

GP regression for arbitrage-free construction of the swaption cube.

Areski Cousin, Adrien Deleplace, Adrien Misko
Nexialog Consulting, Paris, France

September 5, 2022



Abstract

In this paper, we introduce a 3D finite dimensional Gaussian process (GP) regression approach for learning arbitrage-free swaption cubes. Based on possibly noisy observations of swaption prices, the proposed ‘constrained’ GP regression approach is proven arbitrage-free along the strike direction (butterfly and call-spread arbitrages are precluded on the entire 3D input domain). From [Johnson and Nonas \(2009\)](#), the cube is free from static arbitrage along the tenor and maturity directions if swaption prices satisfies an infinite set of in-plane triangular inequalities. We empirically demonstrate that, considering a finite-dimensional weaker form of this condition, is enough for the GP to generate swaption cubes with a negligible proportion of violation points, even for a small training set. Besides, we compare the performance of the GP approach with the SABR model, applied to a data set of payer and receiver OTM swaptions. The constrained GP approach provides better prediction results compared to the SABR approach. In addition, we show that SABR Calibration is better when using the GP cube output as new observations (in term of predictive error and absence of arbitrage). Finally, the GP approach is able to quantify in- and out-of-sample uncertainty through Hamiltonian Monte Carlo simulations, allowing AVA model calculation.

Keywords— Swaption cube, Constrained Kriging, No arbitrage, Option pricing, Hamiltonian Monte Carlo

Contents

| | | |
|----------|--|-----------|
| 1 | Absence of static arbitrages for the swaption cube | 4 |
| 2 | GP Regression to construct an arbitrage free swaption cube | 4 |
| 2.1 | Classical 3-dimensional Gaussian regression | 5 |
| 2.2 | Imposing no arbitrage-free constraints to GP regression | 5 |
| 2.3 | Maximum Likelihood Estimation | 7 |
| 2.4 | The most probable response cube and measurement noises | 7 |
| 2.5 | Sampling | 8 |
| 3 | Empirical Results | 8 |
| 3.1 | Data | 8 |
| 3.1.1 | Absence of arbitrages in the data set | 9 |
| 3.2 | SABR model benchmark | 9 |
| 3.2.1 | SABR model | 9 |
| 3.2.2 | Model calibration and pricing | 10 |
| 3.2.3 | Unobserved parameter estimation | 10 |
| 3.3 | Absence of Arbitrage, Calibration and Prediction performance | 11 |
| 3.3.1 | Implementation of the GP regression approach | 11 |
| 3.3.2 | Calibration and testing methodology | 11 |
| 3.3.3 | Absence of arbitrage in the predicted swaption cube | 12 |
| 3.3.4 | Calibration performance | 15 |
| 3.3.5 | Prediction performance | 15 |
| 3.3.6 | SABR model calibrated on the GP predicted cube | 16 |
| 3.4 | Uncertainty quantification | 17 |
| 4 | Conclusion | 18 |
| 5 | Further research | 19 |
| A | Calibration of the SABR model | 19 |

Introduction

The swaption cube is an important financial object as it is used as an input of many interest-rate derivative pricing engines. It consists in a swaption implied volatility surface, indexed by maturity, strike and tenor. The swaption market is one of the biggest in the interest rate derivatives market. Swaptions are usually used to hedge, and while at-the-money swaptions are observed to be very liquid, away-from-the money swaption prices are often missing or not reliable. The major issue with the swaption cube is that, along its three axes, missing quotes must be interpolated in an arbitrage-free way in order to be used as input of some pricing or hedging models.

For equity options, the volatility surface, indexed by the maturity and the strike of the derivative, is a topic that has been thoroughly studied. Number of machine learning methods for learning pricing functions have emerged (see, e.g., [Ruf and Wang \(2019\)](#) for an in-depth literature review). For these volatility surfaces, absence of arbitrage is guaranteed if there is no Call spread, Butterfly spread and Calendar spread arbitrage opportunities. A Gaussian process regression methodology, constrained on the no arbitrage conditions for equity options, has been introduced in [Chataigner et al. \(2020\)](#). Using this approach, they achieved to produce arbitrage-free interpolated prices, and local volatility surfaces, with their uncertainties. In the literature, several ways of interpolating the missing prices or volatilities of the swaption cube have been introduced. A distinction needs to be made between functional interpolation, which is for example polynomial interpolation, and stochastic interpolation, where stochastic models such as SABR and CEV models are used. The most common method for interpolating in an arbitrage-free way between quoted strikes, as well as extrapolating beyond, is to use the SABR model. However, this construction method requires several observed quotes at different strikes for each considered pair of tenors and maturities. When there is only one quoted strike observation, for a given tenor and maturity, three methods are usually employed in the industry to interpolate the swaption smile. Two of them are based on the cap market. The first approach consists in using caplet SABR parameters as a basis for the calibration of the swaption SABR parameters (see, e.g., [Hagan and Konikov \(2004\)](#)). The second, discussed in [Deloitte \(2019\)](#), consists in mapping underlying caplet volatilities to the swaption volatility by making an assumption in relating the moneyness of the swaption to the moneyness of the underlying caplets. The validity of both approaches decreases when the the tenor of the considered swaptions increases. When cap market data is not available, parameters of the SABR model are taken as the closest SABR parameters that are calibrated with quotes. In order to interpolate w.r.t maturity and tenor, it is possible to use functional interpolation directly on swaptions prices, or in a more elaborate manner, to use functional interpolation on the parameters of stochastic models. As explained in [Johnson and Nonas \(2009\)](#), in addition to the no Call spread and Butterfly spread arbitrage, the interpolated swaption prices must also respect an additional arbitrage condition, named "in-plane triangular".

Inspired by the local volatility construction proposed by [Chataigner et al. \(2020\)](#), the main goal of this paper is to present and study a 3D constrained GP construction method for interpolating arbitrage-free swaption price cubes at a given quotation date. Compared to the equity local volatility construction method proposed in [Chataigner et al. \(2020\)](#), several improvements have been added :

- The constrained GP regression approach has been extended to an input space of dimension 3. Numerical implementation of the 3D 'constrained' kriging problem is much more involved than in the 2D case considered in [Chataigner et al. \(2020\)](#). We show that, considering an anisotropic stationary kernel, allows to benefit from tensorisation of the GP covariance matrix, and significantly reduces numerical complexity.
- Compared to the equity price surface, the no-arbitrage constraints on swaption prices are no more directional. Indeed, the in-plane triangular inequality involves prices of swaptions with different maturities and tenors. This constraint is by far more complex to handle than the no Calendar spreads arbitrage constraint in the equity option case. It can be viewed as a weaker form of the no Calendar spreads arbitrage condition, and requires an infinite number of checks in order to be imposed everywhere on the domain. In our methodology, we consider a weaker version of the in-plane inequality constraint, and we empirically check in Section 3 that the weaker form is enough to ensure no-arbitrage almost everywhere.

- Computation of the GP prior hyperparameters by MLE is much longer in the swaption case than in the equity case (see Subsection 2.3). Search for the maximal likelihood has been considerably improved by explicitly computing the Jacobian of the log-likelihood function.

The paper is organized as follows. In Section 1, we present the conditions on European swaption prices that preclude from any static arbitrage opportunity, at a particular quotation date. These arbitrage-free conditions on the swaption cube translate into a convexity and a monotonicity constraints in strike, and into an "in-plane triangular" constraint in maturity and tenor. In Section 2, we detail the constrained GP regression methodology used to construct swaption price cubes that do not lead to arbitrage opportunities. Finally, in Section 3, we apply the presented GP methodology to a data set of payer and receiver OTM swaptions, and compare the result to a SABR construction approach, which is the standard market practice to parameterize rate volatilities.

1 Absence of static arbitrages for the swaption cube

Before establishing the arbitrage free constraints for the swaption cube, we set some notation.

