The quadratic rough Heston model and the joint S&P 500/Vix smile calibration problem

Fitting SPX and Vix smiles simultaneously is one of the most challenging problems in volatility modelling. A long-standing conjecture is that it may not be possible to jointly calibrate these two quantities using a model with continuous sample paths. Jim Gatheral, Paul Jusselin and Mathieu Rosenbaum present the quadratic rough Heston model as a counterexample to this conjecture. The key idea is the combination of rough volatility with a price-feedback (Zumbach) effect

The volatility index, or Vix, was introduced in 1993 by the Chicago Board Options Exchange (CBOE for short). It was originally designed, according to the CBOE, to 'measure the market's expectation of 30-day volatility implied by at-the-money [Standard & Poor's 100] index option price' (see CBOE 2019). Since 2003, the Vix has been redefined as the square root of the price of a specific basket of options on the Standard & Poor's 500 index (SPX) with maturity 30 days. The basket coefficients are chosen so that, at any time t, the Vix represents the annualised square root of the price of a contract with payout equal to $-2/\Delta \log(S_{t+\Delta}/S_t)$, where $\Delta = 30$ days and S denotes the value of the SPX. Consequently, it can be formally written via risk-neutral expectation in the form:

$$\mathrm{Vix}_t = \sqrt{-\frac{2}{\Delta}\mathbb{E}\left[\log\left(\frac{S_{t+\Delta}}{S_t}\right) \,\middle|\, \mathcal{F}_t\right]} \times 100 \qquad (1)$$

where $(\mathcal{F}_t)_{t \geqslant 0}$ is the natural filtration of the market.

Since 2004, investors have been able to trade Vix futures. To quote CBOE (2019), they 'provide market participants with a variety of opportunities to implement their view using volatility trading strategies, including risk management, alpha generation and portfolio diversification'.[1] Subsequently, in 2006, the CBOE introduced Vix options:

> providing market participants with another tool to manage volatility. Vix options enable market participants to hedge portfolio volatility risk distinct from market price risk and trade based on their view of the future direction or movement of volatility.[2]

Those products are now among the most liquid financial instruments in the world. Indeed, more than 500,000 Vix options are traded each day, with most of the liquidity concentrated on the first three monthly contracts.

Although more vega is now traded in the Vix market than in the SPX market, the wide bid-ask spreads in the former betray its lack of maturity. One of the reasons behind these wide spreads is the market lacks a reliable pricing methodology for Vix options. Since the Vix is, by definition, a derivative of the SPX, any reasonable methodology must necessarily be consistent with the pricing of SPX options. Designing a model that jointly calibrates SPX and Vix options prices is known to be extremely challenging. Indeed, this problem is sometimes considered to be the holy grail of volatility modelling. We will simply refer to it as the joint calibration problem.

The joint calibration problem has been extensively studied by Julien Guyon, who provides a review of various approaches (Guyon 2019b). We can split the different attempts to solve it into three categories. In what is probably the most technical and the most original proposal, as well as the first to have succeeded in obtaining a perfect joint calibration, the problem is interpreted as a model-free constrained martingale transport problem, as initially observed by De Marco & Henry-Labordere (2015). Using this viewpoint, Guyon (2019b) manages to get a perfect calibration of a Vix options smile at time T_1 and a SPX options smile at T_1 and $T_2 = T_1 + 30$ days. As noted by the author, although this methodology can theoretically be extended to any set of maturities, it is much more intricate in practice because of the computational complexity involved.

This drawback is avoided in the second and third approaches Guyon (2019b) outlines, where the models are in continuous time. Such models have the advantage of relying on observable properties of assets, and so allow for practical intuition regarding their dynamics. The second approach involves attempting joint calibration with models where SPX trajectories are continuous (see, in particular, Goutte *et al* 2017). Unfortunately, as of yet, continuous models have not been completely successful in this task. An interpretation for this failure is given by Guyon (2019b), who explains

> the very large negative skew of short-term SPX options, which in continuous models implies a very large volatility of volatility, seems inconsistent with the comparatively low levels of Vix implied volatilities.

