, our model can offer support at managerial level with a market consistent practice for the estimation of correlation values, which is particularly important also in other contexts, such as risk and portfolio management decisions.

The paper is organized as follows. In Section 2, we review the required concepts of Sato processes, Sato subordinators and tempered stable distributions. In Section 3, we develop the multivariate additive process, and study its properties. In Section 4, we design the multi-currency market model and describe how an additive process can be calibrated to such market. The performance of our model is analysed in Section 5, and Section 6 concludes. In the interest of space, all the proofs are deferred to the Appendix.

2 Preliminaries

In this section we review the notions concerning the distributions and processes used in the construction of the proposed multivariate model.

A probability measure η on $\mathbb{R}^d$ is called self-decomposable if, for any $b > 1$, there exists a probability measure ρ_b on $\mathbb{R}^d$ such that

$$\phi_\eta(\boldsymbol{u}) = \phi_\eta(b^{-1}\boldsymbol{u})\phi_{\rho_b}(\boldsymbol{u}), \quad \boldsymbol{u} \in \mathbb{R}^d,$$

where $\phi_\eta(\cdot)$ is the characteristic function of η. If η is self-decomposable, then η is infinitely divisible and $\phi_\eta(\boldsymbol{u}) = \exp\big(\psi_\eta(\boldsymbol{u})\big)$, with $\psi_\eta(\cdot)$ known as the characteristic exponent of η.

A stochastic process $\boldsymbol{S}(t), t \geqslant 0$, on $\mathbb{R}^d$ is a q-self-similar process if, for any $a > 0$, there exists $q > 0$ such that

$$\boldsymbol{S}(at) \overset{\mathrm{d}}{=} a^q \boldsymbol{S}(t), \quad \forall t \geqslant 0, \tag{2.1}$$

where $\stackrel{\mathrm{d}}{=}$ denotes equality in law. If η is a nontrivial self-decomposable distribution on $\mathbb{R}^d$, then, for any $q > 0$, there exists, uniquely in law, a nontrivial q-self-similar additive process $\boldsymbol{S}(t)$, $t \geqslant 0$, such that its time-1 law is η (Theorem 16.1 in Sato, 1999). We call this process Sato process. Using Eq. (2.1), the direct link between the time-1 and the time-t characteristic exponents of $\boldsymbol{S}(t)$ follows as

$$\psi_{\boldsymbol{S}}(\boldsymbol{u}; t) = \psi_{\boldsymbol{S}}(t^q \boldsymbol{u}; 1), \quad \boldsymbol{u} \in \mathbb{R}^d, \ q > 0.$$

If in addition $\boldsymbol{S}(t)$ is a nonnegative nondecreasing process, it is an additive subordinator, that we call Sato subordinator.

An important building block of our model is the well-known class of $\mathbb{R}$-valued (one-side) tempered stable (TS) distributions, which are self-decomposable, therefore infinitely divisible. In details, a one-side TS distribution, denoted by $\eta = TS(\alpha, \xi, \lambda)$, with parameters $\lambda, \xi > 0$, and $\alpha \in (0, 1)$ has characteristic exponent

$$\psi(u) = \lambda \Gamma(-\alpha) \left((\xi - \mathrm{i}u)^\alpha - \xi^\alpha \right), \quad u \in \mathbb{R}, \ \mathrm{i} = \sqrt{-1}$$

($\alpha = 0$ is also admitted, but it is a special case that we do not consider in this paper, see Pérez-Abreu and Stelzer, 2014).

It can be proved (see Li et al., 2016 and Semeraro, 2022) that the Sato process $S(t)$ associated to a one-side $TS(\alpha, \xi, \lambda)$ distribution is a Sato subordinator, named Sato-TS subordinator, with characteristic exponent

$$\psi(u; t) = \lambda \Gamma(-\alpha) \left((\xi - \mathrm{i}t^q u)^\alpha - \xi^\alpha \right), \quad u \in \mathbb{R}, \ t \geqslant 0.$$

We note that the Sato-TS subordinator is described by four parameters: α, ξ and λ which are inherited from the TS distribution, and the parameter q which drives the time inhomogeneity of the process.

3 A new multivariate additive process

In this section we introduce a new multivariate additive process that we name component-wise Sato subordinated Brownian motion (CSB), and study its properties with particular attention paid to the behaviour over time of the moments and the resulting correlation structure.

3.1 The CSB process

The CSB process is a multivariate additive process obtained by means of a factor construction. Thus, it is composed by an idiosyncratic component, formed by independent time changed Brownian motions, and a systematic component with a common stochastic time change. The factors are constructed in such a way that the resulting one dimensional marginal processes are time changed Brownian motions; furthermore, for specific choices of the factors' parameters, the marginal processes have self-decomposable time-1 distribution, e.g., Variance Gamma or Normal inverse Gaussian.

In details, let $\boldsymbol{B}(t)$, $t \geqslant 0$, be an $\mathbb{R}^n$-valued Brownian motion with independent components $B_j(t)$, each with drift $\mu_j \in \mathbb{R}$ and diffusion coefficient $\sigma_j > 0$. Let $\boldsymbol{B}^\rho(t)$, $t \geqslant 0$, be an $\mathbb{R}^n$-valued Brownian motion with drift $\boldsymbol{\mu}^\rho = (\mu_1 \kappa_1, \ldots, \mu_n \kappa_n)$ and diffusion matrix $\boldsymbol{\Sigma}^\rho = (\rho_{ij} \sigma_i \sigma_j \sqrt{\kappa_i} \sqrt{\kappa_j})_{ij}$, $i, j = 1, \ldots n$, with $\kappa_j > 0$, $j = 1, \ldots, n$, and assume $\boldsymbol{B}(t)$ independent of $\boldsymbol{B}^\rho(t)$.

Then, the CSB process $\boldsymbol{Y}(t)$ is defined as

$$\boldsymbol{Y}(t) = \begin{pmatrix} B_1(X_1(t)) + B_1^{\rho}(Z(t)) \\ \dots \\ B_n(X_n(t)) + B_n^{\rho}(Z(t)) \end{pmatrix}, \quad t \geqslant 0, \tag{3.1}$$

where $X_j(t)$ are independent Sato-TS subordinators with parameters $\alpha_j, \xi_j, \lambda_j, q_j$, $j = 1, \dots, n$, independent of $\boldsymbol{B}(t)$ and $\boldsymbol{B}^{\rho}(t)$, and $Z(t)$ is Sato-TS subordinator with parameters $\alpha_Z, \xi_Z, \lambda_Z, q_Z$, independent of $\boldsymbol{B}(t)$, $\boldsymbol{B}^{\rho}(t)$ and $X_j(t)$ for any $j = 1, \dots, n$. It is straightforward to verify that since $X_j(t)$, $j = 1, \dots, n$ and $Z(t)$ are additive, so is $\boldsymbol{Y}(t)$.

The construction (3.1) is such that each margin is a one dimensional time changed Brownian motion: adapting, in fact, the argument in Theorem 5.1 of Luciano and Semeraro (2010) to the time inhomogeneous case, we obtain

$$Y_j(t) = B_j(X_j(t)) + B_j^{\rho}(Z(t)) \stackrel{\mathrm{d}}{=} \tilde{B}_j\left(X_j(t) + \kappa_j Z(t)\right), \quad t \geqslant 0, \tag{3.2}$$

for a Brownian motion $\tilde{B}_j(t)$ with drift $\mu_j \in \mathbb{R}$ and diffusion coefficient $\sigma_j > 0$. A similar construction can be achieved using stable subordinators (see Luciano and Semeraro, 2010), leading to a nonsymmetric generalization of multivariate sub-Gaussian processes (see Samoradnitsky and Taqqu, 1994). However, self-similar stable subordinators are Lévy processes, while with tempered stable subordinators it is possible to allow for time inhomogeneity. Additionally, tempered stable subordinators have finite variance, which can be a desirable property in financial applications.

The factor-based structure of $\boldsymbol{Y}(t)$ allows us to identify an idiosyncratic risk driver $\boldsymbol{B}(\boldsymbol{X}(t))$ and a systematic risk driver $\boldsymbol{B}^{\rho}(Z(t))$ of the process dynamics – specifically, the systematic source is given by the common time change $Z(t)$. Furthermore, Eq. (3.2) clarifies the role played by the parameters κ_j, $j = 1, \dots, n$, as the weights of the common component of the change of time across the one dimensional marginal processes. As such, the parameters κ_j, $j = 1, \dots, n$, allow to modulate the correlation level between the margins, as shown explicitly in Section 3.2.

The process $\boldsymbol{Y}(t)$ admits explicit expression for the time-t characteristic exponent, which is provided in the following proposition.

Proposition 3.1. The time-t characteristic exponent of the process $\boldsymbol{Y}(t)$ defined in Eq. (3.1) is

$$\begin{aligned} \psi_{\boldsymbol{Y}}(\boldsymbol{u}; t) \quad = \quad & \sum_{j=1}^{n} \Gamma(-\alpha_j) \left(\left(\xi_j - t^{q_j} \left(\mathrm{i} u_j \mu_j - \frac{1}{2} u_j^2 \sigma_j^2 \right) \right)^{\alpha_j} - \xi_j^{\alpha_j} \right) \lambda_j \\ & + \Gamma(-\alpha_Z) \left(\left(\xi_Z - t^{q_Z} \left(\mathrm{i} \boldsymbol{u}^{\top} \boldsymbol{\mu}^{\rho} - \frac{1}{2} \boldsymbol{u}^{\top} \boldsymbol{\Sigma}^{\rho} \boldsymbol{u} \right) \right)^{\alpha_Z} - \xi_Z^{\alpha_Z} \right) \lambda_Z, \qquad \boldsymbol{u} \in \mathbb{R}^n. \end{aligned} \tag{3.3}$$

We note that Eq. (3.2) can be proved using the marginal characteristic exponent $\psi_{Y_j}(u_j; t) := \psi_{\boldsymbol{Y}}((0, \dots, u_j, \dots, 0); t)$.

Our CSB model encompasses the Sato-subordinated Brownian motion (SSBM) introduced in Semeraro (2022), and can be considered as a time inhomogeneous generalization of the multivariate Lévy processes introduced in Luciano and Semeraro (2010) (LS model), and the factor construction of Ballotta and Bonfiglioli (2016) applied to subordinated Brownian motions (BB model). Therefore, these models represents the natural benchmarks of our application, and are reviewed in Appendix A

for the reader's convenience.

In particular, if $\alpha_1 = \ldots = \alpha_n = \alpha_Z := \alpha$ and the same parameter set is used across processes, the time-1 laws of the CSB, the SSBM and the LS processes coincide. Furthermore, if also $q_1 = \ldots = q_n = q_Z := q$, our CSB process has the same time-t distribution as the SSBM process.

Although other choices are possible, in what follows we specify a model with a common parameter α, which from Section 3.3 onwards is chosen to ensure that the CSB process has NIG time-1 law. However, we allow the idiosyncratic components and the common component to have different q parameters driving time inhomogeneity in order to improve the fit to real data.

3.2 Moments and behaviour over time

In this section we study the one dimensional moments and the correlation of the time-t distribution of the CSB process and their asymptotic behaviour. This is relevant to gauge the ability of the process to model the implied volatility over both short and long maturities. In what follows, we conveniently set the notation $X(1) := X$, for any generic stochastic process $X(t)$, $t \geqslant 0$.

Let $\mathcal{C}_m(Y_j(t))$ denote the m-th cumulant of the j-th marginal process, then mean, variance, skewness and excess kurtosis of the j-th marginal process are respectively

$$\mathbb{E}[Y_j(t)] = \mathcal{C}_1(Y_j(t)), \quad \mathrm{Var}(Y_j(t)) = \mathcal{C}_2(Y_j(t)), \quad \mathrm{Skew}(Y_j(t)) = \frac{\mathcal{C}_3(Y_j(t))}{\mathcal{C}_2^{3/2}(Y_j(t))}, \quad \mathrm{Ekrt}(Y_j(t)) = \frac{\mathcal{C}_4(Y_j(t))}{\mathcal{C}_2^2(Y_j(t))}.$$

Proposition 3.2. Let $\boldsymbol{Y}(t)$, $t \geqslant 0$, be the process defined in Equation (3.1).

(i) The first four marginal cumulants are given by

$$\begin{aligned}
\mathcal{C}_1(Y_j(t)) &= \mu_j \left(t^{q_j}\mathcal{C}_1(X_j) + \kappa_j t^{q_Z}\mathcal{C}_1(Z)\right), && (3.4)\\
\mathcal{C}_2(Y_j(t)) &= \sigma_j^2 \left(t^{q_j}\mathcal{C}_1(X_j) + \kappa_j t^{q_Z}\mathcal{C}_1(Z)\right) + \mu_j^2 \left(t^{2q_j}\mathcal{C}_2(X_j) + \kappa_j^2 t^{2q_Z}\mathcal{C}_2(Z)\right) && (3.5)\\
\mathcal{C}_3(Y_j(t)) &= 3\mu_j\sigma_j^2 \left(t^{2q_j}\mathcal{C}_2(X_j) + \kappa_j^2 t^{2q_Z}\mathcal{C}_2(Z)\right) + \mu_j^3 \left(t^{3q_j}\mathcal{C}_3(X_j) + \kappa_j^3 t^{3q_Z}\mathcal{C}_3(Z)\right) && (3.6)\\
\mathcal{C}_4(Y_j(t)) &= 3\sigma_j^4 \left(t^{2q_j}\mathcal{C}_2(X_j) + \kappa_j^2 t^{2q_Z}\mathcal{C}_2(Z)\right) + 6\mu_j^2\sigma_j^2 \left(t^{3q_j}\mathcal{C}_3(X_j) + \kappa_j^3 t^{3q_Z}\mathcal{C}_3(Z)\right)\\
&\quad + \mu_j^4 \left(t^{4q_j}\mathcal{C}_4(X_j) + \kappa_j^4 t^{4q_Z}\mathcal{C}_4(Z)\right) && (3.7)
\end{aligned}$$

(ii) The pairwise correlation is given by

$$\begin{aligned}
&\mathrm{Corr}\left(Y_i(t), Y_j(t)\right)\\
&= \frac{\rho_{ij}\sigma_i\sqrt{\kappa_i}\sigma_j\sqrt{\kappa_j}\, t^{q_Z}\mathbb{E}(Z) + \mu_i\kappa_i\mu_j\kappa_j\, t^{2q_Z}\mathrm{Var}(Z)}{\prod\limits_{l\in\{i,j\}} \sqrt{\sigma_l^2 \left(t^{q_l}\mathbb{E}(X_l) + \kappa_l t^{q_Z}\mathbb{E}(Z)\right) + \mu_l^2 \left(t^{2q_l}\mathrm{Var}(X_l) + \kappa_l^2 t^{2q_Z}\mathrm{Var}(Z)\right)}}. \qquad (3.8)
\end{aligned}$$

From Proposition 3.2 we can deduce the asymptotic behaviour for short and long time horizons of skewness, excess kurtosis and pairwise correlation.