Let $S_w(T_s, T_e, K)$ be the price, at a given quotation date, of a European payer (receiver) swaption, with a payer (receiver) swap $S(T_s, T_e, K)$ as underlying asset, where T_s and T_e are respectively the first and last payment date of the swap with fixed interest rate K . For equity options, in order to construct an arbitrage-free price or volatility surface, we must verify that the price of Call spreads, Butterfly spreads and Calendar spreads are positive. In the strike direction, the Butterfly and Call-spread arbitrage conditions also apply to the swaption cube. However the Calendar spread condition does not hold here, as explained in [Johnson and Nonas \(2009\)](#). A non-trivial no arbitrage constraint between European swaptions with the same strike and different option expiries was introduced in this article. Since this condition translates into an inequality between the price of three swaptions with same strikes, it is called "in-plane triangular condition".

Proposition 1 *The European payer (receiver) swaption price cube $(T, t, K) \rightarrow S_w(T, T + t, K)$ is free of static arbitrage if and only if the three following conditions hold:*

- (i) For all $T, t, K \geq 0$, $S_w(T, T + t, K) \geq 0$
- (ii) $K \rightarrow S_w(T, T + t, K)$ is a convex, decreasing (increasing) function
- (iii) The In-plane triangular condition holds:
For all $T, t, h, K \geq 0$, $S_w(T, T + t, K) + S_w(T + t, T + t + h, K) \geq S_w(T, T + t + h, K)$

PROOF: (i) comes from the need of the prices of Call spreads and Butterfly spreads to be positive. A proof of (ii) is given in [Johnson and Nonas \(2009\)](#).

2 GP Regression to construct an arbitrage free swaption cube

In this section, we present a GP based methodology to build an arbitrage-free European payer (receiver) swaption cube.

For an input domain \mathcal{D} of maturities, tenors and strikes, we build at quotation date t_0 , a payer (receiver) swaption price cube

$$S_w : \mathcal{D} \rightarrow \mathbb{R}^+$$

$$(T, t, K) \mapsto S_w(T, T + t, K),$$

satisfying the no arbitrage conditions given in Proposition 1, from n noisy observations $\mathbf{y} = [y_1, \dots, y_n]^\top$ of the function S_w at input points $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]$ where $\mathbf{x}_i = (T_i, t_i, K_i)$, $i = 1, \dots, n$. We assume that the price function S_w is represented as a Gaussian process. The market fit condition is then written as

$$\mathbf{y} = S_w(\mathbf{X}) + \varepsilon \quad (1)$$

where $S_w(\mathbf{X}) := [S_w(\mathbf{x}_1), \dots, S_w(\mathbf{x}_n)]^\top$ is the vector of the values of the GP at the input points \mathbf{X} . The additive noise term $\varepsilon = [\varepsilon_1, \dots, \varepsilon_n]^\top$ is assumed to be a zero-mean Gaussian vector, independent from S_w , and with a homoscedastic covariance matrix given as $\zeta^2 I_n$, where I_n is the identity matrix of size n .

2.1 Classical 3-dimensional Gaussian regression

We consider a zero-mean Gaussian process prior on the mapping $S_w = S_w(\mathbf{x})_{\mathbf{x} \in \mathcal{D}}$ with covariance function c .

Then the output vector $S_w(\mathbf{X})$ has a normal distribution with zero mean and covariance matrix \mathbf{C} with components $\text{cov}(S_w(T_i, T_i + t_i, K_i), S_w(T_j, T_j + t_j, K_j)) = c((T_i, t_i, K_i), (T_j, t_j, K_j)) = c(\mathbf{x}_i), c(\mathbf{x}_j)$. We consider a 3-dimensional anisotropic, stationary covariance kernel given as, for \mathbf{x} and $\mathbf{x}' \in \mathcal{D}$,

$$c(\mathbf{x}, \mathbf{x}') = \sigma^2 R_T(T - T', \theta_T) R_t(t - t', \theta_t) R_K(K - K', \theta_K)$$

where $\theta_T, \theta_t, \theta_K$ are the length scale parameters associated with directions T, t, K , and σ^2 is the marginal variance of the GP prior. The functions R_T, R_t and R_K are kernel correlation functions. As explained in [Rasmussen and Williams \(2006\)](#), $S_w \mid \mathbf{y} = S_w(\mathbf{X}) + \varepsilon$ is again a GP with mean function η and covariance function c^* given respectively by :

$$\eta(\mathbf{x}) = \mathbf{c}(\mathbf{x})^\top (\mathbf{C} + \zeta^2 I_n)^{-1} \mathbf{y}, \quad \mathbf{x} \in \mathcal{D} \quad (2)$$

and

$$c^*(\mathbf{x}, \mathbf{x}') = c(\mathbf{x}, \mathbf{x}') - \mathbf{c}(\mathbf{x})^\top (\mathbf{C} + \zeta^2 I_n)^{-1} \mathbf{c}(\mathbf{x}'), \quad \mathbf{x}, \mathbf{x}' \in \mathcal{D} \quad (3)$$

where $\mathbf{c}(\mathbf{x}) = [c(\mathbf{x}, \mathbf{x}_1), \dots, c(\mathbf{x}, \mathbf{x}_n)]^\top$.

Prediction and uncertainty quantification are made using the conditional distribution $S_w \mid \mathbf{y} = S_w(\mathbf{X}) + \varepsilon$. The best linear unbiased estimator of S_w is given as the conditional mean function (2). The conditional covariance function c^* can be used to obtain confidence bounds around the predicted price surface. The hyper-parameters of the kernel function c , as well as the variance of the noise, can be estimated using a maximum likelihood approach or a cross-validation approach (see [Rasmussen and Williams \(2006\)](#)). The model described in this sub-section corresponds to unconstrained Gaussian process regression, which therefore does not take into account the arbitrage-free conditions as described in Proposition 1.

2.2 Imposing no arbitrage-free constraints to GP regression

Conditional on the inequality constraints given in Proposition 1, the process S_w is no longer Gaussian, this leads to two difficulties. The first is that we depart from the classical framework of GP regression, where the posterior distribution remains Gaussian. The second is that testing the inequality constraints on the entire input domain would require an infinite number of checks. We will use the solution proposed by [Cousin et al. \(2016\)](#) which consists in constructing a finite-dimensional approximation S_w^N of the Gaussian prior S_w , designed in such a way that the constraints can be imposed for the entire domain \mathcal{D} with a finite number of checks. We then recover the constrained posterior distribution by sampling a truncated Gaussian vector. Let us describe the methodology with more details.

We first rescale the input domain \mathcal{D} to $[0, 1]^3$. Without loss of generality, we then consider a discretized version of this rescaled domain as a 3-dimensional regular grid with nodes (T_i, t_j, K_k) , $i = 1, \dots, N_T$, $j = 1, \dots, N_t$, $k = 1, \dots, N_K$, where N_T, N_t, N_K are, respectively, the number of maturities, tenors and strikes we chose for our grid, and $(\delta_T, \delta_t, \delta_K)$ the corresponding constant steps¹. Then, $N = N_T \times N_t \times N_K$

¹This choice has been made for the sake of clarity. It is also possible to consider a grid with heterogeneous steps.

is the total number of nodes of the 3-dimensional grid. Now, each node of the grid is associated with a hat basis function defined as:

$$\phi_{i,j,k}(T, t, K) := \max\left(1 - \frac{|T - T_i|}{\delta_T}, 0\right) \max\left(1 - \frac{|t - t_j|}{\delta_t}, 0\right) \max\left(1 - \frac{|K - K_k|}{\delta_K}, 0\right).$$

Then, the GP prior S_w is approximated on \mathcal{D} by the following finite-dimensional process:

$$S_w^N(T, T + t, K) = \sum_{i=1}^{N_T} \sum_{j=1}^{N_t} \sum_{k=1}^{N_K} S_w(T_i, T_i + t_j, K_k) \phi_{i,j,k}(T, t, K), \quad \text{for all } (T, t, K) \in \mathcal{D} \quad (4)$$

Note that the process S_w^N corresponds to a piecewise linear interpolation of S_w at nodes (T_i, t_j, K_k) , for $i = 1, \dots, N_T$, $j = 1, \dots, N_t$, and $k = 1, \dots, N_K$.