To circumvent this issue, it is then natural to think of rough volatility models, as recently introduced by Gatheral *et al* (2018). However, these models also appear to have been unsuccessful thus far (see Guyon 2018).

The last approach is to allow for jumps in the dynamic of the SPX (see Baldeaux & Badran 2014; Cont & Kokholm 2013; Kokholm & Stisen 2015; Pacati *et al* 2018; Papanicolaou & Sircar 2014). By doing so, one can reconcile the skewness of SPX options with the level of Vix implied volatilities. Nevertheless, besides those of Cont & Kokholm (2013) and Pacati *et al* (2018), existing models with jumps do not really achieve a satisfying accuracy

[1] *See also https://bit.ly/39EsFm0.*
[2] *See https://bit.ly/2X5jjx9.*

for the joint calibration problem. Specifically, most of them fail to reproduce Vix smiles for maturities shorter than one month.

As an aside, even though some models with jumps may satisfactorily resolve the joint calibration problem, they are unsatisfactory in other respects. For example, perfect hedging is not possible in such models; in contrast, under rough volatility derivatives, hedging is fully understood: this is shown by El Euch & Rosenbaum (2018, 2019). Moreover, jumps are conventionally modelled as Lévy jumps, giving rise to unrealistic model time series properties that are at odds with those observed empirically, specifically the clustering of large moves in the underlying. One might imagine trying to fix the latter problem by modelling with self-exciting jump processes, but, in the end, that would lead back to rough volatility models, which can be regarded as special limits of self-exciting jump processes.

In summary, according to Guyon (2019b), despite many efforts:

> so far all the attempts at solving the joint SPX/Vix smile calibration problem [using a continuous time model have] only produced imperfect, approximate fits.

In particular, regarding continuous models, Guyon concludes: 'joint calibration seems out of the reach of continuous-time models with continuous SPX paths'. In this article, we provide a counterexample to this conjecture: namely, a model with continuous SPX and Vix paths that enables us to fit SPX and Vix options smiles simultaneously.

Rough volatility and the Zumbach effect

Recently, rough volatility models – where volatility trajectories, though continuous, are very irregular – have generated a lot of attention. The reason for this success is the ability of these very parsimonious models to reproduce all of the main stylised facts of historical volatility time series and to fit SPX options smiles (see Bayer *et al* 2016; El Euch *et al* 2019b; Gatheral *et al* 2018). One particularly interesting rough volatility model is the rough Heston model introduced in El Euch & Rosenbaum (2019). As its name suggests, it is a rough version of the classical Heston model. This model arises as the limit of natural Hawkes process-based models of price and order flow (see, for example, Jusselin & Rosenbaum 2018). Moreover, there is a quasi-closed-form formula for the characteristic function of the rough Heston model, just as in the classical case. So, fast pricing of European options is possible (see El Euch *et al* (2019b) and the references therein).

Despite these successes, a subtle question raised by Jean-Philippe Bouchaud remains: can a rough volatility model reproduce the so-called Zumbach effect? This is the observation originally due to Gilles Zumbach (see Zumbach 2009, 2010) that financial time series are not time-reversal invariant. To answer this question, we introduce two notions, each of which corresponds to different aspects of the Zumbach effect:

- The weak Zumbach effect:[3] past-squared returns forecast future-integrated volatilities better than past-integrated volatilities forecast future-squared returns. This property is not satisfied in classical stochastic volatility models. However, rough stochastic volatility models are consistent with the weak Zumbach effect: see El Euch *et al* (2019a) for explicit computations using a rough Heston model.
- The strong Zumbach effect: conditional dynamics of volatility with respect to the past depend not only on the past volatility trajectory but also on the historical price path. Specifically, price trends tend to increase volatility (see Zumbach 2010). Such feedback of the historical price path on volatility also occurs on implied volatility, as illustrated by Zumbach (2010). Rough stochastic volatility models such as the rough Heston model are not consistent with the strong Zumbach effect (see El Euch & Rosenbaum 2018).