Proposition 3.3. Let $\boldsymbol{Y}(t)$, $t \geqslant 0$, be the process defined in Equation (3.1).

(i) The long-term marginal skewness and excess kurtosis are

$$\mathrm{Skew}(Y_j(t)) \underset{t\to\infty}{\longrightarrow} \begin{cases} \mathrm{sgn}(\mu_j)\,\mathrm{Skew}(X_j), & q_j > q_Z \\ \mathrm{sgn}(\mu_j)\,\mathrm{Skew}(Z), & q_j < q_Z \\ \mathrm{sgn}(\mu_j)\,\mathrm{Skew}(S_j), & q_j = q_Z \end{cases} \tag{3.9}$$

$$\mathrm{Ekrt}(Y_j(t)) \underset{t\to\infty}{\longrightarrow} \begin{cases} \mathrm{Ekrt}(X_j), & q_j > q_Z \\ \mathrm{Ekrt}(Z), & q_j < q_Z \\ \mathrm{Ekrt}(S_j), & q_j = q_Z, \end{cases} \tag{3.10}$$

where $S_j = X_j + \kappa_j Z$.

(ii) The short and long-term values of the correlation coefficient are

$$\mathrm{Corr}(Y_i(t), Y_j(t)) \underset{t\to 0}{\longrightarrow} \begin{cases} \rho_{ij}, & q_i > q_Z \text{ and } q_j > q_Z \\ 0, & q_i < q_Z \text{ or } q_j < q_Z \\ \mathrm{sgn}(\mu_i\mu_j)\,\frac{\rho_{ij}\sqrt{\kappa_i\kappa_j}\mathbb{E}(Z)}{\sqrt{\mathbb{E}(S_i)\mathbb{E}(S_j)}}, & q_i = q_j = q_Z \end{cases} \tag{3.11}$$

$$\mathrm{Corr}(Y_i(t), Y_j(t)) \underset{t\to\infty}{\longrightarrow} \begin{cases} 0, & q_i > q_Z \text{ or } q_j > q_Z \\ \mathrm{sgn}(\mu_i\mu_j), & q_i < q_Z \text{ and } q_j < q_Z \\ \mathrm{sgn}(\mu_i\mu_j)\,\frac{\kappa_i\kappa_j\mathrm{Var}(Z)}{\sqrt{\mathrm{Var}(S_i)\mathrm{Var}(S_j)}}, & q_i = q_j = q_Z \end{cases} \tag{3.12}$$

Eqs. (3.9)–(3.10) show that the marginal skewness and excess kurtosis of the process converge to nonzero values as the time horizon increases. As such, the CSB process breaks the central limit theorem and consequently it could be a suitable candidate to ensure that implied volatility smiles do not flatten out too quickly over the standard horizons given by traded maturities (see the discussion in Carr and Wu, 2003, for example).

Furthermore, differently from a number of popular models in the literature which typically embed a constant correlation coefficient (e.g., Lévy models), the CSB process is not affected by this restriction. Indeed, as it can be seen from Eq. (3.8), the pairwise correlation of the CSB process is time dependent, which is also verified by the different short and long-term values of Eqs. (3.11)–(3.12).

Finally, Eq. (3.11)–(3.12) show that the short and log-term values of the correlation differ significantly in sign and magnitude based on the values of the time scalers q_i, q_j and q_Z. This implies that the correlation term structure could display a variety of shapes, and therefore the process has the flexibility to describe a variety of market conditions. In Section 3.4 we further study the correlation term structure with a sensitivity analysis under an explicit specification of the subordinators.

3.3 Inverse Gaussian specification

For the purpose of a concrete study, we choose a specification of $\boldsymbol{Y}(t)$ in which the time-1 law of the one dimensional subordinators is the inverse Gaussian distribution. This corresponds to setting

$\alpha_1 = \ldots = \alpha_n = \alpha_Z = 0.5$ in the model parametrization.

In details, a random variable X_{IG} has inverse Gaussian distribution with parameters $C, D > 0$, i.e., $X \sim$IG(C, D), if it is a one-sided $TS(\alpha, \xi, \lambda)$ distribution with $\alpha = 0.5$, $C = -\lambda\Gamma(-\alpha) = 2\sqrt{\pi}\lambda$, and $D = \sqrt{\xi}$. The resulting characteristic exponent is

$$\psi_{\text{IG}}(u) = C\left(D - \sqrt{D^2 - 2\mathrm{i}u}\right), \qquad u \in \mathbb{R}.$$

In the context of the CSB model, let

$$X_j \sim \text{IG}\left(1 - a\sqrt{\kappa_j}, \frac{1}{\sqrt{\kappa_j}}\right), \qquad Z \sim \text{IG}(a, 1), \qquad 0 < a < \frac{1}{\sqrt{\kappa_j}}, \qquad j = 1, \ldots, n, \tag{3.13}$$

be the time-1 laws of the subordinators of Eq. (3.1). Then, by Proposition 3.1, the characteristic exponent of the joint process $\boldsymbol{Y}(t)$ is

$$\begin{aligned}\psi_{\boldsymbol{Y}}(\boldsymbol{u}; t) \quad = \quad & \sum_{j=1}^{n}\left(1 - a\sqrt{\kappa_j}\right)\left(\frac{1}{\sqrt{\kappa_j}} - \sqrt{\frac{1}{\kappa_j} - 2t^{q_j}\left(\mathrm{i}u_j\mu_j - \frac{1}{2}u_j^2\sigma_j^2\right)}\right) \\ & + a\left(1 - \sqrt{1 - 2t^{q_Z}\left(\mathrm{i}\boldsymbol{u}^\top\boldsymbol{\mu}^\rho - \frac{1}{2}\boldsymbol{u}^\top\boldsymbol{\Sigma}^\rho\boldsymbol{u}\right)}\right), \qquad \boldsymbol{u} \in \mathbb{R}^n.\end{aligned} \tag{3.14}$$

This is used for the Fourier-based pricing in Section 5.

From the above expressions, we can identify the set of model parameters as

$$\mathcal{P} = \{\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\kappa}, a, \boldsymbol{\rho}, \boldsymbol{q}\}, \tag{3.15}$$

with cardinality $4n + 2 + n(n-1)/2$ (we note that it is possible to recover the original Normal inverse Gaussian parametrization of Barndorff-Nielsen, 1997, through the relations $\mu_j = \beta_j\delta_j^2$, $\sigma_j = \delta_j$, and $\kappa_j = (\delta_j^2(\gamma_j^2 - \beta_j^2))^{-1}$, for $\delta_j, \gamma_j > 0$, $-\gamma_j < \beta_j < \gamma_j$, $j = 1, \ldots, n$).

3.4 Sensitivity analysis

In order to understand the role of the parameters of the CSB process, we perform a sensitivity analysis under the inverse Gaussian specification introduced in the previous section.

We note that the role of the parameters controlling the marginal distribution, i.e., μ_j, σ_j and κ_j is well known: although interconnected, these parameters are primarily linked to the sign and the level of the marginal skewness, the marginal variance and excess kurtosis, respectively. Furthermore, the role of the parameters a and ρ_{ij} in controlling the dependence structure has been studied by Marena et al. (2018b) for the case of the LS model that shares the same time-1 distribution as our CSB model. Thus, in the following we focus on the time-scaling parameters q_j and q_Z.

Unless otherwise stated, we consider two bidimensional base cases with parameters as specified in Table 1, in which we also report the resulting statistical features. In particular, we note that the set denoted as $\mathcal{P}_{\text{EQ}}$ is characterised by a pronounced negative skewness, which is a typical feature of equity log-returns, whilst the margins under the parametrization denoted as $\mathcal{P}_{\text{FX}}$ have skewness coefficients with mutually different signs, which could reflect for example the more ambiguous behaviour of FX log-rates. As such, we associate the parametrizations $\mathcal{P}_{\text{EQ}}$ and $\mathcal{P}_{\text{FX}}$ with realistic equity-like and FX-like markets, respectively.

	Margin parameters	μ	σ	κ	q	Mean	Variance	Skewness	Exc. Kurt.	Correlation
$\mathcal{P}_{\text{EQ}}$	Y_1	-0.16	0.13	0.45	1	-0.1073	0.0191	-1.5644	5.2755	0.6203
	Y_2	-0.14	0.11	0.53	1	-0.1019	0.0164	-1.7397	6.2196	
	Dependence parameters		a		ρ_{12}	q_Z				
			1		0.8	1				
	Margin parameters	μ	σ	κ	q					
$\mathcal{P}_{\text{FX}}$	Y_1	-0.03	0.13	0.45	1	-0.0201	0.0116	-0.3759	2.2009	0.5351
	Y_2	0.02	0.11	0.53	1	0.0146	0.0090	0.3359	2.3345	
	Dependence parameters		a		ρ_{12}	q_Z				
			1		0.8	1				

Table 1: Sensitivity Analysis: base parameter sets, and corresponding statistics of the time-1 law of the CSB process. Statistics obtained from Eqs. (3.4)–(3.8).

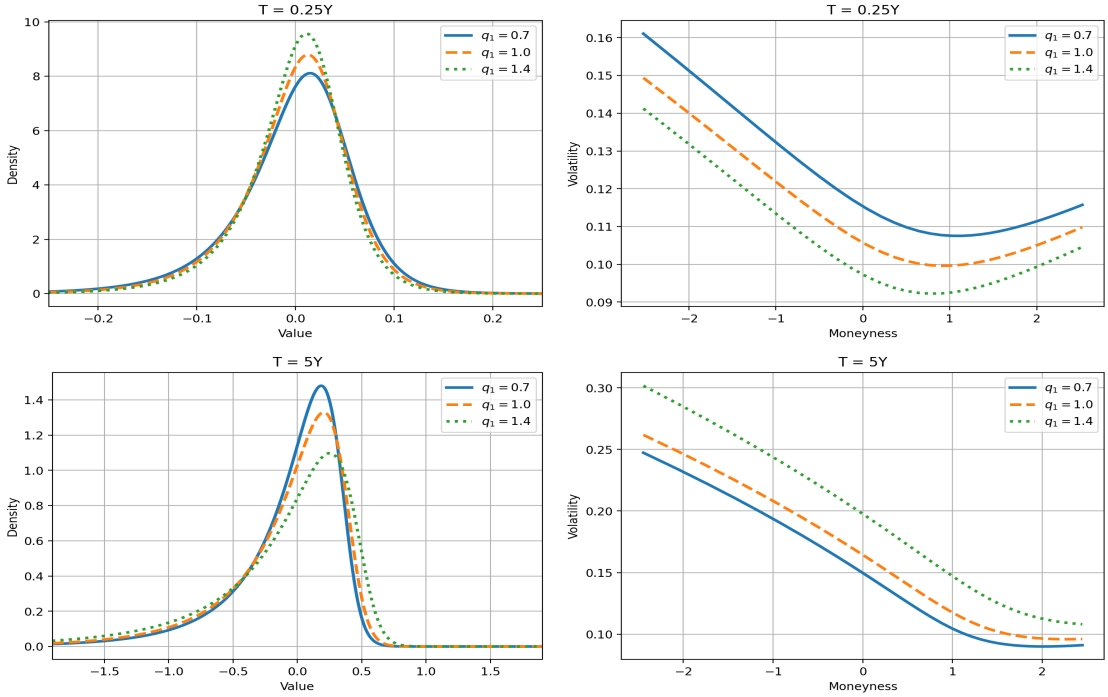

Figure 1: Effect of changes of parameter q_1 on the (centered) marginal densities (left-hand panels) and the implied volatility smiles (right-hand panels) of the CSB process, under short (top panels) and long (bottom panels) time horizons T (3 months and 5 years, respectively). Base-case model parameters: $\mathcal{P}_{\text{EQ}}$ (Table 1). Implied volatilities computed from call option prices – moneyness as in Eq. (3.16), risk-free interest rate and continuously-compounded dividend yield set to 0.

Marginal sensitivity with respect to q_j and q_Z The analysis is carried out by studying both the marginal density of the process and the implied volatility smiles obtained by assuming a log-return dynamics driven by the CSB process, across different levels of moneyness. The link between the log-return densities and the shape of the smiles has been investigated by Backus et al. (2004) among others, who highlight that the log-returns standard deviation controls the level of the implied volatility, whilst skewness and excess kurtosis are reflected in the slope and curvature of the smirk (i.e., the volatility skew).

In order to closely monitor the effect of the time scalers on the skewness we use the equity base case $\mathcal{P}_{\text{EQ}}$, in which the distribution asymmetry is pronounced. Furthermore, we consider two different time horizons, namely 3 months and 5 years, to better visualize the changes in densities and corresponding smiles over both short and long periods.

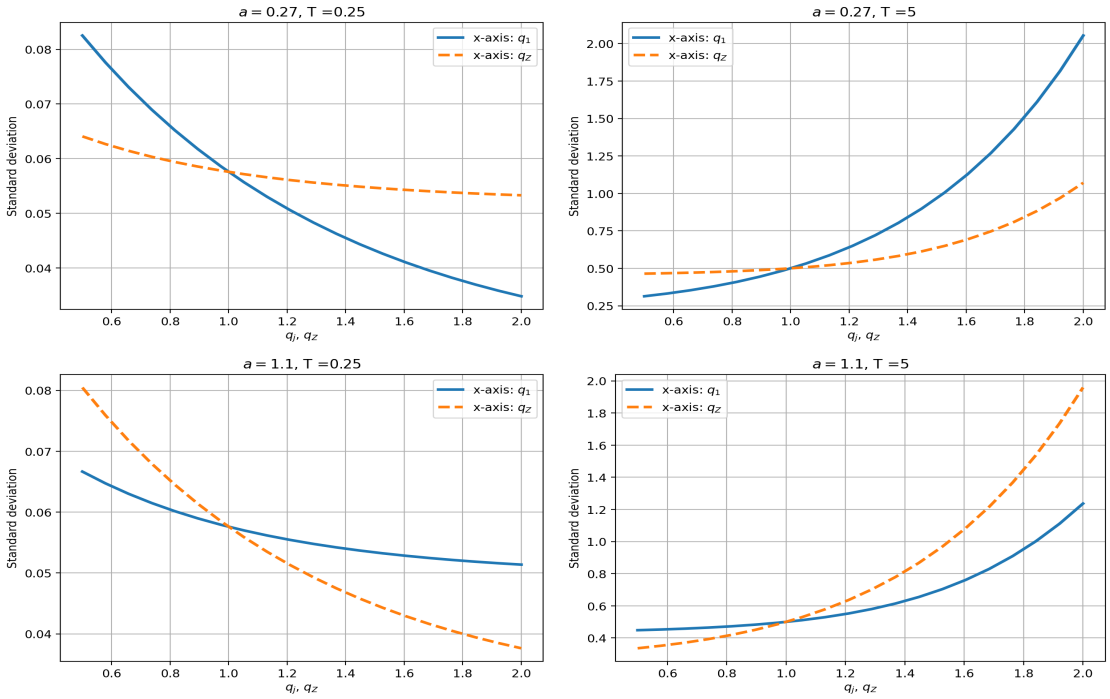

Figure 2: Effect of parameters q_j and q_Z on the CSB marginal standard deviation (see Eq. (3.5)), for different values of parameter a (0.27 and 1.1) and different time horizons T (3 months and 5 years). Base-case model parameters: $\mathcal{P}_{\text{EQ}}$ (Table 1).