If we denote $S_w(T_i, T_i + t_j, K_k)$ as $\xi_{i,j,k}$, then the vector $\boldsymbol{\xi} = [\xi_{1,1,1}, \dots, \xi_{i,j,k}, \dots, \xi_{N_T, N_t, N_K}]^\top$ is a zero-mean Gaussian vector with $N \times N$ covariance matrix Γ^N with components $\Gamma_{\mathbf{x}, \mathbf{x}'}^N = c(\mathbf{x}, \mathbf{x}')$ for \mathbf{x} and \mathbf{x}' two nodes of the grid. Let us define $\boldsymbol{\phi}(T, t, K)$ as the N dimensional vector given as:

$$\boldsymbol{\phi}(T, t, K) = [\phi_{1,1,1}(T, t, K), \dots, \phi_{i,j,k}(T, t, K), \dots, \phi_{N_T, N_t, N_K}(T, t, K)].$$

Then the finite-dimensional GP prior defined in 4 can be restated in matrix form: The equality (4) can then be written as:

$$S_w^N(T, T + t, K) = \boldsymbol{\phi}(T, t, K) \cdot \boldsymbol{\xi}$$

Then if we define $\boldsymbol{\Phi}(\mathbf{X})$ as the $n \times N$ matrix of basis functions such that each row i corresponds to the vector $\boldsymbol{\phi}(\mathbf{x}_i)$, we have that $S_w^N(\mathbf{X}) = \boldsymbol{\Phi}(\mathbf{X}) \cdot \boldsymbol{\xi}$, with $S_w^N(\mathbf{X}) := [S_w^N(\mathbf{x}_1), \dots, S_w^N(\mathbf{x}_n)]^\top$.

From now on, we will refer to $\boldsymbol{\Phi}(\mathbf{X})$ as $\boldsymbol{\Phi}$.

Proposition 2 *The following statements hold for European Payer swaptions.*

- (i) *The finite-dimensional process S_w^N uniformly converges to S_w on \mathcal{D} as $N_T \rightarrow \infty$, $N_t \rightarrow \infty$ and $N_K \rightarrow \infty$, almost surely,*
- (ii) *S_w^N is a decreasing function of K on \mathcal{D} if and only if $\xi_{i,j,k} \geq \xi_{i,j,k+1}$*
- (iii) *S_w^N is a convex function of K on \mathcal{D} if and only if $\xi_{i,j,k+2} - \xi_{i,j,k+1} \geq \xi_{i,j,k+1} - \xi_{i,j,k}$.*

PROOF: These properties can be proved using the same methodology used in [Maatouk and Bay \(2017\)](#) for proving monotonicity and convexity. Note that these inequality constraints are linear in the components of $\boldsymbol{\xi}$.

Remark 1 *For an European receiver swaption, the above properties are the same for (i) and (iii). However (ii) is not the same. We have that: (ii) S_w^N is an increasing function of K if and only if $\xi_{i,j,k} \leq \xi_{i,j,k+1}$*

However, it is not possible to impose the in-plane triangular condition everywhere on the domain \mathcal{D} with a finite number of checks. In empirical Section 3, we consider a weaker version of this constraint by only imposing its validity for time steps of size 1 year.

This constraint is set using a "fictitious grid" of maturities $T^* = [T_1^*, T_2^*, \dots, T_{N_T^*}^*]$, where $T_1^* = T_1$, $T_{N_T^*}^* = T_{N_T}$ and the step $\delta_T^* = \frac{1}{T_1 - T_{N_T}}$. Therefore it is a rescaled 1 year step.

Formally, we consider the following weaker version of the in-plane triangular condition:

$$S_w^N(T_l^*, T_{l+1}^*, K) + S_w^N(T_{l+1}^*, T_{l+2}^*, K) - S_w^N(T_l^*, T_{l+2}^*, K) \geq 0, \quad l = 1, \dots, N_T^* \quad (5)$$

This is equivalent to:

$$\boldsymbol{\phi}(T_l^*, T_{l+1}^*, K) \cdot \boldsymbol{\xi} + \boldsymbol{\phi}(T_{l+1}^*, T_{l+2}^*, K) \cdot \boldsymbol{\xi} - \boldsymbol{\phi}(T_l^*, T_{l+2}^*, K) \cdot \boldsymbol{\xi} \geq 0 \quad (6)$$

Then (6) are linear in $\boldsymbol{\xi}$ and can be added to the ones defined constraints to the ones defined in Proposition 2.

Given property (i) of Proposition 2, if we define \mathcal{M} as the set of 3-dimensional continuous functions which are increasing and monotonic w.r.t. K , and respecting the in-plane triangular inequality (for steps of size 1 year, defined as (5)), then our construction problem consists in finding the conditional distribution of S_w^N such that

$$\begin{cases} \mathbf{y} = S_w^N(\mathbf{X}) + \varepsilon \\ S_w^N \in \mathcal{M}. \end{cases}$$

Given properties (ii) and (iii) of Proposition 2, our process S_w^N satisfies the no-arbitrage conditions in the strike direction on the entire domain \mathcal{D} when these constraints are satisfied at the nodes. The problem above can then be restated as follows

$$\begin{cases} \mathbf{y} = \Phi \cdot \xi + \varepsilon \\ \xi \in \mathcal{C}_{ineq}. \end{cases}$$

Indeed, $S_w^N \in \mathcal{M} \Leftrightarrow \xi \in \mathcal{C}_{ineq}$, where \mathcal{C}_{ineq} is a set of linear inequality constraints as defined in (ii), (iii) of Proposition 2 and the relaxed in-plane triangular inequality constraint (6).

2.3 Maximum Likelihood Estimation

The parameters of the kernel function c and the noise parameter ς can be given as an input or estimated. Let $\Theta = [\theta_T, \theta_t, \theta_K, \sigma, \varsigma^2]^\top$. In our approach, we estimate Θ as the hyperparameter vector that maximizes the marginal log likelihood $\mathcal{L}(\Theta)$ for the process S_w^N .

The unconstrained Gaussian Likelihood is written as $\mathbb{P}(\mathbf{y} = \Phi \cdot \xi + \varepsilon | \Theta)$. Under the finite dimensional approximation, the marginal log likelihood can be expressed as:

$$\mathcal{L}(\Theta) = -\frac{1}{2} \mathbf{y}^\top (\Phi \Gamma^N \Phi^\top + \varsigma^2 I_n)^{-1} \mathbf{y} - \frac{1}{2} \log(\det(\Phi \Gamma^N \Phi^\top + \varsigma^2 I_n)) - \frac{n}{2} \log(2\pi). \quad (7)$$

This marginal log likelihood also has a closed form derivative w.r.t each parameter Θ_i (see [Rasmussen and Williams \(2006\)](#)), given as:

$$\frac{\partial \mathcal{L}(\Theta)}{\partial \Theta_i} = \frac{1}{2} \mathbf{y}^\top A^{-1} \frac{\partial A}{\partial \Theta_i} \mathbf{y} - \frac{1}{2} \text{tr}(A^{-1} \frac{\partial A}{\partial \Theta_i}). \quad (8)$$

Here $A = (\Phi \Gamma^N \Phi^\top + \varsigma^2 I_n)$ and $\frac{\partial A}{\partial \Theta_i}$ the element-wise derivative of A w.r.t Θ_i .

This closed form can be used in order to shorten the computation time of the MLE in practice. This is useful because the computation of the MLE is longer for the swaption cube than for the equity surface, since the product $\Phi \Gamma^N \Phi^\top$ depends on the dimension of the problem.