The quest for a rough volatility model consistent with the strong Zumbach effect and the empirical success of quadratic Hawkes process-based models documented by Blanc *et al* (2017) led to the development of super-Heston rough volatility models (Dandapani *et al* 2019). These extensions of the rough Heston model arise as limits of quadratic Hawkes process-based microstructural models, just as the rough Heston model arises as the continuous-time limit of a linear Hawkes process-based microstructural model.

The idea of using super-Heston rough volatility models to solve the joint calibration problem came after a presentation by Julien Guyon at École Polytechnique in March 2019. In this talk, Guyon highlighted a necessary condition for a continuous model to fit SPX and Vix smiles simultaneously: the inversion of convex ordering between volatility and the local volatility implied by option prices (see Guyon 2019a). The intuition behind this condition could be reinterpreted as some kind of strong Zumbach effect. It was therefore natural for us to investigate the ability of super-Heston rough volatility models to solve the joint calibration problem.

The quadratic rough Heston model

The quadratic rough Heston model we consider is essentially a special case of the super-Heston rough volatility models of Dandapani *et al* (2019). The joint dynamics of the asset S (here, the SPX) and its spot variance V satisfy:

$$\begin{aligned} \mathrm{d}S_t &= S_t \sqrt{V_t}\,\mathrm{d}W_t \\ V_t &= a(Z_t - b)^2 + c \end{aligned}$$

where W is a Brownian motion and a, b and c are positive constants. This model is of rough Heston type, in the sense that weighted past price returns are drivers of the volatility dynamics. More precisely:

$$\begin{aligned} Z_t = {} & \int_0^t (t-s)^{\alpha-1} \frac{\lambda}{\Gamma(\alpha)} (\theta_0(s) - Z_s)\,\mathrm{d}s \\ & + \int_0^t (t-s)^{\alpha-1} \frac{\lambda}{\Gamma(\alpha)} \eta \sqrt{V_s}\,\mathrm{d}W_s \end{aligned} \tag{2}$$

with $\alpha \in (1/2, 1)$, $\lambda > 0$, $\eta > 0$ and θ_0 a deterministic function. In this special case of a super-Heston rough volatility model, the asset S and its volatility depend on the history of only one Brownian motion. The model is thus a pure feedback model; volatility is driven only by the price dynamics, with no additional source of randomness. In general, of course, the volatility process does not need to depend only on the Brownian motion driving the asset price S. For simplicity, we will refer to (2), a pure feedback version of a super-Heston rough volatility model, as the quadratic rough Heston model.

[3] *This is typically considered in the econophysics literature; see Zumbach (2009).*

As in the general case of super-Heston rough volatility models, because the effect of past returns on Z cannot be reduced to an influence of past volatility dynamics on Z, the quadratic rough Heston model also exhibits the strong Zumbach effect (see Dandapani *et al* (2019) for more details).

The quadratic rough Heston process. The process Z_t may be understood as a weighted moving average of past price log returns. Indeed, from El Euch & Rosenbaum (2018, lemma A.1), we have:

$$Z_t = \int_0^t f^{\alpha,\lambda}(t-s)\theta_0(s)\,\mathrm{d}s + \int_0^t f^{\alpha,\lambda}(t-s)\eta\sqrt{V_s}\,\mathrm{d}W_s$$

where $f^{\alpha,\lambda}(t)$ is the Mittag-Leffler density function defined for $t \geqslant 0$ as:

$$f^{\alpha,\lambda}(t) = \lambda t^{\alpha-1} E_{\alpha,\alpha}(-\lambda t^{\alpha})$$

with:

$$E_{\alpha,\beta}(z) = \sum_{n\geqslant 0} \frac{z^n}{\Gamma(\alpha n+\beta)}$$

The variable Z_t is therefore path-dependent: a weighted average of past returns of the type typically considered in path-dependent volatility models. As explained by Guyon (2014), modelling with path-dependent variables is a natural way to reproduce the fact that volatility depends on recent price changes. However, the kernels used to model this dependency are typically exponential. Here, a crucial idea – motivated by our previous work (Dandapani *et al* 2019) – is to use a rough kernel: more precisely, to use the Mittag-Leffler density function. Thanks to this kernel, the 'memory' of Z decays as a power law, and Z is highly sensitive to recent returns since:

$$f^{\alpha,\lambda}(t) \underset{t\to+\infty}{\sim} \frac{\alpha}{\lambda\Gamma(1-\alpha)} t^{-\alpha-1}$$

and

$$f^{\alpha,\lambda}(t) \underset{t\to 0^+}{\sim} \frac{\lambda}{\Gamma(\alpha)} t^{\alpha-1}$$