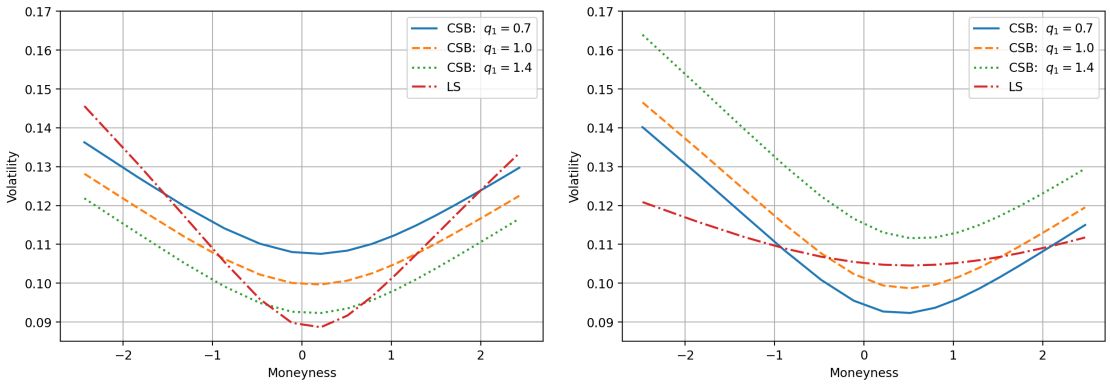

Figure 3: Short and long term implied volatility smiles of the CSB and the LS (Luciano and Semeraro, 2010) processes. Base-case model parameters: $\mathcal{P}_{\text{FX}}$ (Table 1) for both CSB and LS models, up to the time-scaling parameters. Time horizons: 3 months (left hand panel) and 5 years (right hand panel). Implied volatilities of the CSB model for different levels of q_1. Moneyness as in Eq. (3.16).

Moneyness is measured as

$$\frac{\ln(K/F_T)}{\overline{\sigma}\sqrt{T}}, \tag{3.16}$$

where K is the strike, T is the time to maturity, F_T is the T-forward price of the asset, and $\overline{\sigma}$ is an indicator of the average implied volatility of the asset of interest. As observed by Carr and Wu (2003), this choice of the moneyness allows for a fairer comparison between smiles associated to different expiries and to different financial instruments.

The results are illustrated in Figure 1 for the case of the first margin (results for the second margin are similar and available upon request). Specifically, we notice that changes in the time-scaling parameter q_1 impact primarily on the distribution's variance and translate into parallel shifts of the implied volatility. This is consistent with the findings in Section 3.2.

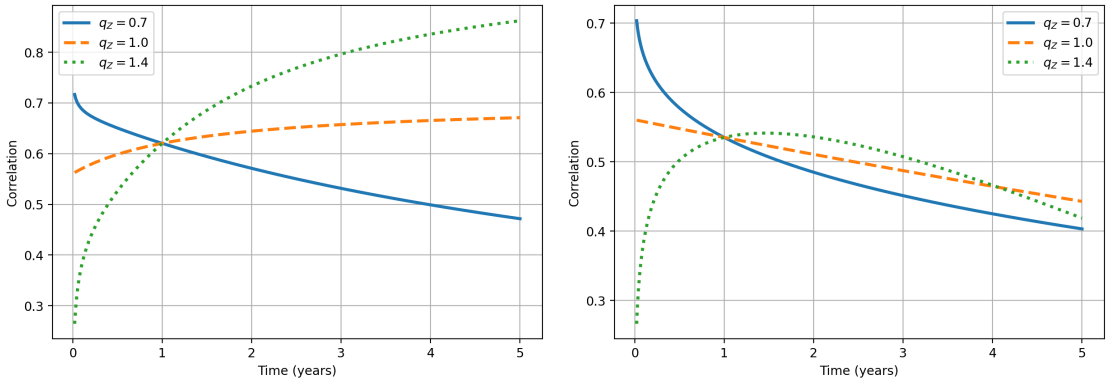

Figure 4: Correlation term structure of the CSB process (see Eq. (3.8)) as a function of time and q_Z. Base-case model parameters: $\mathcal{P}_{\text{EQ}}$ (left hand panel) and $\mathcal{P}_{\text{FX}}$ (right hand panel) as in Table 1.

In details, Eq. (3.5) highlights that the dominant term in the variance of the log-return distribution is of order t^{q_1} over short time horizons. This specific feature explains the inverse relation between q_1, and the implied volatility parallel shift over short maturities. However, for longer periods of time, the dominant term is of order t^{2q_1}, which implies the generally higher levels of the implied volatility, and the positive relation between the time-scaling parameter and the parallel shifts of the smirk.

From the expressions of the analytical moments we note that q_Z produces similar effects. However, the relative impact of q_1 and q_Z on the marginal dynamics depends on the moments of the idiosyncratic and the systematic subordinators. With the inverse Gaussian subordinators defined in Eq. (3.13), it is easy to check that the parameter a regulates this trade-off. This is shown in Figure 2, in which we analyse the behaviour of the process standard deviation for varying time-scaling parameters q_1 and q_Z: larger values of a correspond to a higher relative impact of q_Z with respect to q_1.

Behaviour over time As discussed in the previous sections, the proposed multidimensional CSB process is characterised by non-vanishing skewness and excess kurtosis over the long period. This break from the central limit theorem implies non-flattening implied volatilities for long maturity (see Carr and Wu, 2003, for example).

We illustrate this behaviour in Figure 3, in which we display volatility smiles generated by our model for different maturities and levels of the time scaler q_1. In order to gain a comprehensive view, we compare against the volatility smiles generated by the LS process of Luciano and Semeraro (2010) which we choose as representative model for multivariate Lévy processes due to the fact that they share the same time-1 distribution. To ensure a fair comparison, we assign the same parameter values to the two processes, which then possess equal time-1 laws, and use the – more symmetric – parametrization $\mathcal{P}_{\text{FX}}$ to better control for excess kurtosis, and therefore the curvature of the smiles.

The results visually confirm the rapid flattening of the smile in the LS model on the one hand, and the 'skew persistence' in the smile for the case of the CSB model on the other hand.

Correlation sensitivity To conclude the sensitivity analysis, we study the effect of the time scaling parameters on the joint dynamics of the CSB process. As Eqs. (3.8) and (3.11)–(3.12) show that the signs of μ_1 and μ_2 impact on the correlation values, we consider both parameter sets $\mathcal{P}_{\text{EQ}}$, in which $\text{sgn}(\mu_1) = \text{sgn}(\mu_2)$, and $\mathcal{P}_{\text{FX}}$, in which $\text{sgn}(\mu_1) \neq \text{sgn}(\mu_2)$. For sake of simplicity, we analyze only the cases in which the systematic time scaling parameter q_Z is either larger, or smaller or equal to both the idiosyncratic time scalers q_1 and q_2.

Results are shown in Figure 4, which confirms the richness of the shapes of the correlation term structure that our model can cater.

3.5 Monte Carlo simulation of the CSB process

As the characteristic function of the CSB process is available in closed form, pricing of plain-vanilla options and an increasing number of exotic options can be recovered via inverse Fourier transform methods in the same fashion as for Lévy and stochastic volatility models. However, for many applications including pricing of multidimensional structures and risk management, it might be desirable to implement Monte Carlo simulation methods. Whilst there is a vast literature regarding the simulation of Lévy processes and stochastic volatility models, sampling the trajectories of additive processes is a more recent area of study. The difficulty with additive processes stems from the fact that the increments follow mutually different laws over different time horizons due to time inhomogeneity. Furthermore, the underlying distribution could be non trivial to sample from.

In light of the above, in this section we give a concise overview of how the CSB process can be efficiently simulated exploiting the knowledge of its characteristic function. Pioneered by Broadie and Kaya (2006) for the Heston model, this methodology of combining Monte Carlo with sampling from the characteristic function of a process via Fourier inversion has been adopted among others by Ballotta and Kyriakou (2014) for the CGMY process, Kang et al. (2017) for the Wishart multidimensional model, and more recently by Azzone and Baviera (2023) for one dimensional additive processes.

Our simulation strategy exploits the mutual independence of the increments of both the Sato subordinators and the CSB process, and also the factor structure of the CSB process. Therefore, the scheme is simple and modular as it is applicable under any assumption concerning the time-1 law of the Sato subordinators. Thus, the algorithm is organized as follows.

The first step is articulated in two parts. Given a time interval $[s,t]$, we recover numerically the distribution of the increments of the one dimensional Sato subordinators composing the CSB process by inversion of the relevant characteristic function. This is achieved by means of the COS method (see Fang and Oosterlee, 2009). Then, this numerical distribution is used to sample the increments of the univariate subordinators. The second step consists in the simulation of the relevant Brownian motions over the (random) time scale given by the increments of the Sato subordinators obtained in the previous step. Finally, the increments of the CSB process are generated using Eq. (3.1). We refer to this scheme as the COS-Monte Carlo, henceforth COS-MC.

In more details, let α and β denote the truncation range of the numerical characteristic function, and N be the number of expansion terms of the Fourier cosine series. Then, the COS approximation of the probability density function (PDF) of the increments of a generic Sato subordinator $X^{[s,t]} := X(t) - X(s)$, $0 \leqslant s < t < \infty$, is

$$f_X(x) \approx \hat{f}_x = \sum_{h=0}^{N-1}{}' C_h \cos(u_h(x-\alpha)) \tag{3.17}$$

with

$$C_h = \frac{2}{\beta-\alpha}\mathrm{Re}\left(\phi_X(u_h)e^{-\mathrm{i}u_h\alpha}\right), \quad u_h = \frac{h\pi}{\beta-\alpha}, \quad h = 0,\dots,N-1$$

for some evaluation point x. The symbol Σ' indicates that the first summation term is halved.

The numerical CDF is recovered from Eq. (3.17) for a fixed grid $\boldsymbol{x} := (x_0, x_1, \ldots, x_{M-1})$ as

$$F_X(x_m) \approx \hat{F}_{x_m} = \sum_{h=0}^{N-1}{}' \, C_h \, \zeta_h(x_m)$$

with

$$\zeta_h(x_m) = \begin{cases} \left(\sin\left(h\pi \frac{x_m - \alpha}{\beta - \alpha} \right) - \sin\left(h\pi \frac{x_0 - \alpha}{\beta - \alpha} \right) \right) \frac{\beta - \alpha}{h\pi}, & h \neq 0 \\ x_m - x_0 & h = 0. \end{cases}$$

The selection of suitable values for α and β requires some attention, as a narrow interval $[\alpha, \beta]$ leads to a large truncation error, while a too large interval can produce a large discretization error, if N is not adapted accordingly. Based on extensive testing on a range of Lévy processes, Oosterlee and Grzelak (2019) suggest to select the integration range based either on the cumulants of the target quantity $X^{[s,t]}$ or, in a simplified version, only on the time length $[s, t]$.

In the following we adopt the latter rule and, due to the fact that the Sato subordinators are nonnegative processes, we set the integration range as

$$[\alpha, \beta] = \left[0, \; G\sqrt{t - s} \right], \tag{3.18}$$

where G regulates the trade-off between the truncation and the discretization error. Therefore, the time dependence of the above range is consistent with the proportional increase of the variation of the increment distribution as the time length increases.

Once $\hat{\boldsymbol{F}} = (\hat{F}_{x_0}, \ldots, \hat{F}_{x_{M-1}})$ has been approximated on the grid $\boldsymbol{x} := (x_0, x_1, \ldots, x_{M-1})$, we can sample from this law via interpolation: although other choices are possible, we use linear interpolation as in Glasserman and Liu (2010). Therefore, assume that $\boldsymbol{x}$ is equally spaced with intervals of length $\bar{x}$. Then, given a sample deviate $\breve{U} \sim \text{Uniform}(0, 1)$, the corresponding sample deviate from the CDF $F_X(x)$ is obtained as

$$\breve{X} = \begin{cases} \frac{x_m \hat{F}_{x_{m+1}} - x_{m+1} \hat{F}_{x_m} + \bar{x}\breve{U}}{\hat{F}_{x_{m+1}} - \hat{F}_{x_m}}, & \text{if } \; \hat{F}_{x_m} \leqslant \breve{U} < \hat{F}_{x_{m+1}} \\ x_0, & \text{if } \; \breve{U} < \hat{F}_{x_0} \\ x_{M-1}, & \text{if } \; \breve{U} \geqslant \hat{F}_{x_{M-1}}. \end{cases}$$

The accuracy of this method clearly relies on a good approximation of the CDF, which in turns depends on the choice of the grid $\boldsymbol{x}$. In our tests, we find that a good selection of the range of $\boldsymbol{x}$ should follow similar criteria as for the choice of the truncation range for the numerical Fourier transform. Specifically, we set

$$[x_0, x_{M-1}] = \left[0, \; D\sqrt{t - s} \right],$$

where D has the same order of magnitude of G in Eq. (3.18). We formalize our Monte Carlo scheme for the CSB process in Algorithm 1.

We now examine the validity of such a scheme by simulating the CSB process (using the $\mathcal{P}_{\text{EQ}}$ parametrization, see Table 1) over a 1-year horizon. Following extensive numerical experiments, we set $N = 2^{12}$, $G = 20$, $M = 2^{11}$, $D = 20$, where N and M are sufficiently large so that no significant improvement in the accuracy of the CDF approximation can be further achieved. To assess the accuracy of the scheme, we compare the first four empirical standardized marginal moments and the

Algorithm 1 Monte Carlo scheme of the CSB process $\boldsymbol{Y}(t)$ (see Eq. (3.1)). CSB parameters: see Table 1. Parameters of the COS method for density recovery: N (number of expansion terms); M (number of evaluation points of the numerical CDF); G and D (arbitrary coefficients to regulate the truncation ranges of the characteristic function and the CDF, respectively). Time grid of the simulation: $(t_0, t_1, \ldots, t_K)$. Number of Monte Carlo samples: L.