Remark 2 *The reason for the choice of maximizing the unconstrained likelihood instead of maximizing the constrained likelihood $\mathbb{P}(\mathbf{y} = \Phi \cdot \xi + \varepsilon | \xi \in \mathcal{C}_{ineq}, \Theta)$ is discussed in [Bachoc et al. \(2019\)](#). In this paper, it is explained that constraining the GP increases the computational burden of the maximization, with a negligible impact on the resulting MLE.*

2.4 The most probable response cube and measurement noises

In the constrained GP regression case, we consider the mode of the truncated Gaussian process as an estimator for the cube instead of the mean. For the unconstrained GP regression case, mean and mode coincide because of the Gaussian profile of the conditional distribution, but the mode is easier to compute.

The Maximum a posteriori (MAP) of S_w^N given the constraints satisfies the constraints on the entire domain. In the sense of Bayesian statistics, it coincides with the mode of the truncated Gaussian process. Its expression is given in [Cousin et al. \(2016\)](#),

$$\mathbf{M}_{S_w^N}(T, t, K) := \sum_{i=1}^{N_T} \sum_{j=1}^{N_t} \sum_{k=1}^{N_K} \hat{\xi}^{i,j,k} \phi_{i,j,k}(T, t, K) \quad (9)$$

Where $\hat{\xi} = \left(\hat{\xi}^{(1,1,1)}, \dots, \hat{\xi}^{(i,j,k)}, \dots, \hat{\xi}^{(N_T, N_t, N_K)} \right)^\top$ is the MAP of the Gaussian coefficients ξ , which satisfies the inequality constraints. In order to locate the largest noise terms, we proceed as in [Chataigner et al. \(2020\)](#). We compute the joint MAP $(\hat{\xi}, \hat{\varepsilon})$ of the truncated Gaussian vector (ξ, ε) . This joint MAP $(\hat{\xi}, \hat{\varepsilon})$ is solution of:

$$\operatorname{argmax}_{\vartheta, e} \mathbb{P}(\xi \in [\vartheta, \vartheta + d\vartheta], \varepsilon \in [e, e + de] \mid \Phi \cdot \xi + \varepsilon = \mathbf{y}, \xi \in \mathcal{C}_{ineq}).$$

The fact that (ξ, ε) is Gaussian with mean 0 with covariance matrix the block-diagonal matrix

$$\begin{pmatrix} \Gamma^N & 0 \\ 0 & \varsigma^2 I_n \end{pmatrix}$$

implies that $(\hat{\xi}, \hat{\varepsilon})$ is a solution to the quadratic problem:

$$\operatorname{argmin}_{\Phi \cdot \vartheta + e = \mathbf{y}, \vartheta \in \mathcal{C}_{ineq}} \left(\vartheta^\top (\Gamma^N)^{-1} \vartheta + e^\top (\varsigma^2 I_n)^{-1} e \right) \quad (10)$$

The most probable response cube is then $M_{S_w^N}(\mathbf{x}) = \Phi(\mathbf{x}) \cdot \hat{\xi}$, for $\mathbf{x} \in \mathcal{D}$, and the most probable measurement noise is $\hat{\varepsilon}$.

2.5 Sampling

Contrary to constrained spline interpolation approaches, the GP regression is a truly probabilistic approach. Sampling of the posterior distribution can be used for uncertainty quantification. Sampling paths of the swaption price cube consists in sampling ξ truncated on \mathcal{C}_{ineq} , such that $\Phi \cdot \xi + \varepsilon = \mathbf{y}$. The conditional distribution of $\xi \mid \mathbf{y} = \Phi \cdot \xi + \varepsilon$ is multivariate normal with mean vector μ_{cond} and covariance matrix Σ_{cond} defined as follows:

$$\mu_{\text{cond}} = \Gamma^N \Phi^\top (\Phi \Gamma^N \Phi^\top + \varsigma^2 I_n)^{-1} \mathbf{y} \quad (11)$$

and

$$\Sigma_{\text{cond}} = \Gamma^N \Phi^\top (\Phi \Gamma^N \Phi^\top + \varsigma^2 I_n)^{-1} \Phi \Gamma^N. \quad (12)$$

In practice we will sample this truncated distribution by using Hamiltonian Monte Carlo, as presented in [Pakman and Paninski \(2012\)](#). The starting vector we must choose has to satisfy the constraints. Thus $\hat{\xi}$, the MAP of ξ defined in Section 2.4, is an ideal candidate.

3 Empirical Results

This section compares the performance of the "constrained" GP regression with a calibrated SABR approach.

3.1 Data

For our empirical study, our data set consists of, at a quotation date t_0 , 112 OTM payer swaptions, 28 ATM swaptions and 112 OTM receiver swaptions with maturities $T = [5, 10, 15, 20]$ and tenors $t = [1, 2, 5, 10, 15, 20, 30]$. There is then $N = 252$ input prices. For OTM payer swaptions, quotes are provided for the relative strikes $K_{relPayer} = [0.0025, 0.005, 0.01, 0.02]$. For the OTM receiver swaptions, we have $K_{relReceiver} = [-0.02, -0.01, -0.005, -0.0025]$. At the money, the prices of payer and receiver swaptions coincide. Each option of our data set is listed with market price, market implied normal volatility, forward swap rate, absolute strike and annuity. The annuities are computed from a yield curve, built from the observed swap rates.

3.1.1 Absence of arbitrages in the data set

The first step of our analysis is to check if the data set itself is exempted from arbitrages, as defined in Proposition 1. The in-plane triangular condition compares swaptions with the same absolute strike and different options expiries. Since the options in our data set are quoted for relative strikes, we do not have multiple observations for the same absolute strike, then we lack the data to test the in-plane triangular constraint. We can however for each pair (Maturity, Tenor) := (T, t) test the monotonicity and convexity constraint in the strike direction for both payer and receiver swaptions. In this data set, there are 28 pairs of maturities and tenors.

| Pairs (T,t) respecting the AOA constraints | Monotonicity constraint | Convexity constraint |
|--|-------------------------|----------------------|
| Input data | in K | in K |
| Payer | 100% | 67.86% |
| Receiver | 100% | 100% |

Table 1 Percent of pairs (T,t) respecting the absence of arbitrage in strike constraints defined in Proposition 1

As stated in Table 1, there is no violations of the monotonicity constraint, therefore the data set is exempted from Call spreads arbitrage. The data set is not exempted from Butterfly spreads arbitrage because, for the payer swaptions, only 67.86% of the pairs (T, t) respect the convexity constraint. Nevertheless, arbitrable observations are not removed from the data set.

3.2 SABR model benchmark

In all our study, we will use the SABR model as a benchmark regarding the prediction performance for our Gaussian process model. The SABR model has become a market standard for interpolation of swaption volatilities in the strike direction.

In this subsection, we briefly explain the calibration of the SABR model and how we estimate the parameters of the model.

3.2.1 SABR model

SABR is a dynamic model in which the forward process F and the volatility process $\hat{\sigma}$ satisfy the following system of SDE:

$$\begin{cases} d\hat{F}(t) = \hat{\sigma}(t)\hat{F}(t)^\beta dW(t), & \hat{F}(0) = f \\ d\hat{\sigma}(t) = \alpha\hat{\sigma}(t)dZ(t), & \hat{\sigma}(0) = \sigma \\ dW(t)dZ(t) = \rho dt \end{cases}$$

Where W, Z are Wiener processes with correlation $\rho \in [-1, 1]$, $\alpha \geq 0$, $0 \leq \beta \leq 1$ and $\sigma > 0$. In our study we fix $\beta = 1$. In order to construct the swaption price cube via the SABR model, we consider as many SABR models as underlying swap, i.e., pair (T, t) .