This essentially means long periods of trends or sudden upwards or downwards moves of the price generate large values for $|Z|$ and, thus, high volatility, particularly when Z is negative. Such a link is clearly observable in the data: see figure 1, where the Vix spikes almost instantaneously after large negative returns of the SPX, before decreasing slowly afterwards. In figure 2, we plot an example of sample paths of the SPX and the Vix in our model. The feedback of negative price trends on volatility is very well reproduced. Finally, the choice of $f^{\alpha,\lambda}$ as kernel ensures the volatility process is rough, with a Hurst parameter equal to:

$$H = \alpha - 1/2$$

As shown by Gatheral *et al* (2018), this enables us to reproduce the behaviour of both the historical volatility time series and the SPX implied volatility surface, provided H is taken to be of order 0.1.

As explained above, an immediate consequence of the feedback effect is negative price trends generating high volatility levels. However, such trends also impact the instantaneous variance of volatility in our model. To see this, consider the classical case with $\alpha = 1$. In that case, an application of Itô's formula gives this up to a drift term:

$$\mathrm{d}V_t = 2a(Z_t - b)\lambda\eta\sqrt{V_t}z\,\mathrm{d}W_t$$

Thus, the 'variance of instantaneous variance' coefficient is proportional to $a(Z_t - b)^2$, which up to c is equal to the variance of $\log S$. Thus, when volatility is high, the volatility of volatility is also high. In particular, conditional on a large downwards move in SPX, we would expect V to be high along with the volatility of V. This explains why our model generates upwards-sloping Vix smiles.

We remark that incorporating the influence of price trends on volatility and on the instantaneous variance of volatility is the main motivation underlying the model of Goutte *et al* (2017). That model, although not solving the joint calibration problem, is probably the best of the continuous models introduced so far. In this switching model, the price follows a classical Heston dynamic, where the parameters can change depending on the value of a hidden Markov chain with three states. It is motivated by a 100-day rolling calibration of the classical Heston model performed by the authors (see Goutte *et al* 2017, figure 2). This rolling calibration suggests volatility, leverage and volatility of volatility are higher in periods of crisis. Hence, Goutte *et al* introduce a Markov chain to trigger crisis phases and switch the parameters of the Heston model depending on the situation. The three