Inputs: $\{t_k\}_{k=0}^K$. COS parameters: N, G, M, D. CSB parameters. #MC samples: L.

Initialize an $L \times n$ matrix $\boldsymbol{Y}$ with entries equal to zero
for $k = 0, \ldots, K-1$ **do**
  Set $\alpha = 0$ and $\beta \leftarrow G\sqrt{t_{k+1} - t_k}$
  Set $\boldsymbol{x} \leftarrow (x_0, \ldots, x_{M-1})$, where $x_0 = 0$, $x_{M-1} = D\sqrt{t_{k+1} - t_k}$, $x_{m+1} - x_m = \bar{x}, \forall m$.
  for $X = X_1, \ldots, X_n, Z$ **do**
    Get the characteristic function of the law of the increment $[t_k, t_{k+1}]$:
    $\phi_X := \phi_X(t_k, t_{k+1}) \leftarrow \phi_X(t_{k+1})/\phi_X(t_k)$
    Get numerical CDF $\hat{\boldsymbol{F}}^{(X)}$ on grid $\boldsymbol{x}$:

$$\hat{\boldsymbol{F}}^{(X)} \leftarrow \left\{ \frac{x_m}{\beta} + \sum_{h=1}^{N-1} \frac{2}{h\pi} \mathrm{Re}\left(\phi_X\left(\frac{h\pi}{\beta} \right) \right) \sin\left(h\pi \frac{x_m}{\beta} \right) \right\}_{m=0}^{M-1}$$

    Generate i.i.d. random variables $\{\breve{U}_l\}_{l=1}^L \sim$ Uniform(0,1)
    Compute the $\hat{\boldsymbol{F}}^{(X)}$-distributed sample $\{\breve{X}_l\}_{l=1}^L$ as

$$\breve{X}_l \leftarrow \frac{x_m \hat{F}_{x_{m+1}} - x_{m+1}\hat{F}_{x_m} + \bar{x}\breve{U}_l}{\hat{F}_{x_{m+1}} - \hat{F}_{x_m}}, \text{ if } \hat{F}_{x_m} \leqslant \breve{U}_l < \hat{F}_{x_{m+1}}$$

$$\breve{X}_l \leftarrow x_0, \text{ if } \breve{U}_l < \hat{F}_{x_0}$$

$$\breve{X}_l \leftarrow x_{M-1}, \text{ if } \breve{U}_l \geqslant \hat{F}_{x_{M-1}}$$

  end for
  Generate $\{\breve{B}_{l,j}\}_{L\times n} \sim \mathcal{N}(0,1)$
  Generate $\{\breve{B}^\rho_l\}_L \sim \mathcal{N}(\boldsymbol{0}, \tilde{\Sigma})$, where $\tilde{\Sigma} = \{\rho_{ij}\}_{i,j=1,\ldots,n}$
  Compute $\breve{\boldsymbol{Y}}^{[t_k, t_{k+1}]} \leftarrow \left\{ \mu_j \breve{X}_{l,j} + \sigma_j \sqrt{\breve{X}_{l,j}} \breve{B}_{l,j} + \mu_j \kappa_j \breve{Z}_l + \sigma_j \sqrt{\kappa_j \breve{Z}_l} \breve{B}^\rho_{l,j} \right\}_{L\times n}$
  Update $\boldsymbol{Y} \leftarrow \boldsymbol{Y} + \breve{\boldsymbol{Y}}^{[t_k, t_{k+1}]}$
end for
return $\boldsymbol{Y}$

empirical correlation with the corresponding values of the analytical distribution obtained in Section 3.2. We repeat the test with one, two, and four simulation steps, in order to show that the results are robust to the specific increments being simulated, provided that the terminal (i.e., time-1) distribution is the same. Results are presented in Table 2.

Due to a slight loss in accuracy as the number of time steps increases, we plot the sampled time-1 marginal PDF of the 4-step Monte Carlo against its direct estimate given by Eq. (3.17). It can be observed in the left-hand panel of Fig. 5 that the simulated PDF well fits the reference density.

As a further test, we plot a sample of 3000 terminal values of the 4-step Monte Carlo in the right-hand panel of Fig. 5, verifying that the observed values reflect the moments of Table 2. In particular, it is confirmed in the scatter plot that the margins are negatively skewed, due to the presence of large negative values, and positively correlated as well.

In addition, we use MC simulation to price a range of call options, and compare results with the semi-analytical COS method for option pricing. Results can be read in Table 3: we note that the standard errors are the same up to the forth decimal place regardless of the time discretization used.

4 Multi-currency additive market

In this section we develop a multi-currency market model under the assumption that FX log-rates follow the CSB process introduced above.

	Analytic		MC (1s)		MC (2s)		MC (4s)	
	Y_1	Y_2	Y_1	Y_2	Y_1	Y_2	Y_1	Y_2
Mean	0	0	-0.0002	-0.0001	-0.0002	-0.0001	-0.0001	-0.0002
Variance	0.0191	0.0164	0.0191	0.0164	0.0191	0.0164	0.0190	0.0163
Skewness	-1.5644	-1.7397	-1.5691	-1.7378	-1.5857	-1.7522	-1.5407	-1.6995
Exc.Kurtosis	5.2755	6.2196	5.2601	6.1502	5.6576	6.4202	5.0743	5.7681
Correlation	0.6203		0.6136		0.6198		0.6167	
CPU time			1.04		2.01		3.82	

Table 2: First four standardized moments and correlation of the (centered) CSB process: analytical values (see Eqs. (3.2), (3.8)) and the corresponding MC estimates. CSB parameters: $\mathcal{P}_{\text{EQ}}$ from Table 1. Time horizon: 1 year. Equally-spaced time grids with 1, 2, and 4 steps, respectively. Number of Monte Carlo samples: 10^6. Experiments are performed on a MacBook Pro with a 1.4 GHz Quad-Core Intel Core i5 processor and 8GB of RAM. CPU time in seconds.

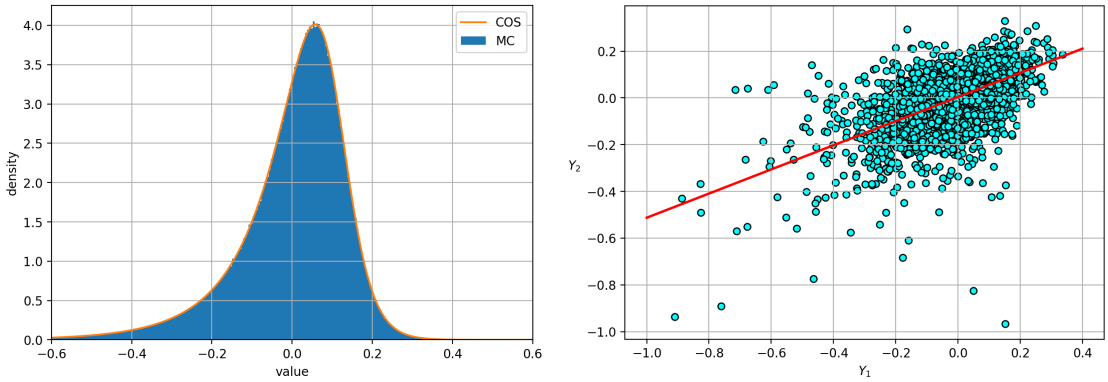

Figure 5: Left-hand panel: time-1 density of the (centered) marginal CSB process recovered via COS inverse transform method and the COS-MC sampling method of Algorithm 1 (number of Monte Carlo samples: 10^6). Right-hand panel: 1-year terminal values of the bidimensional (centered) CSB process simulated with the COS-MC method and estimated regression line of the sample (number of Monte Carlo samples: 3000). For both panels: equally-spaced time grids with 4 steps. CSB parameters: $\mathcal{P}_{\text{EQ}}$ (see Table 1).

Strike	Ref. price	MC price (1s)	MC price (2s)	MC price (4s)	SE
80	20.6453	20.6302	20.6344	20.6285	0.0111
90	11.8326	11.8218	11.8242	11.8181	0.0092
100	4.8143	4.8055	4.8086	4.8022	0.0063
110	1.0717	1.0687	1.0660	1.0676	0.0031
120	0.1411	0.1401	0.1401	0.1401	0.0012

Table 3: Call option prices computed with the marginal CSB process. Ref. prices: call prices recovered through COS method (Fang and Oosterlee, 2009). MC prices: call prices recovered through the COS-MC sampling method of Algorithm 1. Equally-spaced time grids with 1, 2, and 4 steps, respectively. Number of Monte Carlo samples: 10^6. SE: standard error. CSB parameters: $\mathcal{P}_{\text{EQ}}$ in Table 1. Time horizon: 1 year.

4.1 General setting

Consider a financial market $\mathcal{M}$ composed by $n+1$ traded assets with price processes

$$\hat{A}_0(t) = e^{r_d t}, \qquad \hat{A}_{f_j,d}(t) = e^{r_{f_j} t}\, A_{f_j,d}(t), \qquad t \geqslant 0, \qquad j = 1, \dots, n,$$

with $A_{f_j,d}(t)$ denoting the FX rate that expresses the amount of domestic currency d needed to buy one unit of foreign currency f_j. The parameters r_d and r_{f_j} denote the continuously compounded risk free rates of interest – assumed to be constant – in the domestic and foreign market, respectively. We assume that $\mathcal{M}$ is free of arbitrage, so that $\hat{A}_{f_j,d}(t)/\hat{A}_0(t)$, $t \geqslant 0$, $j = 1, \dots, n$, are $\mathbb{Q}^d$-martingales, where $\mathbb{Q}^d$ is a risk neutral measure in the domestic market.

For the construction of our market model, we highlight the well-known triangulation property of currency markets, which any FX log-rate process needs to satisfy. Let $A_{f_k,f_l}(t)$, for some k, l, be the FX rate between two foreign currencies, henceforth the cross rate between f_k and f_l. Then, equality

$$A_{f_k,f_l}(t) = A_{f_k,d}(t)/A_{f_l,d}(t), \qquad t \geqslant 0,$$

holds. Note that by analogy with the d-denominated market $\mathcal{M}$, $\hat{A}_{f_k,f_l}(t) = e^{r_{f_k}t} A_{f_k,f_l}(t)$ is an asset in a f_l-denominated market.

4.2 Additive specification

Consider a currency triangle denoted as $f_k/f_l/d$. We are interested in calibrating the joint process $\boldsymbol{A}(t) = \big(A_{f_k,d}(t), A_{f_l,d}(t)\big)^\top$, $t \geqslant 0$, assuming that FX log-rates follow additive dynamics. Information on the margins $A_{f_k,d}(t)$ and $A_{f_l,d}(t)$ can be extracted from liquid single-asset options written on these instruments. The dependence structure of $\boldsymbol{A}(t)$ can then be recovered by means of options written on the cross rate $A_{f_k,f_l}(t)$, which we assume to be liquidly traded in the f_l-denominated market.

Let $\boldsymbol{Y}(t) = \big(Y_{f_k,d}(t), Y_{f_l,d}(t)\big)^\top$, $t \geqslant 0$, be the additive process that drives the dynamics of $\boldsymbol{A}(t)$. Then, the marginal processes are of the form

$$A_{f_j,d}(t) = A_{f_j,d}(0)\, \exp\left((r_d - r_{f_j})t - \psi^d_{Y_{f_j,d}}(-\mathrm{i};t) + Y_{f_j,d}(t)\right), \qquad j = k, l; \tag{4.1}$$

in the above expression, $\psi^d_{Y_{f_j,d}}(\cdot;t)$ denotes the characteristic exponent, under $\mathbb{Q}^d$, of the time-inhomogeneous process $Y_{f_j,d}(t)$ evaluated at time t. We note that the proposed market model is incomplete and consequently the risk neutral measure is not unique. Hence, we follow standard practice for incomplete markets and fix the risk neutral measure through the prices of derivative contracts traded in the market. By construction, the cross rate $A_{f_k,f_l}(t)$ is of the form

$$A_{f_k,f_l}(t) = A_{f_k,f_l}(0)\, \exp\left((r_{f_l} - r_{f_k})t - \psi^{f_l}_{Y_{f_k,f_l}}(-\mathrm{i};t) + Y_{f_k,f_l}(t)\right), \tag{4.2}$$

where $Y_{f_k,f_l}(t) = Y_{f_k,d}(t) - Y_{f_l,d}(t)$.

In order to extract the joint dynamics of FX rates, it is crucial to connect the domestic and the foreign pricing measures. This is formalized in Proposition 4.1, which is a direct application of Theorem 1 in Geman et al. (1995).

Proposition 4.1. Consider an asset $\hat{A}_{f_l,d}(t)$, for some l, belonging to market $\mathcal{M}$. Then there exists a probability measure $\mathbb{Q}^{f_l} \sim \mathbb{Q}^d$ defined by the density process

$$\left.\frac{\mathrm{d}\mathbb{Q}^{f_l}}{\mathrm{d}\mathbb{Q}^d}\right|_{\mathcal{F}_t} = \frac{e^{r_{f_l}t} A_{f_l,d}(t)}{e^{r_d t} A_{f_l,d}(0)}, \qquad t \geqslant 0, \tag{4.3}$$

such that $\hat{A}_{f_j,d}(t)/\hat{A}_{f_l,d}(t)$, $t \geqslant 0$, $\forall j$, are $\mathbb{Q}^{f_l}$-martingales. Moreover, $\mathbb{Q}^{f_l}$ is a risk neutral measure for the f_l-denominated market.

In virtue of both the triangulation property of FX rates and Proposition 4.1, we can price any derivative written on $A_{f_k,f_l}(t)$, traded in the l-denominated market, through the $\mathbb{Q}^d$-dynamics of $\big(A_{f_k,d}(t), A_{f_l,d}(t)\big)$. Inspired by the results presented in Eberlein et al. (2009), the following proposition shows the connection between the $\mathbb{Q}^d$-dynamics of $\boldsymbol{Y}(t)$ and the $\mathbb{Q}^{f_l}$-dynamics of $Y_{f_k,f_l}(t)$, which is crucial to calibrate the dependence structure of $\boldsymbol{Y}(t)$.