For a pair (T, t) , the calibration procedure is explained in Section 3.3.4. To accommodate with negative rates, we use the so-called Shifted SABR model defined as follows:

$$\begin{cases} d\hat{F}(t) = \hat{\sigma}(t)(\hat{F}(t) + s)^\beta dW(t), & \hat{F}(0) = f \\ d\hat{\sigma}(t) = \alpha\hat{\sigma}(t)dZ(t), & \hat{\sigma}(0) = \sigma \\ dW(t)dZ(t) = \rho dt \end{cases}$$

Where s is the shift parameter.

3.2.2 Model calibration and pricing

As it is explained in [Hagan and Lesniewski \(2002\)](#), an approximated analytical solution for the implied volatility exists. The volatility of a swaption with maturity T_{ex} is given as the following asymptotic approximation, where we consider $\beta = 1$:

$$\sigma_N(K) = \sigma \frac{f - K}{\log \frac{f}{K}} \cdot \left(\frac{z}{x(z)} \right) \cdot \left[1 + \left(-\frac{1}{24}\sigma^2 + \frac{1}{4}\rho\sigma\alpha + \frac{1}{24}(2 - 3\rho^2)\alpha^2 \right) T_{ex} \right]$$

Where,

$$z = \frac{\alpha}{\sigma} \log \frac{f}{K}, \quad x(z) = \log \left(\frac{\sqrt{1 - 2\rho z + z^2} - \rho + z}{1 - \rho} \right)$$

This model is free of arbitrage except for very low strikes and large maturities T_{ex} , for which the smile $\sigma_N(K)$ leads to negative probabilities, and therefore arbitrages. This behaviour of the smile for such strikes and maturities has been empirically shown in [Hagan et al. \(2014\)](#). The normal volatilities of the shifted SABR model are obtained by replacing in the formula above f by $f + s$ and K by $K + s$, where s is the shift parameter, chosen according to the data set. In our study we consider $s = 0.025$.

In this approach, a set of parameters $(\beta, \sigma, \alpha, \rho)$ is mapped to each distinct pair (T, t) observed. Once we fixed $\beta = 1$, the other parameters (σ, α, ρ) are calibrated by minimizing the following across the m strikes observed for the observed pair (T, t) :

$$\sum_{k=1}^m [\sigma_N(K_k; \sigma, \alpha, \rho) - \sigma^{MARKET}(K_k)]^2, \quad m > 0 \quad (13)$$

Then, once the parameters are calibrated for each observed pair (T, t) , the price of European payer and receiver swaptions $S_w(T, t, K)$ is obtained using the normal SABR volatility $\sigma_N(K)$ as input in Bachelier's pricing formula.

3.2.3 Unobserved parameter estimation

The previous section explained how to calibrate the parameters of the shifted SABR model for each pair (T, t) , where we have at least one observation of swaptions prices. In order to reconstruct the cube, a set of parameters $(\beta, \sigma, \alpha, \rho)$ needs to be mapped to each unobserved couple (T, t) . As previously, β is set to 1. Then for each unobserved pairs (T, t) , each parameter of the set (σ, α, ρ) is chosen as the weighted average of the corresponding parameter of the 4 nearest calibrated pair (T, t) . The weight is the inverse of the distance between the unobserved pair and the four nearest calibrated pairs. This method of interpolation and extrapolation has been chosen because it offered better results than a cubic spline approach, especially when extrapolating in-sample. From now on, this method will be referred to as "barycentric mean" approach. One problem of such method of interpolation and extrapolation is that the SABR model parameters $(\beta, \sigma, \alpha, \rho)$, for each unobserved pair (T, t) , are estimated independently in contrary to the classical calibration. For an unobserved pair (T, t) , there is then no guarantee that the resulting smile will be free of arbitrage. The absence of arbitrage of the swaption cube obtained using this method is investigated in Subsection 3.3.3.

3.3 Absence of Arbitrage, Calibration and Prediction performance

In this section, we compare the violation of AOA, calibration and prediction performance for the cubes constructed from the constrained GP approach and the SABR benchmark approach introduced in Section 3.2.

3.3.1 Implementation of the GP regression approach

Since OTM payer and receiver swaptions are different contracts, with different domains of strike (absolute and relative), they need to be considered separately. For the OTM payer swaptions, the limits of the 3-dimensional grid are defined by \mathcal{D}_{Payer} where $T : [5, 20]$, $t : [1, 30]$ and $K : K_{Payer}$, and the limits of the 3-dimensional grid for OTM receiver swaptions are defined by $\mathcal{D}_{Receiver}$ where $T : [5, 20]$, $t : [1, 30]$ and $K : K_{Receiver}$. For both payer and receiver swaptions, we consider a GP method with matern 5/2 kernels² over a $[0, 1]^3$ rescaled domain. The parameters of the kernels $[\theta_T, \theta_t, \theta_K, \sigma]$ and the homoscedastic noise variance are fitted through MLE as described in Section 2.3. Each different set of observed data will lead to a different metamodel. The MLE is maximized with the function 'minimize', with the method 'L-BFGS-B', from the Python library 'Scipy.optimize'. With the purpose of not ending up with a local minimum, the algorithm is restarted 20 time with random starting points. In order to speed up the computation time the Jacobian of the log likelihood is given as input of the optimizer routine.

The grid of basis functions for constructing the finite-dimensional process S_w^N consists of 6 nodes in the absolute strike direction, 7 nodes in the tenor direction and 6 nodes in the maturity direction.

In order to solve the quadratic problem (10), and thereby find the MAP, the Python library 'quadprog' is used. Sampling of the posterior swaption cubes distribution is done by Hamiltonian Monte Carlo (see Section 2.5). Its implementation is inspired by the algorithm described in [Pakman and Paninski \(2012\)](#).

3.3.2 Calibration and testing methodology

Let us now explain the methodology used for calibration and testing in practice, for our GP regression model and for the SABR model.

GP regression models and the SABR model are fitted on the same randomly generated subsets of the whole data set. Each random subset is divided in two equally sized subsets: one with prices of OTM payer swaptions and ATM swaptions and the other with prices of OTM receiver swaptions and ATM swaptions. For the prediction performance, the whole dataset is chosen as testing dataset.

For the SABR model to calibrate correctly, each pair (T, t) needs at least one ATM swaption observation, one OTM payer swaption observation, and one OTM receiver swaption observation. When a small random training set is chosen, it is possible that some pairs (T, t) end up with one or two observations, then the calibration could lead to absurd results. To address this issue, we decided that each random training set will always contains one observation of the ATM price, and at least one observation of a payer and receiver OTM price for the pairs $(5, 5)$, $(5, 10)$, $(10, 5)$ and $(10, 10)$. Those points are very liquid in the market and quotes are then expected in practice. The details of the SABR calibration procedure are explained in Annex A.

For the sake of comparability, the observed data for the (T, t) pairs $(5, 5)$, $(5, 10)$, $(10, 5)$ and $(10, 10)$ are also included in the training data for the GP.

All results below are produced from the following experimental procedure. We first choose a percentage p . Then the training set is constituted of two random subsets S_1 and S_2 . S_1 includes a random $p\%$ of the ATM swaptions and OTM payer swaptions, and S_2 a random $p\%$ of the ATM swaptions and OTM receiver swaptions. The SABR model is fitted on the union of S_1 and S_2 (there is one set of parameters $(\beta, \sigma, \alpha, \rho)$ for each pair (T, t) , obtained through calibration or interpolation and extrapolation). The GP regression model is fitted two times, once on S_1 and once on S_2 , and then the results for the payer and receiver swaptions are gathered together. For each ATM value, the value kept is the average of the values obtained for the swaptions payer and receiver ATM if there is zero or two observations of this value. If

²A 1-dimensional Matern 5/2 kernel is given as $c(\mathbf{x}, \mathbf{x}') = \sigma^2(1 + \frac{\sqrt{5}|\mathbf{x}-\mathbf{x}'|}{\theta} + \frac{5(\mathbf{x}-\mathbf{x}')}{3\theta^2})\exp(-\frac{\sqrt{5}|\mathbf{x}-\mathbf{x}'|}{\theta})$.

there is only one observation, then the predicted value is taken from the output of the GP which was fitted on this observation. For each percentage of training data p we consider, random sampling, fitting, and testing is repeated 200 times. Then the performance results are averaged and presented in the Tables of Subsection 3.3.3, 3.3.4 and 3.3.5.