1 SPX (in blue) and Vix (in red) from November 25, 2004 to November 25, 2019

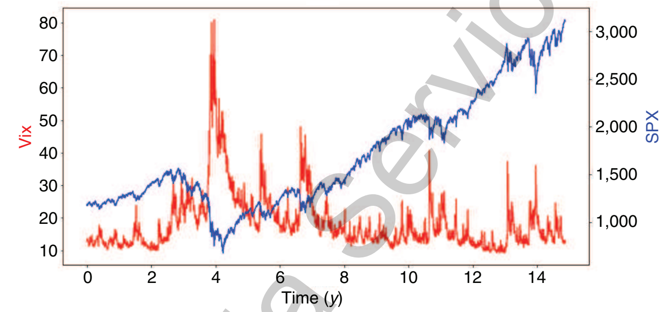

2 SPX (in blue) and Vix (in red) from simulation of the quadratic rough Heston model

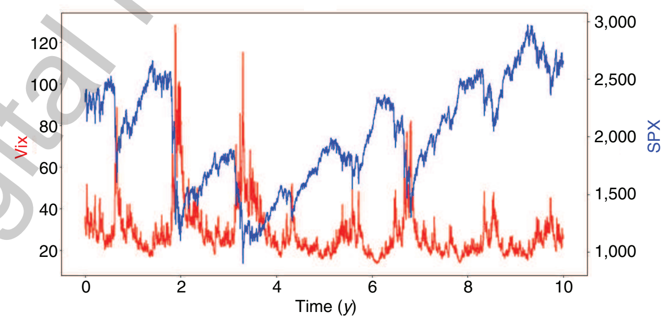

3 Implied volatility on SPX options for May 19, 2017

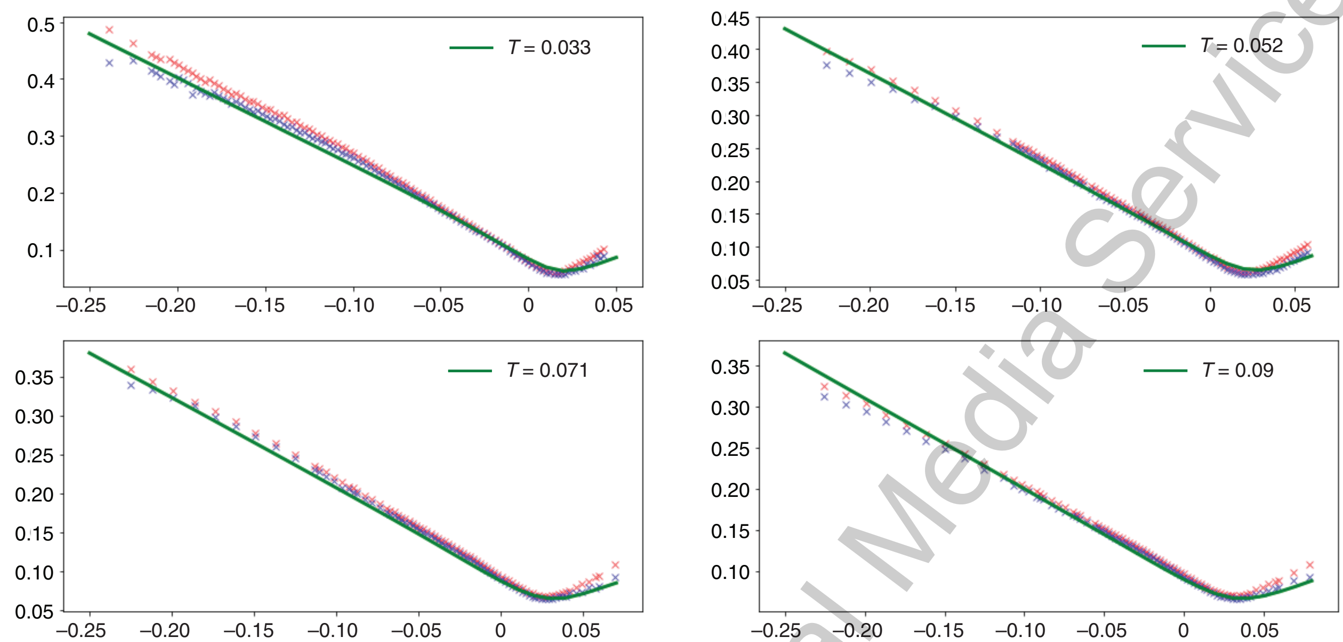

The blue and red points are, respectively, the bid and ask of market implied volatilities. The implied volatility smiles from the model are green. The strikes are in logmoneyness, and T is time to expiry in years

possible states of the chain can therefore be interpreted as corresponding to the following situations:

- flat or increasing SPX;
- transition phase between flat SPX and crisis;
- crisis with dramatically decreasing SPX.

The Markov chain of Goutte *et al* (2017) can therefore be seen as an *ad hoc* version of the process Z in the quadratic rough Heston model.

Parameter interpretation. The parameters a, b and c in the specification:

$$V_t = a(Z_t - b)^2 + c$$

can be interpreted in the following way.