Proposition 4.2. Let $\boldsymbol{Y}(t) = \big(Y_{f_k,d}(t), Y_{f_l,d}(t)\big)^\top$, $t \geqslant 0$, be the $\mathbb{R}^2$-valued additive process that drives the dynamics of $\boldsymbol{A}(t)$. Then, the process $Y_{f_k,f_l}(t)$ that drives the dynamics of the cross rate $A_{f_k,f_l}(t)$ in Eq. (4.2) is an additive process under $\mathbb{Q}^{f_l}$ with characteristic exponent

$$\psi^{f_l}_{Y_{f_k,f_l}}(u;t) = \psi^d_{\boldsymbol{Y}}(-\mathrm{i}\bar{\boldsymbol{h}} + \bar{\boldsymbol{u}}; t) - \psi^d_{\boldsymbol{Y}}(-\mathrm{i}\bar{\boldsymbol{h}}; t), \quad \bar{\boldsymbol{h}} = (0,1)^\top, \ \bar{\boldsymbol{u}} = (u,-u)^\top, \ u \in \mathbb{R}.$$

Due to the invariance with respect to linear transformations and the invariance with respect to the Esscher change of measure, additive processes are symmetric with respect to inversion and triangulation (see definition in De Col et al., 2013). As such, they can be consistently calibrated to FX markets, which is what Section 5 is devoted to.

5 Empirical analysis

In this section, we analyze the ability of the CSB market model presented in Section 4 to fit market data, and show the non stationary nature of multi-currency markets. For our analysis, we consider six currency triangles, the details of which are reported in Table 4. In order to test a variety of market conditions, we have included both developed and developing countries. Market quotes for the three implied volatility surfaces associated to each triangle are extracted from Bloomberg. The volatility smiles composing these surfaces are quoted at six expiries: 1, 2, 3, 6, 12, and 18 months.

Following FX conventions, the implied volatilities are associated to the quoted option delta neutral calls and puts, and the 10 and 25 delta calls and puts, instead of a strike price. The conversion to strike terms is operated by means of the routine of Bossens et al. (2010). Finally, we ensure that all volatility surfaces are arbitrage free applying the SVI method of Gatheral and Jacquier (2014) when required.

For practical purposes, we assume that the CSB process has time-1 multivariate inverse Gaussian distribution, so that the characteristic exponent of the joint risk drivers of the FX log-rates is given by Eq. (3.14). Similarly, we assume inverse Gaussian subordinators for any other benchmark model used in the analysis. Finally, all calibrations are performed with the trust region reflective algorithm provided by the SciPy package in Python.

Calibration of the option price surfaces To highlight the importance of modelling a rich time inhomogeneous dynamics in FX markets, we compare the calibration performances of the CSB model and its closest competitors, i.e., the SSBM model of Semeraro (2022), the BB model and the LS model for Lévy processes by Ballotta et al. (2017) and Luciano and Semeraro (2010) (see Appendix A).

In what follows, we keep the notation $f_k/f_l/d$ for a generic currency triangle composed by a domestic currency d and foreign currencies f_k and f_l. Let $\{C^m_x, m \in \mathbb{M}_x, x \in \mathcal{X}\}$, for $\mathcal{X} := \{(f_k,d),(f_l,d),(f_k,f_l)\}$, be the market prices of the call options associated to the implied volatility surfaces of each leg of the triangles, and denote by $\hat{C}^m_x(\mathcal{P})$ the corresponding model prices computed with the COS method of Fang and Oosterlee (2009).

We denote with $\mathbb{M}$ the sets of traded instruments used in the calibration, which are identified by certain strikes and maturities. For the case of the surface calibration described in this paragraph, we consider all the plain vanilla options in our dataset, namely 30 options for each FX rate. Let $\mathcal{P}$ be the set of model parameters and $\overline{\mathcal{P}}$ a model-specific feasible region of solutions.

Following Cont and Tankov (2004), we calibrate prices scaled by Black-Scholes vegas, denoted with

	FX rates		
Triangles	Leg 1	Leg 2	Leg 3 (cross rate)
USD (United States dollar) EUR (Euro) CHF (Swiss frank)	USDCHF	EURCHF	USDEUR
AUD (Australian dollar) CNY (Chinese yuan) GBP (British pound)	AUDGBP	CNYGBP	AUDCNY
JPY (Japanese yen) CAD (Canadian dollar) USD (United States dollar)	JPYUSD	CADUSD	JPYCAD
ZAR (South African rand) GBP (British pound) EUR (Euro)	ZAREUR	GBPEUR	ZARGBP
INR (Indian rupee) GBP (British pound) ZAR (South African rand)	INRZAR	GBPZAR	INRGBP
GBP (British pound) ILS (Israeli new shekel) USD (United States dollar)	GBPUSD	ILSUSD	GBPILS

Table 4: Currency Triangles. First column: currencies composing each triangle. Columns 2-4: FX rates associated to each triangle (e.g., the triangle USD/EUR/CHF includes the FX rate between U.S. dollar and Swiss frank, the FX rate between Euro and Swiss frank, and the FX rate between U.S. dollar and Euro, respectively).

Triangle: USD/EUR/CHF				
Idiosyncratic fit			Systematic fit	
	USDCHF	EURCHF		
μ_j	-0.0176	-0.0616	a	1.5094
σ_j	0.0961	0.0752	ρ	0.4907
κ_j	0.4072	0.345	q_Z	1.1529
q_j	0.7624	0.6817		

Triangle: AUD/CNY/GBP				
Idiosyncratic fit			Systematic fit	
	AUDGBP	CNYGBP		
μ_j	-0.0129	0.0134	a	0.8618
σ_j	0.0998	0.0928	ρ	0.6357
κ_j	0.5516	0.4641	q_Z	1.3644
q_j	0.8993	0.933		

Triangle: JPY/CAD/USD				
Idiosyncratic fit			Systematic fit	
	JPYUSD	CADUSD		
μ_j	0.0065	-0.0349	a	1.0438
σ_j	0.1041	0.0729	ρ	0.2736
κ_j	0.916	0.5949	q_Z	1.1769
q_j	1.3732	0.8852		

Triangle: ZAR/GBP/EUR				
Idiosyncratic fit			Systematic fit	
	ZAREUR	GBPEUR		
μ_j	-0.1417	-0.0591	a	0.4331
σ_j	0.1821	0.0791	ρ	0.999
κ_j	0.3597	0.2931	q_Z	1.4455
q_j	0.9819	1.1532		

Triangle: INR/GBP/ZAR				
Idiosyncratic fit			Systematic fit	
	INRZAR	GBPZAR		
μ_j	0.0985	0.0685	a	1.7918
σ_j	0.2086	0.1964	ρ	0.9638
κ_j	0.2107	0.2854	q_Z	1.0364
q_j	1.0539	0.8643		

Triangle: GBP/ILS/USD				
Idiosyncratic fit			Systematic fit	
	GBPUSD	ILSUSD		
μ_j	-0.1207	-0.0696	a	0.665
σ_j	0.1183	0.1168	ρ	0.999
κ_j	0.2554	0.7111	q_Z	0.9801
q_j	1.032	0.7018		

Table 5: Calibrated parameters of the component-wise Sato-subordinated Brownian motion for each currency triangle, under full-surface calibration (5.1). Triangles illustrated in Table 4.

V, as opposed to implied volatilities. Then, we solve the triangle-based calibration problem

$$\mathcal{P}^* = \underset{\mathcal{P}\in\overline{\mathcal{P}}}{\operatorname{argmin}} \sum_{x\in\mathcal{X}} \sum_{m\in\mathbb{M}_x} \left(\frac{\hat{C}_x^m(\mathcal{P}) - C_x^m}{V_x^m} \right)^2. \tag{5.1}$$

The calibrated parameters of the CSB process are reported in Table 5, while the mean absolute errors of the calibrations of all models are displayed in Fig. 6. We also show the smile fit of the CSB and SSBM models for the case of GBPILS rate in Fig. 7. We clearly see that the CSB performs better than the benchmark models in the analyzed cases. It also possesses a parsimonious parametrization.

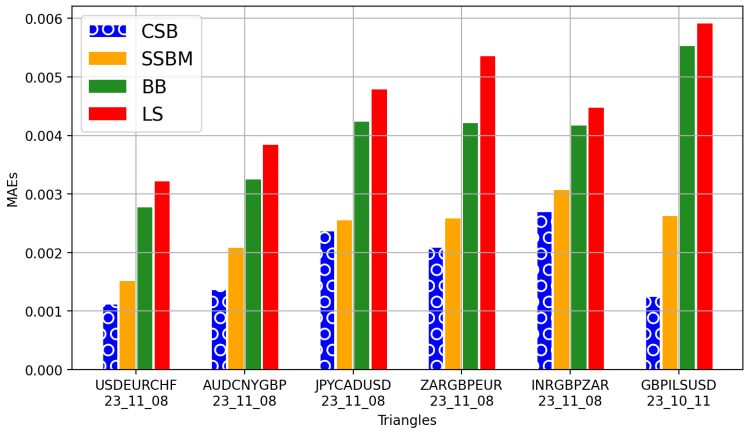

Figure 6: Mean Absolute Errors of the full-surface calibrations (5.1) of each model. Errors computed as differences between model and market prices (scaled by Black-Scholes vegas) of the options on the FX rates composing each triangle. Models of interest: Component-wise Sato-subordinated Brownian motion (CSB), SSBM model (Semeraro, 2022), BB model (Ballotta et al., 2017), and LS model (Luciano and Semeraro, 2010). Triangles illustrated in Table 4.

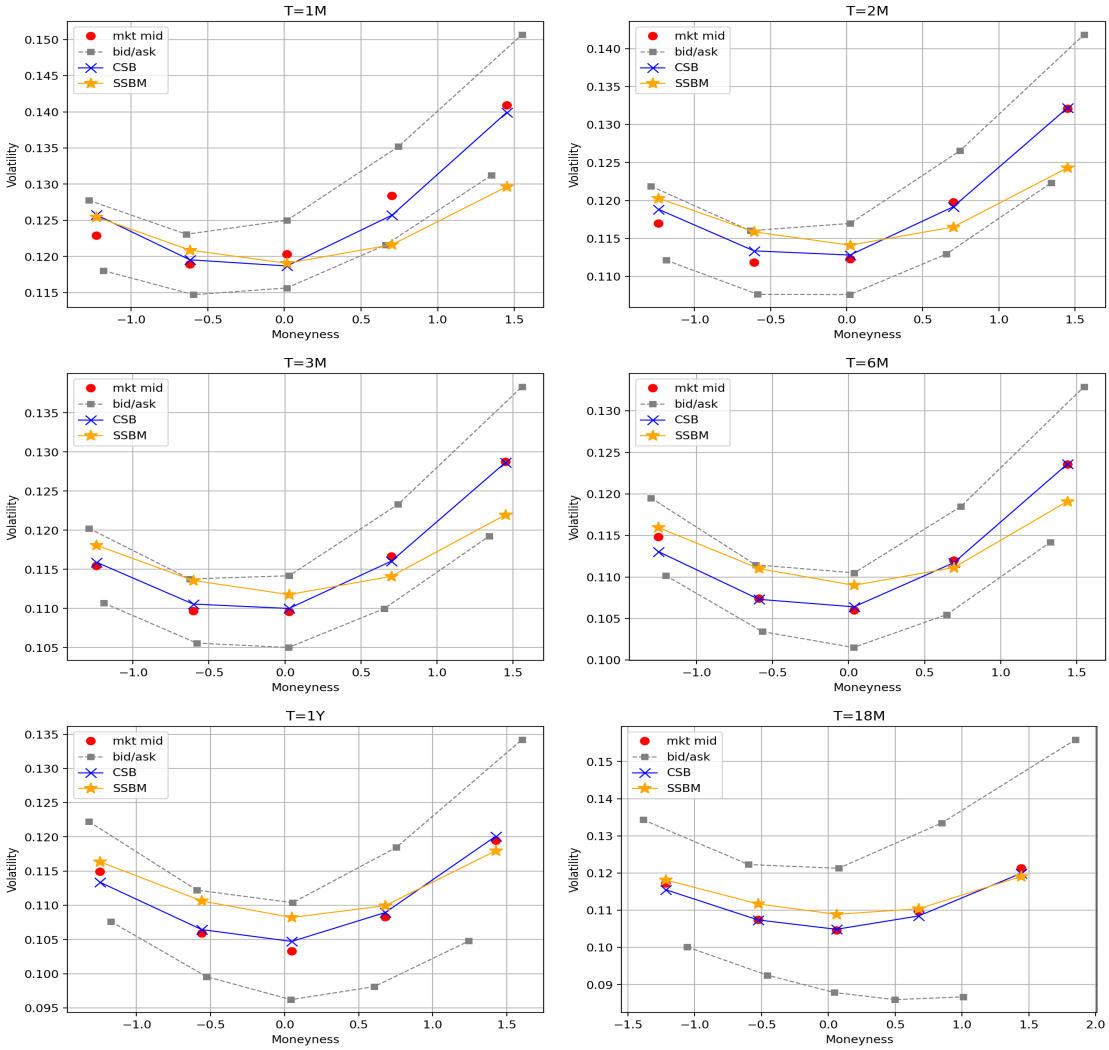

Figure 7: Calibrated implied volatility smiles of CSB and SSBM models under the full-surface fit (5.1) in the case of FX rate GBPILS (triangle GBP/ILS/USD, see Table 4).