3.3.3 Absence of arbitrage in the predicted swaption cube

In this Subsection, we compare the performances of the constrained GP regression model and the SABR model with respect to the arbitrage-free constraints defined in Proposition 1. The grid we use for testing is a grid of maturities, tenors such that, $T = [5, 6, \dots, 19, 20]$, $t = [1, 2, \dots, 29, 30]$ and of 9 strikes K evenly spaced between the minimum and the maximum of K_{Payer} when testing the payer swaption cube, and between the minimum and maximum of $K_{Receiver}$ when testing the receiver swaption cube. More precisely, for the monotonicity and convexity in strike constraints we check how many pairs (T, t) among the 480 pairs of the testing grid respect the constraints. For the in-plane triangular constraint, we check the constraints on 5000 randomly selected T_1, T_2, T_3, K values in the the testing grid. Tables 2 and 3 (respectively Tables 4 and 5) show the results for payer (respectively receiver) swaptions.

| Respect of the AOA constraints GP Payer | Decrease in Strike (Constrained) | Convexity in Strike (Constrained) | In-plane triangular inequality (Constrained) | Decrease in Strike (Unconstrained) | Convexity in Strike (Unconstrained) | In-plane triangular inequality (Unconstrained) |
|---|--|---|--|--|---|--|
| p = 15% | 100% | 100% | 100% | 87.84% | 8.18% | 99.58% |
| p = 30% | 100% | 100% | 100% | 93.79% | 13.22% | 99.97% |
| p = 50% | 100% | 100% | 100% | 94.37% | 10.24% | 99.98% |
| p = 70% | 100% | 100% | 100% | 99.55% | 5.98% | 99.97% |
| p = 90% | 100% | 100% | 100% | 92.17% | 1.51% | 99.92% |

Table 2 Respect of the no-arbitrage constraints for the predicted OTM payer swaptions priced by Constrained and Unconstrained GP regression, for increasing sizes of training set.

| Respect of the AOA constraints SABR Payer | Decrease in Strike | Convexity in Strike | In-plane triangular inequality |
|---|-----------------------|------------------------|-----------------------------------|
| p = 15% | 99.00% | 99.70% | 98.09% |
| p = 30% | 99.49% | 99.82% | 97.96% |
| p = 50% | 99.82% | 99.91% | 99.186% |
| p = 70% | 99.96% | 99.99% | 99.90% |
| p = 90% | 100% | 100% | 100% |

Table 3 Respect of the no-arbitrage constraints for the predicted OTM payer swaptions priced by the chosen SABR model with barycentric mean parameter interpolation, for increasing sizes of training set.

| Respect of the AOA constraints GP Receiver | Increase in Strike (Constrained) | Convexity in Strike (Constrained) | in-plane triangular inequality (Constrained) | Increase in Strike (Unconstrained) | Convexity in Strike (Unconstrained) | In-plane triangular inequality (Unconstrained) |
|--|--|---|--|--|---|--|
| p = 15% | 100% | 100% | 99.51% | 31.46% | 82.67% | 97.10% |
| p = 30% | 100% | 100% | 99.08% | 49.29% | 82.08% | 97.37% |
| p = 50% | 100% | 100% | 99.53% | 56.49% | 86.41% | 98.00% |
| p = 70% | 100% | 100% | 99.78% | 60.03% | 93.69% | 97.95% |
| p = 90% | 100% | 100% | 99.94% | 56.36% | 91.74% | 97.87% |

Table 4 Respect of the no-arbitrage constraints for the predicted OTM receiver swaptions priced by Constrained and Unconstrained GP regression, for increasing sizes of training set.

| Respect of the AOA constraints SABR Receiver | Increase in Strike | Convexity in Strike | In-plane triangular inequality |
|--|-----------------------|------------------------|-----------------------------------|
| p = 15% | 99.77% | 98.73% | 97.67% |
| p = 30% | 99.82% | 99.22% | 97.62% |
| p = 50% | 99.94% | 99.67% | 99.05% |
| p = 70% | 100% | 99.94% | 99.88% |
| p = 90% | 100% | 100% | 100% |

Table 5 Respect of the no-arbitrage constraints for the predicted OTM receiver swaptions priced by the SABR model with barycentric mean parameter interpolation, for increasing sizes of training set.


We can observe that, the prices produced by the SABR model, for a training set of size $p = 50\%$ of the total input data size or lower, do not respect the monotonicity and convexity constraints. Still, the number of pairs (T, t) violating the constraints compared to the total number of pairs is very low. When we check for which pairs (T, t) the constraints are not respected, we see that our method for avoiding absurd parameters values during the calibration of the SABR model does not prevent arbitrage for the observed pairs. For instance, pairs (T, t) with less than 3 observations or no observations of one side of the smile, can produce parameters leading to arbitrage. Arbitrage is also violated for unobserved pairs, for which the parameters are interpolated or extrapolated by barycentric mean. For training sets of size higher than 50% of the total input size, we can consider that the SABR model produces free of arbitrage prices.

For payer and receiver swaptions, the monotonicity and convexity constraints in strike are always ensured for prices estimated through constrained Gaussian process regression (CGPR), unlike prices predicted through unconstrained Gaussian process regression (UGPR). We can remark that receiver swaptions prices estimated through UGPR achieve 91.74% on the convexity in strike test, while payer swaptions achieve only 1.51% even when the 90% of the data set is given as training data. This difference in performance could be explained by the fact that the payer swaption database is not convex everywhere in strike.


In Section 2.2 we explain that we implement the in-plane triangular inequality constraints only for time steps of length 1 year. For payer swaptions we see that the in-plane triangular constraint is always respected by the CGPR, and we can consider that it is also respected by the UGPR. The receiver swaptions Table (Table 4) shows the impact of the 1-year step constraint. While the UGPR does not respect the in-plane triangular constraint everywhere, the CGPR produces a negligible error. Although this error is

negligible, it is proof that this 1-year step constraint for the GPR is not enough to ensure that the in-plane triangular inequality is respected everywhere in the estimated price cube. In [Johnson and Nonas \(2009\)](#) it is mentioned that practitioners only considers the monotonicity and convexity of Proposition 1. Our empirical results supports their decision. Even though our 1-year step constraint does not totally prevent in-plane triangular arbitrage, it could still be helpful for extrapolation.


Figure 1 below presents a few graphical examples of price curves for fixed tenors and maturities, produced by constrained Gaussian process regression and unconstrained Gaussian process regression. We can easily observe the impact of the monotonicity and convexity along the strike direction constraint. In figure (e) and (f) below, we can notice that the mean relative error of the unconstrained GP regression is lower than the one of the constrained GP regression. If we observe those two graphs we can see that the unconstrained GP price crosses the option data price with the higher relative strike while the constrained GP price does not. The constrained GP regression forces the path of the mode of the GP to be convex and increasing/decreasing in strike. Since the payer data set is not convex everywhere, it is not surprising that the convexity constraint increases the testing error.