- c represents the minimal instantaneous variance. When calibrating the model, we use c to shift the smiles of SPX options upwards or downwards.
- $b > 0$ encodes the asymmetry of the feedback effect. Indeed, for the same absolute value of Z, the volatility is higher when Z is negative than when it is positive. Such asymmetry aims at reproducing the empirical behaviour of the Vix. This is illustrated in figure 1, where we can observe that the Vix spikes when the SPX tumbles down, but not after it goes up. From a calibration point of view, the higher b, the more SPX options smiles are shifted to the right.
- a is the sensitivity of the volatility to the feedback of price returns. The greater a, the greater the role of feedback in the model, and the higher the volatility of volatility. Consistent with this, SPX smiles become more extreme as a increases.

Infinite-dimensional Markovian representation. Although the quadratic rough Heston model is not Markovian in the variables (S, V), it does admit an infinite-dimensional Markovian representation. Inspired by the computations of El Euch & Rosenbaum (2018), we obtain that, for any t and t_0 positive:

$$\begin{aligned} Z_{t_0+t} = {} & \int_0^t (t-s)^{\alpha-1} \frac{\lambda}{\Gamma(\alpha)} (\theta_{t_0}(s) - Z_{t_0+s}) \, \mathrm{d}s \\ & + \int_0^t (t-s)^{\alpha-1} \frac{\lambda}{\Gamma(\alpha)} \eta \sqrt{V_{t_0+s}} \, \mathrm{d}W_{t_0+s} \end{aligned} \tag{3}$$

where θ_{t_0} is a $\mathcal{F}_{t_0}$-measurable function. More precisely, θ_{t_0} is given by:

$$\begin{aligned} \theta_{t_0}(u) = {} & \theta_0(t_0+u) + \frac{\alpha}{\lambda \Gamma(1-\alpha)} \\ & \times \int_0^{t_0} (t_0 - v + u)^{-1-\alpha} Z_v z \, \mathrm{d}v \end{aligned}$$

Equation (3) implies the law of $(S_t, V_t)_{t \geqslant t_0}$ only depends on S_{t_0} and θ_{t_0}. In view of (1), and using the same methodology as in El Euch & Rosenbaum (2018), it means we can express the Vix at time t as a function of θ_t and S_t. Consequently, we can write the price of any European option with payout depending on the SPX and the Vix as a function of time, S and θ.

Numerical results

In this section, we illustrate how successfully we can fit both SPX and Vix smiles on May 19, 2017, one of the dates considered by El Euch *et al* (2019b)