In fact, in this bivariate setting the numbers of parameters of the studied models are $|\mathcal{P}(\text{CSB})| = 4n + 2 + n(n-1)/2 = 11$, $|\mathcal{P}(\text{SSBM})| = 3n + 2 + n(n-1)/2 = 9$, $|\mathcal{P}(\text{BB})| = 4n + 3 = 11$, and

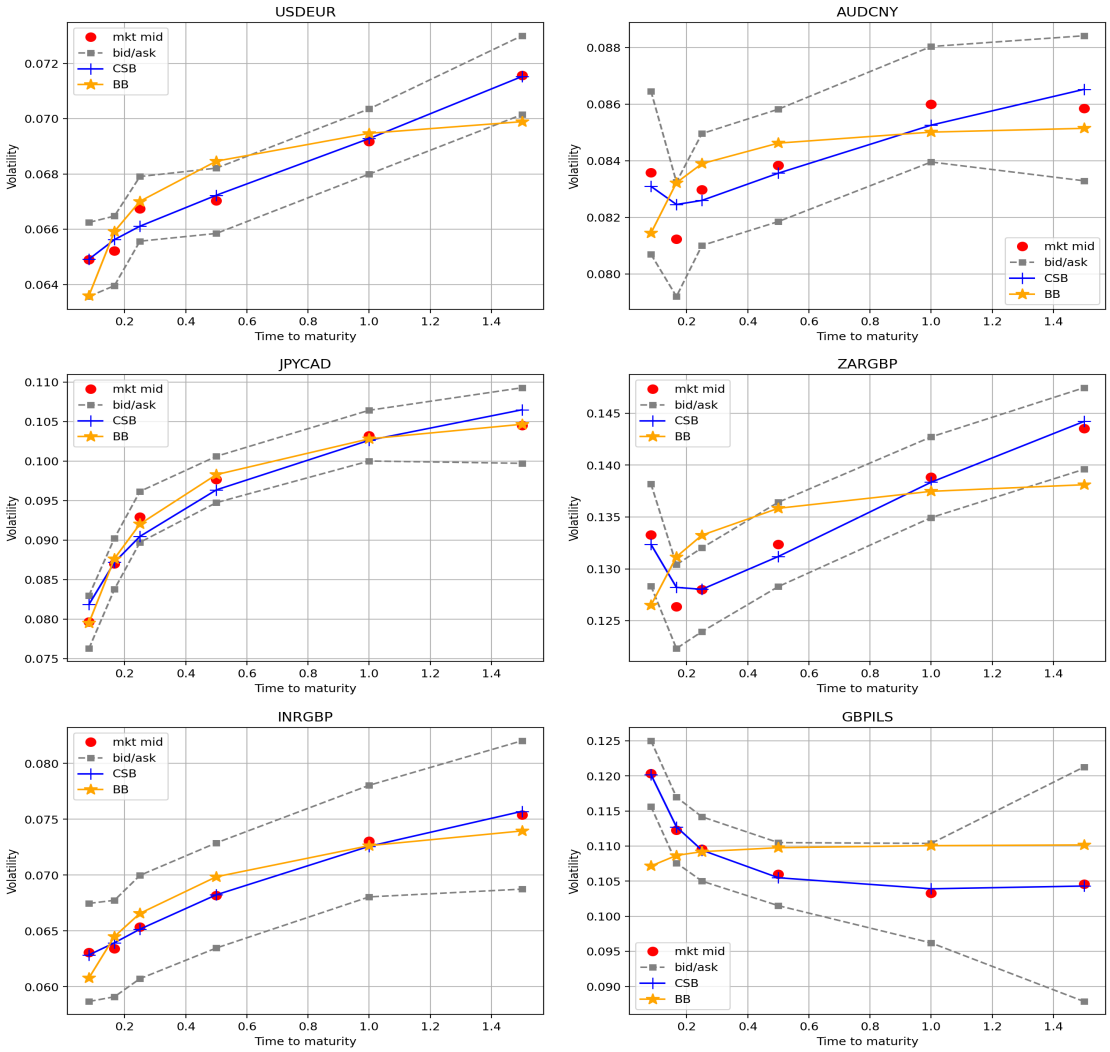

Figure 8: Calibrated ATM implied volatilities of the CSB and BB models – calibration performed only on the six ATM options of each cross rate (see Table 4), separately (calibration problem (5.2)).

$|\mathcal{P}(\mathrm{LS})| = 3n+1+n(n-1)/2 = 8$. As such, the only significant difference between these constructions is the greater ability of the CSB process to model time inhomogeneity.

In the following, we design a number of sub-calibration exercises aimed at gaining in-depth understanding of the impact of time inhomogeneity on the calibration performance.

Limitations of the stationarity property As a first test, we analyze how the stationarity of Lévy processes limits their fit in a multi-maturity calibration. To this aim, we compare the performance of the CSB model against the BB model (that is, the Lévy model that enjoys the best fit according to Fig. 6), by calibrating the two constructions to the ATM option prices on the cross rate of each triangle, separately. Thus, the calibration problem for any cross rate is formulated as

$$\mathcal{P}^* = \operatorname*{argmin}_{\mathcal{P}\in\overline{\mathcal{P}}} \sum_{m\in\mathbb{M}^{\mathrm{ATM}}_{f_k,f_l}} \left(\frac{\hat{C}^m_{f_k,f_l}(\mathcal{P}) - C^m_{f_k,f_l}}{V^m_{f_k,f_l}} \right)^2, \tag{5.2}$$

where $\mathbb{M}^{\mathrm{ATM}}$ denotes the set of quoted ATM call options for all the available maturities.

Even with such a little calibration size, the BB model cannot replicate the shape of market implied volatilities, and it almost lies out of the bid-ask spread, as can be observed in Fig. 8. As a matter

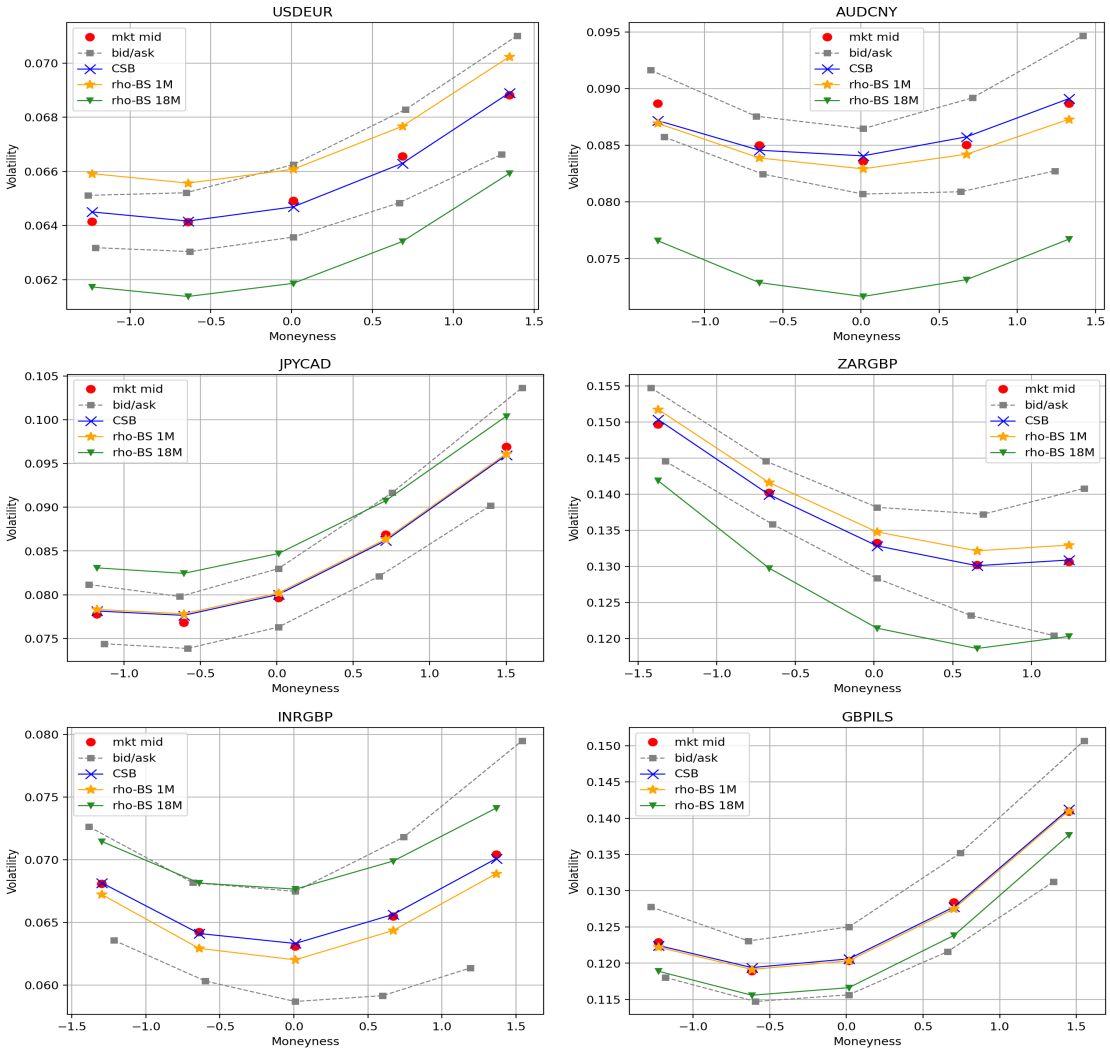

Figure 9: Implied volatility smiles of options written on the cross rates (see Table 4), under 1-month calibrations of the CSB model with different dependence calibrations (calibration problem (5.3)). CSB line: CSB smile obtained by calibrating the model dependence structure to all the available option quotes on the cross rate (triangle-based calibration). rho-BS 1M and rho-BS 18M lines: CSB smile obtained by calibrating the model correlation to a Black-Scholes implied correlation from an ATM option price at a reference maturity (1 month and 18 months, respectively).

of fact, it displays roughly the same shape in all the analyzed cases, despite the different market conditions. As expected, the CSB model is more accurate. This test remarks that stationary models are not suited to calibrate market data with different expiries. On the other hand, it is interesting to see the remarkable improvement offered by the parameters controlling time inhomogeneity.

Time dependence of the correlation Besides stationarity, another common restriction for multivariate pricing models is that the resulting correlation is constant over time. Thus, we show the consequences of calibrating the dependence structure of a model under this restriction.

Assume that we are interested in calibrating a pricing model to the 1-month option prices only. Also, consider a situation in which options on the cross rate are not liquid. In this case, market providers sometimes quote an ATM Black-Scholes implied correlation at some reference maturity.

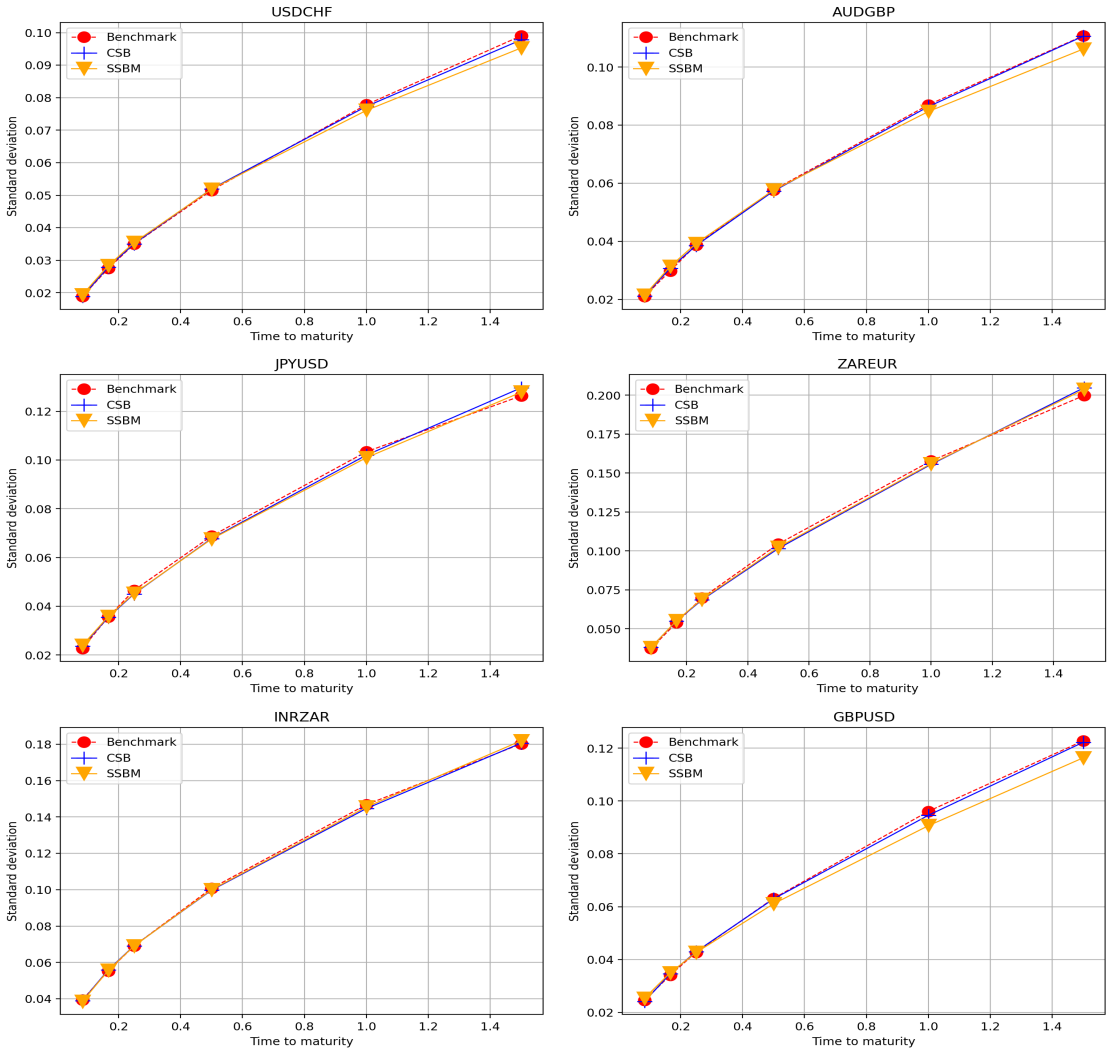

Figure 10: Standard deviation (derived from Eq. (3.5)) of the calibrated CSB and SSBM models. CSB and SSBM lines: model standard deviations calibrated in the full-surface fit – standard deviation of the first leg of each triangle (see Table 4). Benchmark line: proxies of the market standard deviation, computed from highly accurate model calibrations to single-maturity sets of options.

With this information available, calibration is of the form

$$\mathcal{P}^* = \underset{\mathcal{P}\in\overline{\mathcal{P}}}{\operatorname{argmin}} \sum_{x\in\mathcal{X}'} \sum_{m\in\mathbb{M}_x^{1\mathrm{M}}} \left(\frac{\hat{C}_x^m(\mathcal{P}) - C_x^m}{V_x^m} \right)^2 + \left(\rho_{1\mathrm{M}}^{\mathrm{CSB}}(\mathcal{P}) - \rho_\tau^{\mathrm{BS}} \right)^2, \tag{5.3}$$

where $\mathcal{X}' = \{(f_k, d), (f_l, d)\}$, $\mathbb{M}^{1\mathrm{M}}$ denotes the set of traded calls with 1-month expiry, and ρ_τ^{BS} is the Black-Scholes implied correlation at maturity τ. For the purpose of our analysis, we compare the calibration performance of the CSB model in the case in which τ is set to 1 month and 18 months, respectively. We also calibrate the CSB model under a triangle-based fit (with calls in the sets $\mathbb{M}_{f_k,d}^{1\mathrm{M}}$, $\mathbb{M}_{f_l,d}^{1\mathrm{M}}$, and $\mathbb{M}_{f_k,f_l}^{1\mathrm{M}}$) as a reference, resembling a situation of full liquidity of all FX options.