(a) Constrained GP regression
 $T = 15, t = 1, p = 10\%$




(b) Unconstrained GP regression
 $T = 15, t = 1, p = 10\%$




(c) Constrained GP regression
 $T = 5, t = 1, p = 50\%$



(d) Unconstrained GP regression
 $T = 5, t = 1, p = 50\%$



(e) Constrained GP regression
 $T = 15, t = 1, p = 70\%$



(f) Unconstrained GP regression
 $T = 15, t = 1, p = 70\%$

Figure 1 Constrained and unconstrained GP regression price curves side by side

3.3.4 Calibration performance

Table 6 shows the calibration performance of our models.

| Average calibration error | Constrained GP Mean relative error | SABR Mean relative error | Constrained GP RMSE | SABR RMSE |
|---------------------------|------------------------------------|--------------------------|---------------------|-----------|
| p = 15% | 1.77% | 1.61% | 2.68 | 12.65 |
| p = 30% | 2.29% | 1.47% | 2.78 | 10.56 |
| p = 50% | 2.67% | 2.21% | 3.43 | 11.67 |
| p = 70% | 3.02% | 2.87% | 3.88 | 13.85 |
| p = 90% | 3.25% | 3.22% | 4.11 | 14.59 |

Table 6 Mean relative calibration error and calibration RMSE of the constrained GP regression and SABR model with barycentric mean parameter interpolation, for increasing sizes of training set.

Based on the mean relative error metric, the SABR model calibrates better than the constrained GP model. That was to be expected because the constrained GP regression assumes that the observations are noisy in order to avoid over-fitting and to respect the constraints of no arbitrage. Since the SABR model is not constrained, it gives better mean relative error. We can observe that the mean relative error increases with the size of the training set, and that the error of the GP regression model converges to the error of the SABR model. However, the constrained GP model average RMSE is lower than the average RMSE of the SABR model on all percentages p tested. This means that the estimated standard deviation of the distribution of the absolute errors of the SABR model is in average higher than the one of the constrained GP model.

3.3.5 Prediction performance

Table 7 shows the prediction performance of our models.

| Average prediction error | Constrained GP RMSE | SABR RMSE |
|--------------------------|---------------------|-----------|
| p = 15% | 29.01 | 172.42 |
| p = 30% | 27.83 | 58.55 |
| p = 50% | 9.98 | 33.25 |
| p = 70% | 6.30 | 20.57 |
| p = 90% | 4.71 | 16.54 |

Table 7 RMSE of the constrained GP regression and SABR model with barycentric mean parameter interpolation, for increasing sizes of training set.

Based on the mean RMSE metric, the constrained GP regression performs better than the SABR model for all percentages. For larger and larger training sets, the two models perform better and better

in prediction with respect to the RMSE.

The complexity of the GP regression is dominated by the inversion of the matrix of size (N_{obs}, N_{obs}) , where N_{obs} is the number of observations in S_1 when we fit the payer swaption cube, or the the number of observation in S_2 otherwise. The total complexity when producing the cube is then a $O((N_{obs}^3_{payer} + N_{obs}^3_{receiver}) \times L_1)$, where L_1 is the computation cost of the function 'minimize' with 20 restarts and with the input of the Jacobian matrix of the MLE, defined in (8), as it significantly speeds up the minimization. The complexity of the calibration of the SABR model is a $O((N_{obs}^3_{payer} + N_{obs}^3_{receiver}) \times L_2)$ where N_{obs} is the total number of observed data and L_2 the computation cost of the function 'minimize' with 20 restarts and without inputting a Jacobian. Thus, one of the main drawback of GP based interpolation is the computer time that can be quite substantial for large dataset.

3.3.6 SABR model calibrated on the GP predicted cube

Instead of calibrating the SABR model directly on the observations, it's possible to use the GP-predicted cube as input. Table 8 shows a significant improvement w.r.t the RMSE between the two methods, especially for low percentage of input data. Indeed for these ones, the SABR model cannot be calibrated properly because the smile is not well enough described. In fact, we need at least for each tenor and maturity, one receiver OTM and one payer OTM to define the skew rho and the convexity alpha, plus a point near the ATM to define the model volatility sigma. As the percentage decreases, this condition is less satisfied. Thus, by using the GP cube, we provide an idea of the whole smile and get around the issue. We can see that the gain for $p = 15\%$ is lower than for $p = 30\%$. This can be explained by the fact that the performance of the GP model deteriorates when there are too few observations.

| | | | | | |
|----------------------------------|---------|---------|---------|---------|---------|
| Percentage of data used as input | p = 15% | p = 30% | p = 50% | p = 70% | p = 90% |
| RMSE relative gain | 63% | 68% | 59% | 35% | 19% |

Table 8 Average gain in RMSE when calibrating the SABR model with the output of the GP regression model.

| Respect of the AOA constraints Payer | Decrease in Strike (SABR) | Convexity in Strike (SABR) | in-plane triangular inequality (SABR) | Decrease in Strike (GP-SABR) | Convexity in Strike (GP-SABR) | In-plane triangular inequality (GP-SABR) |
|--------------------------------------|---------------------------|----------------------------|---------------------------------------|------------------------------|-------------------------------|--|
| p = 15% | 99.45% | 99.80% | 98.51% | 99.90% | 99.98% | 100% |
| p = 30% | 99.68% | 99.98% | 96.63% | 100% | 100% | 100% |
| p = 50% | 99.75% | 99.90% | 99.37% | 100% | 100% | 100% |
| p = 70% | 100% | 100% | 100% | 100% | 100% | 100% |
| p = 90% | 100% | 100% | 100% | 100% | 100% | 100% |

Table 9 Respect of the no-arbitrage constraints for the predicted OTM payer swaptions priced by the SABR model calibrated on observations or on the output of the GP model, for increasing sizes of training set.

| Respect of the AOA constraints Receiver | Increase in Strike (SABR) | Convexity in Strike (SABR) | in-plane triangular inequality (SABR) | Increase in Strike (GP-SABR) | Convexity in Strike (GP-SABR) | In-plane triangular inequality (GP-SABR) |
|--|---------------------------------|----------------------------------|---|------------------------------------|-------------------------------------|--|
| p = 15% | 99.87% | 99.20% | 98.22% | 99.99% | 99.53% | 99.52% |
| p = 30% | 99.88% | 99.40% | 96.43% | 100% | 99.76% | 99.97% |
| p = 50% | 99.88% | 99.70% | 99.10% | 100% | 99.80% | 100% |
| p = 70% | 100% | 100% | 100% | 100% | 99.80% | 100% |
| p = 90% | 100% | 100% | 100% | 100% | 99.80% | 100% |

Table 10 Respect of the no-arbitrage constraints for the predicted OTM receiver swaptions priced by the SABR model calibrated on observations or on the output of the GP model, for increasing sizes of training set.

3.4 Uncertainty quantification

By sampling the truncated Gaussian distribution described in section 2.5, it is possible to obtain uncertainties for all the prices estimated through GP regression. The sampling is done by Hamiltonian Monte Carlo, using the mode $\hat{\xi}$ as starting vector. In practice we sample the distribution 5000 times. Therefore we obtain 5000 GP paths of the swaption price cube, all respecting the constraints. These sampling could be used to generate quantiles and calculate model error due to uncertainty. Besides, uncertainties due to interpolation/extrapolation could also be propagated into rate models to see the impact on exotic products pricing. For example, we could calibrate a HJM model with "a percentile 90" cube swaption and "a percentile 10" cube swaption, then price a basket of bermuda options with these two models and study the differences (AVA model).