We plot the calibrated smile of the cross rates under the different calibration types in Fig. 9. Not surprisingly, the calibration of the full-liquidity CSB specification is very accurate. Furthermore, although a single Black-Scholes implied correlation is not expected to capture the whole dependence structure between log-rates, matching the CSB correlation to the 1-month ATM Black Scholes implied correlation provides a good fit. This changes dramatically if we consider τ set at 18 months. As the latter case produces implied volatilities out of the bid-ask spread, our tests shows that assuming a

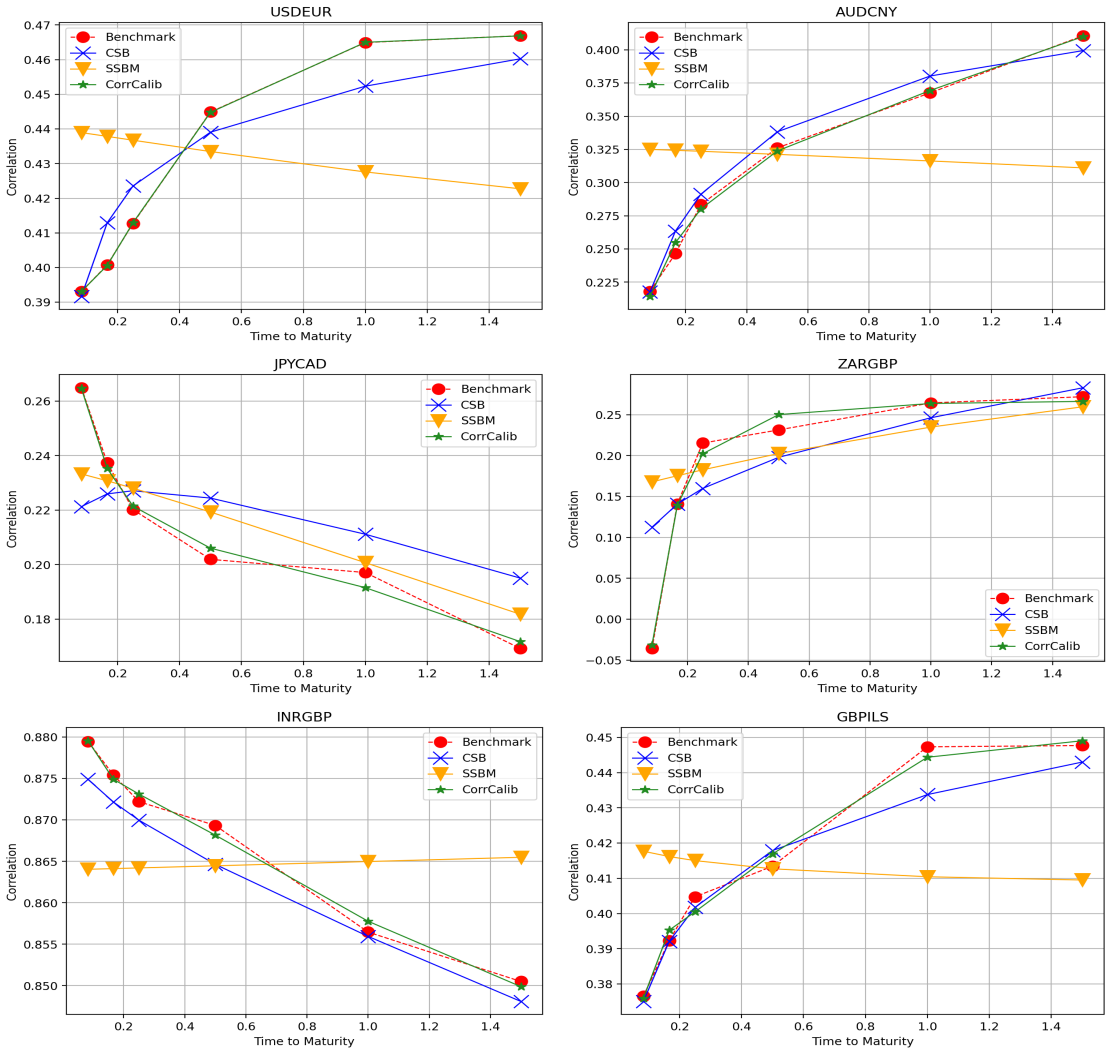

Figure 11: Correlation (from Eq. (3.8)) of the calibrated CSB and SSBM models. CSB and SSBM lines: model correlations calibrated in the full-surface fit. Benchmark line: proxies of the market correlations, computed with highly accurate model calibrations to single-maturity sets of options. CorrCalib line: CSB correlations – model fitted only to the benchmark term structure of correlations, disregarding quoted option prices (calibration problem (5.4)).

constant correlation can lead to arbitrage when the time horizons are not consistent.

Impact of model statistics on calibration We study the impact of the model standard deviation and correlation on the cross-rate surface calibrations: the intuition is that the closer a model statistic is to its benchmark, the more this statistic is a core driver of the full-surface calibration performance showcased in Fig. 6.

This exercise is organized as follows. First, we compute from Eqs. (3.5) and (3.8) the term structures of standard deviation and correlations, which we use as proxies for their market counterparts. These correspond to the model statistics resulting from highly accurate single-maturity calibrations. Secondly, we compute the term structures of standard deviation and correlations resulting from the full-surface calibration, under both the CSB and SSBM models. As the SSBM model has only one time scaler, it is interesting to check whether this parameter is mostly devoted to fit either the market standard deviation, the market correlation, or both.

Results are shown in Figs. 10 and 11. In details, from the quality of the fit observed in Fig. 10, it is clear that standard deviation retains a major role in the volatility surface calibration. However,

Fig. 11 shows that an accurate fit of the term structure of correlation is hardly achieved, highlighting a relatively less important role of correlation compared to the standard deviation. This is particularly visible for the SSBM model: due to its less sophisticated correlation structure in which time inhomogeneity is controlled by a common parameter (namely, Eq. (3.8) in the special case $q_1 = \ldots = q_n = q_Z$), the calibration procedure ends up 'sacrificing' the correlation fit in the interest of the global fit. Our results show that in order to truly capture the observed term structure of correlations, we need a component-wise parametrization to model time inhomogeneity, as the CSB model allows for.

To conclude, since the benchmark term structure of correlations is not always matched by the CSB correlations, we verify that this is not due to restrictions in the admissible range of values of the parameters. Thus, we perform a simple calibration of the dependence structure, designed as

$$\mathcal{P}^* = \underset{\mathcal{P}\in\overline{\mathcal{P}}}{\operatorname{argmin}} \sum_{\tau\in\mathcal{T}} \left(\rho_\tau^{\mathrm{CSB}}(\mathcal{P}) - \rho_\tau^{\mathrm{bench}}\right)^2. \tag{5.4}$$

where $\mathcal{T}$ is the set of available maturities.

We can see from Fig. 11 that the CSB model has, in principle, a sufficiently rich parametrization to fit the benchmark almost exactly. Therefore, the relatively poor fit of the surface-calibrated correlation with respect to the standard deviation is in fact due to the greater relevance held by the model standard deviation for the global fit of the implied volatility surface.

6 Conclusions

This paper introduces a new multivariate pure jump process by additive subordination, namely the CSB process, which is analytically tractable, parsimonious and able to model time inhomogeneity both at a marginal and joint level. Furthermore, the parameters have a clear role in controlling the statistics of the underlying distribution, the dependence structure and the time inhomogeneity, as shown by the sensitivity analysis carried out in this paper.

The results from the empirical study based on market quotes of options on currency triangles confirm the ability of the CSB model to produce an accurate fit of the implied volatility surfaces also for the inferred cross-rates. In particular, our empirical exercises confirm that the success of the CSB is mainly due to its component-wise non-stationary design. Furthermore, our results suggest that a model marginal standard deviation which accurately reproduces its market counterpart is more relevant than the fit of the correlation term structure for the overall accuracy of the calibration, although the component-wise time inhomogeneity of the CSB successfully manages to fit this latter feature as well.

We conclude with some remarks on possible directions for future research which stem from this paper. In terms of option pricing, the FX market would be the ideal context to further test the robustness of the CSB model by increasing the number of triangles associated to any given currency, i.e., the number of dimensions, due to the high liquidity which characterizes it. Furthermore, the valuation of options on cryptocurrencies, and related triangles can be of great interest in order to assess the calibration performance of the CSB in more volatile markets.

Furthermore, given the ability of our model to generate a term structure for the correlation (and other dependence features as well, although not studied in this paper), the CSB process could find a direct application in problems of portfolio optimization. Thanks to the factor construction underlying the CSB process, the estimation of its parameters under the real world measure could be carried out

using the approach of Ballotta et al. (2019a), offering a viable alternative to portfolio models based on copulas as for example in Babaei et al. (2015). As correlation is a fundamental ingredient of wrong way risk, the CSB process could be a useful model for counterparty credit risk quantification as well (see Ballotta et al., 2019b, and reference therein).

A related application of interest for the CSB process which looks beyond the area of finance could be in the context of inventory management policies. A realistic multivariate model for spot prices could indeed support the selection process of an optimal procurement strategy involving baskets of dependent commodities in the spirit of Chiarolla et al. (2015).

Moreover, we note that tempered stable distributions, such as the ones used in this work, have also been used in command and control problems by Kalloniatis et al. (2020) for decision making in organizations such as the military and emergency response services, but also traffic control and energy and safety operations rooms. In Kalloniatis et al. (2020), the tempered stable distributed jump process reflects the 'leaping ahead' in the decision process of experienced decisions makers to anticipate emerging situations during time of crisis. In this context, we envisage a potential for the application of our multivariate framework for coordinated decisions between more organizations of this kind.

Finally, it is of interest to pursue more theoretical lines of research focused on generalizing our factor-based construction to processes beyond the class of additive semimartingales, following Jarrow et al. (2009) for example. We note though that for the application to multi-FX markets, it is fundamental to verify first that the symmetries with respect to inversion and triangulation hold.

References

Amici, G., Brandimarte, P., Messeri, F., and Semeraro, P. (2023). Multivariate Lévy Models: Calibration and Pricing. *arXiv preprint arXiv:2303.13346*.

Azzone, M. and Baviera, R. (2023). A fast Monte Carlo scheme for additive processes and option pricing. *Computational Management Science*, 20(1):31.

Babaei, S., Sepehri, M. M., and Babaei, E. (2015). Multi-objective portfolio optimization considering the dependence structure of asset returns. *European Journal of Operational Research*, 244(2):525–539.

Backus, D. K., Foresi, S., and Wu, L. (2004). Accounting for biases in Black-Scholes. *Available at SSRN 585623*.

Ballotta, L. and Bonfiglioli, E. (2016). Multivariate asset models using Lévy processes and applications. *The European Journal of Finance*, 22(13):1320–1350.

Ballotta, L., Deelstra, G., and Rayée, G. (2017). Multivariate FX models with jumps: Triangles, quantos and implied correlation. *European Journal of Operational Research*, 260(3):1181–1199.

Ballotta, L., Fusai, G., Loregian, A., and Perez, M. F. (2019a). Estimation of multivariate asset models with jumps. *Journal of Financial and Quantitative Analysis*, 54(5):2053–2083.

Ballotta, L., Fusai, G., and Marazzina, D. (2019b). Integrated structural approach to Credit Value Adjustment. *European Journal of Operational Research*, 272(3):1143–1157.

Ballotta, L. and Kyriakou, I. (2014). Monte Carlo simulation of the CGMY process and option pricing. *Journal of Futures Markets*, 34(12):1095–1121.

Barndorff-Nielsen, O. E. (1997). Processes of normal inverse Gaussian type. *Finance and Stochastics*, 2(1):41–68.

Bossens, F., Rayée, G., Skantzos, N. S., and Deelstra, G. (2010). Vanna-volga methods applied to FX derivatives: from theory to market practice. *International Journal of Theoretical and Applied Finance*, 13(08):1293–1324.

Brigo, D., Pisani, C., and Rapisarda, F. (2019). The multivariate mixture dynamics model: shifted dynamics and correlation skew. *Annals of Operations Research*, 299(1–2):1411–1435.

Broadie, M. and Kaya, Ö. (2006). Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations Research*, 54(2):217–231.

Carr, P., Geman, H., Madan, D. B., and Yor, M. (2002). The fine structure of asset returns: An empirical investigation. *Journal of Business*, 75(2):305–332.

Carr, P., Geman, H., Madan, D. B., and Yor, M. (2007). Self-decomposability and option pricing. *Mathematical Finance*, 17(1):31–57.

Carr, P. and Wu, L. (2003). The finite moment log stable process and option pricing. *The Journal of Finance*, 58(2):753–777.

Chiarolla, M. B., Ferrari, G., and Stabile, G. (2015). Optimal dynamic procurement policies for a storable commodity with Lévy prices and convex holding costs. *European Journal of Operational Research*, 247(3):847–858.

Clark, P. K. (1973). A subordinated stochastic process model with finite variance for speculative prices. *Econometrica*, 41(1):135–155.

Cont, R. and Tankov, P. (2004). Nonparametric calibration of jump-diffusion option pricing models. *The Journal of Computational Finance*, 7:1–49.

De Col, A., Gnoatto, A., and Grasselli, M. (2013). Smiles all around: FX joint calibration in a multi-Heston model. *Journal of Banking & Finance*, 37(10):3799–3818.

Eberlein, E., Glau, K., and Papapantoleon, A. (2010). Analysis of Fourier transform valuation formulas and applications. *Applied Mathematical Finance*, 17(3):211–240.

Eberlein, E. and Madan, D. B. (2009). Sato processes and the valuation of structured products. *Quantitative Finance*, 9(1):27–42.

Eberlein, E., Papapantoleon, A., and Shiryaev, A. N. (2009). Esscher transform and the duality principle for multidimensional semimartingales. *The Annals of Applied Probability*, 19(5):1944–1971.

Engle, R. F. (2000). The econometrics of ultra-high-frequency data. *Econometrica*, 68(1):1–22.

Engle, R. F. and Russell, J. R. (1998). Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica*, 66(5):1127–1162.

Fang, F. and Oosterlee, C. W. (2009). A novel pricing method for European options based on Fourier-cosine series expansions. *SIAM Journal on Scientific Computing*, 31(2):826–848.

Fontana, C., Gnoatto, A., and Szulda, G. (2022). CBI-time-changed Lévy processes for multi-currency modeling. *Annals of Operations Research*, pages 1–26.

Gatheral, J. and Jacquier, A. (2014). Arbitrage-free SVI volatility surfaces. *Quantitative Finance*, 14(1):59–71.

Geman, H., El Karoui, N., and Rochet, J.-C. (1995). Changes of numeraire, changes of probability measure and option pricing. *Journal of Applied probability*, 32(2):443–458.

Glasserman, P. and Liu, Z. (2010). Sensitivity estimates from characteristic functions. *Operations Research*, 58(6):1611–1623.

Guillaume, F. (2018). Multivariate option pricing models with Lévy and Sato VG marginal processes. *International Journal of Theoretical and Applied Finance*, 21(02):1850007.