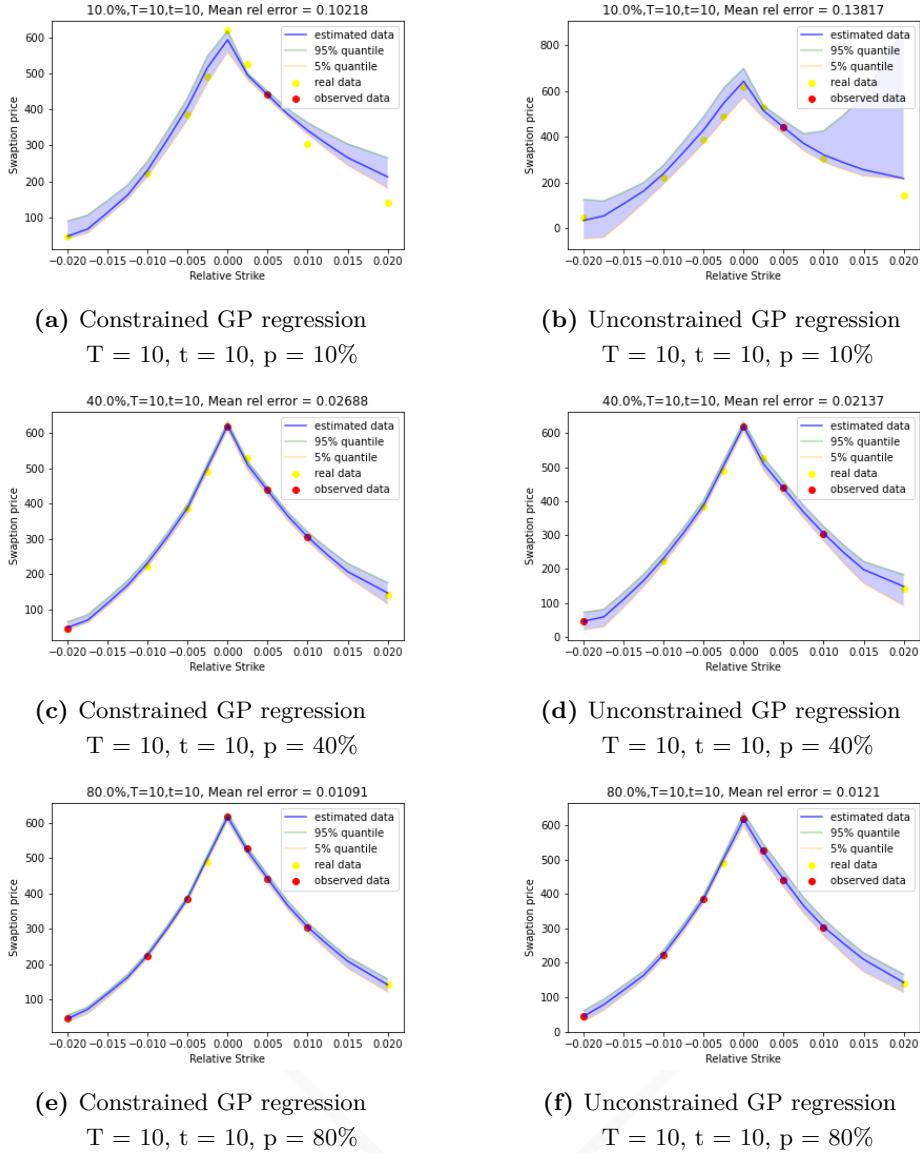


Figure 2 Constrained and unconstrained GP regression price curves, with uncertainties, side by side

Figure 2 shows examples of price curves produced by constrained and unconstrained GP regression, alongside their empirical pointwise quantiles. We can observe that adding observations reduces the size of the pointwise confidence intervals. We can also observe that unconstrained GP regression leads to wider pointwise confidence intervals. In figure (b) the 5% empirical quantile of the posterior of the unconstrained GP is even negative, which is not allowed for prices.

4 Conclusion

In this paper, we introduced a methodology to construct a free of arbitrage swaption price cube based on GP regression. We applied this methodology on a data set of ATM, OTM payer and OTM receiver swaptions. We showed that the constrained Gaussian process regression leads to a swaption cube free of arbitrage along the strike direction for any size of the training data. Even though the practitioners were usually satisfied by the fact that there is no Call spreads and Butterfly spreads, we tried to go further by using in practice the in-plane triangular inequality introduced by [Johnson and Nonas \(2009\)](#).

We then compared the constrained Gaussian process model to a state of the art SABR model. In Section 3.3, we obtained that the constrained Gaussian process performed better in prediction than the SABR model. Besides, the predictive performance of SABR seems to improve significantly when the model is calibrated on the GP-predicted cube. Finally, we show that constraints reduce the error and uncertainty for extrapolations.

5 Further research

A topic left to be discussed is the updating of the GP model when a new quote streams in. A straightforward solution could be to use the last estimation of the maximum log-likelihood as the new starting point of the optimizer routine. However, since the computational cost of the GP regression increases when the number of observations increases, it could be too time consuming to be used in practice. For this problem, Bayesian Optimization could be considered. As explained in [Shahriari et al. \(2016\)](#), Bayesian Optimization consists in selecting the $n + 1$ th observation, based on an acquisition function computed with the current posterior distribution. Article [Yin et al. \(2019\)](#) presents an other way of selecting the $n + 1$ th observation, based on a MSE criteria.

A Calibration of the SABR model

The calibration procedure we applied is as follows:

1. Calibration of the parameters of the SABR model for the pairs (5, 5), (5, 10), (10, 5) and (10, 10).
2. For all the other pairs, take as a starting value for the parameters α , ρ and σ the value of the calibrated parameters of the nearest neighbour among the pairs (5, 5), (5, 10), (10, 5) and (10, 10).
3. Then for the pairs (T, t) calibrate:
 - (a) α , ρ and σ if there is at least 3 observations.
 - (b) α only if there is one or two observations of payer swaptions (or receiver).
 - (c) σ only if the only observation is the price of a swaption ATM.
 - (d) ρ and α if there is an observation of a payer swaption, and an observation of a receiver swaption.
 - (e) α and σ if we observed one swaption ATM and one payer or receiver swaption.

The calibration of the parameters of the model is made through the function 'minimize' introduced in 3.3.1, with the same number of random restarts.

References

- Bachoc, Lagnoux, and López-Lopera (2019). Maximum likelihood estimation for gaussian processes under inequality constraints. *Electronic Journal of Statistics* 13(2), 2921–2969.
- Chataigner, M., A. Cousin, S. Crépey, M. Dixon, and D. Gueye (2020). Beyond surrogate modeling: Learning the local volatility via shape constraints.
- Cousin, A., H. Maatouk, and D. Rullière (2016). Kriging of financial term-structures. *European Journal of Operational Research* 255(2), 631–648.
- Da Veiga, S. and A. Marrel (2012). Gaussian process modeling with inequality constraints. *Annales de la Faculté des sciences de Toulouse : Mathématiques Ser. 6*, 21(3), 529–555.
- Deloitte (2019). The perfect smile.
- Hagan and Lesniewski (2002). Managing smile risk. *Wilmott*.
- Hagan, P. and M. Konikov (2004). Interest rate volatility cube: Construction and use.
- Hagan, P., D. Kumar, A. Lesniewski, and D. Woodward (2014). Arbitrage-free sabr. *Wilmott 2014*.
- Johnson, S. and B. Nonas (2009). Arbitrage-free construction of the swaption cube. *Commerzbank Corporates and Markets*.
- López-Lopera, Bachoc, Durrande, and O. Roustant (2017). Finite-dimensional gaussian approximation with linear inequality constraints.
- Maatouk, H. and X. Bay (2017). Gaussian process emulators for computer experiments with inequality constraints. *Mathematical Geosciences, Springer Verlag, 2017*, 49(5), 557–582.
- Mazzoni, T. (2018). Slicing the swaption cube. *Wilmott* (94), 38–49.
- Pakman, A. and L. Paninski (2012). Exact hamiltonian monte carlo for truncated multivariate gaussians. *Journal of Computational and Graphical Statistics* 23, 518 – 542.
- Rasmussen and Williams (2006). Gaussian processes for machine learning. *The MIT Press*.
- Ruf, J. and W. Wang (2019). Neural networks for option pricing and hedging: A literature review. *PSN: Other Political Methods: Quantitative Methods (Topic)*.
- Shahriari, B., K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas (2016). Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE* 104(1), 148–175.
- Yin, H., J. Ma, K. Dong, Z. Peng, P. Cui, and C. Yang (2019). Model updating method based on kriging model for structural dynamics. *Shock and Vibration*.