Jarrow, R. A., Protter, P., and Sayit, H. (2009). No arbitrage without semimartingales. *The Annals of Applied Probability*, 19(2):596–616.

Kalloniatis, A. C., McLennan-Smith, T. A., and Roberts, D. O. (2020). Modelling distributed decision-making in command and control using stochastic network synchronisation. *European Journal of Operational Research*, 284(2):588–603.

Kang, C., Kang, W., and Lee, J. M. (2017). Exact simulation of the Wishart multidimensional stochastic volatility model. *Operations Research*, 65(5):1190–1206.

Li, J., Li, L., and Mendoza-Arriaga, R. (2016). Additive subordination and its applications in finance. *Finance and Stochastics*, 20(3):589–634.

Luciano, E. and Schoutens, W. (2006). A multivariate jump-driven financial asset model. *Quantitative Finance*, 6(5):385–402.

Luciano, E. and Semeraro, P. (2010). Multivariate time changes for Lévy asset models: Characterization and calibration. *Journal of Computational and Applied Mathematics*, 233(8):1937–1953.

Madan, D. B., Carr, P., and Chang, E. (1998). The Variance Gamma process and option pricing. *European Finance Review*, 2:79–105.

Marena, M., Romeo, A., and Semeraro, P. (2018a). Multivariate factor-based processes with Sato margins. *International Journal of Theoretical and Applied Finance*, 21(01):1850005.

Marena, M., Romeo, A., and Semeraro, P. (2018b). Pricing multivariate barrier reverse convertibles with factor-based subordinators. *Journal of Computational Finance*, 21(5).

Mendoza-Arriaga, R. and Linetsky, V. (2016). Multivariate subordination of markov processes with financial applications. *Mathematical Finance*, 26(4):699–747.

Oosterlee, C. W. and Grzelak, L. A. (2019). *Mathematical modeling and computation in finance: with exercises and Python and MATLAB computer codes*. World Scientific.

Papapantoleon, A. (2006). *Applications of semimartingales and Lévy processes in finance: duality and valuation*. PhD thesis, Citeseer.

Pérez-Abreu, V. and Stelzer, R. (2014). Infinitely divisible multivariate and matrix gamma distributions. *Journal of Multivariate Analysis*, 130:155–175.

Samoradnitsky, G. and Taqqu, M. S. (1994). *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance (1st ed.)*. Chapman and Hall.

Sato, K. (1999). *Lévy processes and infinitely divisible distributions*. Cambridge University Press.

Semeraro, P. (2022). Multivariate tempered stable additive subordination for financial models. *Mathematics and Financial Economics*, 16(4):685–712.

Sun, Y., Mendoza-Arriaga, R., and Linetsky, V. (2017). Marshall–Olkin distributions, subordinators, efficient simulation, and applications to credit risk. *Advances in Applied Probability*, 49(2):481–514.

A Overview of other factor-based jump processes

A.1 LS model

The factor-based Lévy process introduced by Luciano and Semeraro (2010) is constructed as follows. Let $\boldsymbol{B}(t)$, $t \geqslant 0$, and $\boldsymbol{B}^{\rho}(t)$, $t \geqslant 0$, be $\mathbb{R}^n$-valued Brownian motions defined as in Section 3. Then, the log-return process reads

$$\boldsymbol{Y}^{L}(t) = \begin{pmatrix} B_1(X_1(t)) + B_1^{\rho}(Z(t)) \\ \dots \\ B_n(X_n(t)) + B_n^{\rho}(Z(t)) \end{pmatrix}, \qquad t \geqslant 0, \tag{A.1}$$

where $X_j(t)$, $j = 1, \dots, n$, and $Z(t)$ are Lévy tempered stable subordinators with parameters α, ξ_j, λ_j and α, ξ_Z, λ_Z, respectively. The subordinators are mutually independent and they are independent of $\boldsymbol{B}(t)$ and $\boldsymbol{B}^{\rho}(t)$.

Using the inverse Gaussian specification, the time-1 laws of the subordinators are equal to those of Eq. (3.13). Therefore, the characteristic exponent of $\boldsymbol{Y}(t)$ becomes

$$\begin{aligned} \psi_{\boldsymbol{Y}^L}(\boldsymbol{u}; t) = t \sum_{j=1}^{n} \left(1 - a\sqrt{\kappa_j}\right) \left(\frac{1}{\sqrt{\kappa_j}} - \sqrt{\frac{1}{\kappa_j} - 2\left(\mathrm{i}u_j\mu_j - \frac{1}{2}u_j^2\sigma_j^2\right)} \right) \\ + ta \left(1 - \sqrt{1 - 2\left(\mathrm{i}\boldsymbol{u}^{\top}\boldsymbol{\mu}^{\rho} - \frac{1}{2}\boldsymbol{u}^{\top}\boldsymbol{\Sigma}^{\rho}\boldsymbol{u}\right)} \right), \qquad \boldsymbol{u} \in \mathbb{R}^n, \end{aligned}$$

with parametrization $\mathcal{P} = \{\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\kappa}, a, \boldsymbol{\rho}\}$ and cardinality $3n + 1 + n(n-1)/2$.

A.2 SSBM model

The Sato-Subordinated Brownian Motion (denoted as $\boldsymbol{Y}^{S}(t)$) of Semeraro (2022) preserves a similar structure of the LS model, so that the log-return dynamics can be still represented in the form (A.1) with the Brownian motions defined in the same way. However, the processes $X_j(t)$, $j = 1, \dots, n$ are the Sato-TS subordinators with $TS(\alpha, \xi_j, \lambda_j)$ time-1 distribution and $Z(t)$ is a Sato-TS subordinators with $TS(\alpha, \xi_Z, \lambda_Z)$ time-1 distribution, and they are mutually independent. The inverse Gaussian specification of the process assigns the same time-1 laws of Eq. (3.13) to the subordinators. Thus, the characteristic exponent becomes

$$\begin{aligned} \psi_{\boldsymbol{Y}^S}(\boldsymbol{u}; t) = \sum_{j=1}^{n} \left(1 - a\sqrt{\kappa_j}\right) \left(\frac{1}{\sqrt{\kappa_j}} - \sqrt{\frac{1}{\kappa_j} - 2t^q\left(\mathrm{i}u_j\mu_j - \frac{1}{2}u_j^2\sigma_j^2\right)} \right) \\ + a \left(1 - \sqrt{1 - 2t^q\left(\mathrm{i}\boldsymbol{u}^{\top}\boldsymbol{\mu}^{\rho} - \frac{1}{2}\boldsymbol{u}^{\top}\boldsymbol{\Sigma}^{\rho}\boldsymbol{u}\right)} \right), \qquad \boldsymbol{u} \in \mathbb{R}^n, \end{aligned}$$

where the set of model parameters is $\mathcal{P} = \{\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\kappa}, a, \boldsymbol{\rho}, q\}$ and the cardinality is $3n + 2 + n(n-1)/2$.

A.3 BB model

The BB model introduced in Ballotta and Bonfiglioli (2016) has a factor-based design given by

$$\boldsymbol{Y}^B(t) = \begin{pmatrix} X_1(t) + b_1 Z(t) \\ \dots \\ X_n(t) + b_n Z(t) \end{pmatrix}, \quad t \geqslant 0, \quad b_j > 0, \ \ j = 1, \dots, n, \tag{A.2}$$

where $X_j(t)$, $j = 1, \dots, n$, and $Z(t)$ are $\mathbb{R}$-valued Lévy processes, and $b_1, \dots, b_n$ are $\mathbb{R}$-valued loadings. In the interest of comparability to the aforementioned models, it is possible to express the 1-dimensional processes in Eq. (A.2) as Brownian motions subordinated to inverse Gaussian processes. The characteristic exponent of $\boldsymbol{Y}^B(t)$ under this specification is then given by

$$\begin{aligned} \psi_{\boldsymbol{Y}^B}(\boldsymbol{u}; t) = t \sum_{j=1}^{n} \left(\frac{1}{\sqrt{\kappa_{X_j}}} - \sqrt{\frac{1}{\kappa_{X_j}} - 2\left(\mathrm{i} u_j \mu_{X_j} - \frac{1}{2} u_j^2 \sigma_{X_j}^2 \right)} \right) \\ + t \left(\frac{1}{\sqrt{\kappa_Z}} - \sqrt{\frac{1}{\kappa_Z} - 2 \sum_{j=1}^{n} b_j \left(\mathrm{i} u_j \mu_Z - \frac{1}{2} u_j^2 \sigma_Z^2 \right)} \right), \qquad \boldsymbol{u} \in \mathbb{R}^n, \end{aligned}$$

with parametrization $\mathcal{P} = \{\boldsymbol{\mu_X}, \boldsymbol{\sigma_X}, \boldsymbol{\kappa_X}, \mu_Z, \sigma_Z, \kappa_Z, \boldsymbol{b}\}$ and cardinality $4n + 3$.

B Proofs

B.1 Proof of Proposition 3.1

Let us define $\phi_j(u) = \mathbb{E}[\exp(\mathrm{i} u B_j(1))]$ and $\phi^\rho(\boldsymbol{u}) = \mathbb{E}[\exp(\mathrm{i}\langle \boldsymbol{u}, \boldsymbol{B}^\rho(1)\rangle)]$, and let $Y_j^I(t) = B_j(X_j(t))$ in Eq. (3.1). Then, by conditioning on $X_j(t)$ we have

$$\phi_{Y_j^I}(u; t) = \mathbb{E}\left[\mathbb{E}\left[\exp\left(\mathrm{i} u Y_j^I(t) \right) \middle| X_j(t) \right] \right]$$

Since $X_j(t)$ is independent of $B_j(t)$,

$$\begin{aligned} \mathbb{E}\left[\exp\left(\mathrm{i} u Y_j^I(t) \right) | X_j(t) = s \right] &= \mathbb{E}\left[\exp\left(\mathrm{i} u B_j(s) \right) | X_j(t) = s \right] \\ &= \mathbb{E}\left[\exp\left(s \psi_j(u) \right) \right], \end{aligned}$$

for $\psi_j(u) = \mathrm{i} u \mu_j - u^2 \sigma_j^2 / 2$, and consequently

$$\begin{aligned} \phi_{Y_j^I}(u; t) &= \mathbb{E}\left[\exp\left(X_j(t) \psi_j(u) \right) \right] \\ &= \phi_{X_j}(-\mathrm{i} t^{q_j} \psi_j(u); 1), \end{aligned} \tag{B.1}$$

where the last equality follows because $X_j(t)$ is a Sato process. Let $\boldsymbol{Y}^C(t) = \boldsymbol{B}^\rho(Z(t))$ in Eq. (3.1). Using similar steps as above, we obtain

$$\phi_{\boldsymbol{Y}^C}(\boldsymbol{u}; t) = \phi_Z(-\mathrm{i} t^{q_Z} \psi^\rho(\boldsymbol{u}); 1),$$

for $\psi^{\rho}(\boldsymbol{u}) = \mathrm{i}\boldsymbol{u}^{\top}\boldsymbol{\mu}^{\rho} - \frac{1}{2}\boldsymbol{u}^{\top}\boldsymbol{\Sigma}^{\rho}\boldsymbol{u}$. By mutual independence of the processes, it follows that

$$\phi_{\boldsymbol{Y}}(\boldsymbol{u};t) = \left(\prod_{j=1}^{n}\phi_{Y_j^I}(u_j;t)\right)\phi_{\boldsymbol{Y}^C}(\boldsymbol{u};t).$$

The thesis follows by substituting the explicit expression of the characteristic function of the Sato subordinator.

B.2 Proof of Proposition 3.2

Expressing the cumulants in terms of the (central) moments,

$$\mathcal{C}_1(Y_j(t)) = \mathbb{E}\left[Y_j(t)\right], \qquad \mathcal{C}_2(Y_j(t)) = \mathbb{E}\left[\left(Y_j(t) - \mathbb{E}\left[Y_j(t)\right]\right)^2\right],$$
$$\mathcal{C}_3(Y_j(t)) = \mathbb{E}\left[\left(Y_j(t) - \mathbb{E}\left[Y_j(t)\right]\right)^3\right], \qquad \mathcal{C}_4(Y_j(t)) = \mathbb{E}\left[\left(Y_j(t) - \mathbb{E}\left[Y_j(t)\right]\right)^4\right] - 3\,\mathcal{C}_2(Y_j(t))^2,$$

and making standard computations we get result (i), where the explicit expression in terms of the time scalers follows by applying the relations $\mathcal{C}_m(X_j(t)) = t^{m\cdot q_j}\mathcal{C}_m(X_j)$, and $\mathcal{C}_m(Z(t)) = t^{m\cdot q_Z}\mathcal{C}_m(Z)$. Part (ii) is proved following analogous step as in part (i).

B.3 Proof of Proposition 3.3

Part (i) holds by expressing the skewness and excess kurtosis in terms of the cumulants given in Eqs. (3.4)–(3.7) and checking the long-run dominant terms according to the higher-order infinite terms t^{q_j} and t^{q_Z}, and the relative values of q_j and q_Z. Part (ii) is proved in a similar way.

B.4 Proof of Proposition 4.2

Since

$$Y_{f_k,f_l}(t) = Y_{f_k,d}(t) - Y_{f_l,d}(t) = (1,-1)\boldsymbol{Y}(t), \quad t \geqslant 0, \tag{B.2}$$

and additive processes are invariant under linear transformations (see Proposition 2.10 in Papapantoleon, 2006), $Y_{f_k,f_l}(t)$ is an additive process under $\mathbb{Q}^d$.

Additivity is also preserved under Esscher changes of measure (see Theorem 2.26 in Papapantoleon, 2006). Therefore, if $Y_{f_k,f_l}(t)$ is additive under $\mathbb{Q}^d$, it is also an additive process under $\mathbb{Q}^{f_l}$, which is an Esscher transformation of $\mathbb{Q}^d$ with parameter $\bar{\boldsymbol{h}} = (0,1)^{\top}$.

As (B.2) implies that $\psi_{Y_{f_k,f_l}}(u;t) = \psi_{\boldsymbol{Y}}(\bar{\boldsymbol{u}};t)$, for $\bar{\boldsymbol{u}} = (u,-u)^{\top}$, $u \in \mathbb{R}$, it is easy to show that the characteristic exponent of $\boldsymbol{Y}(t)$ under $\mathbb{Q}^{f_l}$ is

$$\psi^{f_l}_{Y_{f_k,f_l}}(u;t) = \psi^{d}_{\boldsymbol{Y}}(-\mathrm{i}\bar{\boldsymbol{h}} + \bar{\boldsymbol{u}};t) - \psi^{d}_{\boldsymbol{Y}}(-\mathrm{i}\bar{\boldsymbol{h}};t).$$

The assert is proved.