4 Implied volatility on Vix options for May 19, 2017

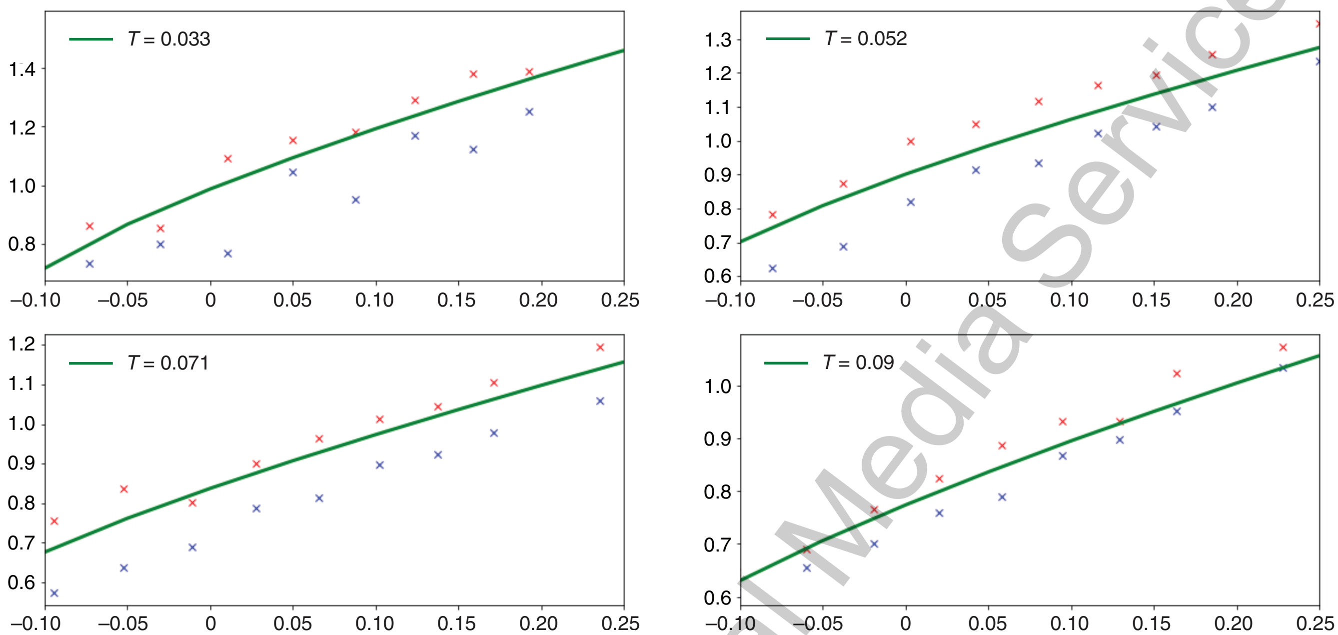

The blue and red points are, respectively, the bid and ask of market implied volatilities. The implied volatility smiles from the model are green. The strikes are in logmoneyness, and T is time to expiry in years

that is otherwise randomly chosen.[4] We focus on short expirations, from two to five weeks, where the bulk of Vix liquidity is found. Moreover, short-dated smiles are typically fitted poorly by conventional models.

In the quadratic rough Heston model, the function $\theta_0(\cdot)$ needs to be calibrated to market data. In the rough Heston model, there is a simple bijection between $\theta_0(\cdot)$ and the forward variance curve. In the quadratic rough Heston model, this connection is more intricate; for simplicity, then, we choose the following restrictive parametric form for Z:

$$Z_t = Z_0 - \int_0^t (t-s)^{\alpha-1} \frac{\lambda}{\Gamma(\alpha)} Z_s \, \mathrm{d}s + \int_0^t (t-s)^{\alpha-1} \frac{\lambda}{\Gamma(\alpha)} \eta \sqrt{V_s} \, \mathrm{d}W_s$$

which is equivalent to taking:

$$\theta_0(t) = \frac{Z_0}{\lambda \Gamma(1-\alpha)} t^{-\alpha}$$

Allowing $\theta_0(\cdot)$ to belong to a larger space would obviously lead to even better results, but it would require a more complex calibration methodology. Thus, we are left to calibrate the parameters $\nu = (\alpha, \lambda, a, b, c, Z_0)$. We use the following objective function:

$$F(\nu) = \frac{1}{\#\mathcal{O}^{\mathrm{SPX}}} \sum_{o \in \mathcal{O}^{\mathrm{SPX}}} (\sigma^{o,\mathrm{mid}} - \sigma^{o,\nu})^2 + \frac{1}{\#\mathcal{O}^{\mathrm{Vix}}} \sum_{o \in \mathcal{O}^{\mathrm{Vix}}} (\sigma^{o,\mathrm{mid}} - \sigma^{o,\nu})^2,$$

where $\mathcal{O}^{\mathrm{SPX}}$ is the set of SPX options; $\mathcal{O}^{\mathrm{Vix}}$ is the set of Vix options; $\sigma^{o,\mathrm{mid}}$ denotes the market 'mid' implied volatility for the option o; and $\sigma^{o,\nu}$ is the implied volatility of the option o in the quadratic rough Heston model, with parameter ν obtained by Monte Carlo simulations. To calibrate the model, we minimise the function F over a grid centred around an initial guess ν_0.

We obtain the following parameters:[5]

$$\begin{aligned} \alpha &= 0.51 \\ \lambda &= 1.2 \\ a &= 0.384 \\ b &= 0.095 \\ c &= 0.0025 \\ Z_0 &= 0.1 \end{aligned} \tag{4}$$

The corresponding SPX and Vix options smiles are plotted in figures 3 and 4.

Despite the fact that our calibration methodology is suboptimal and we only have six parameters, Vix smiles generated by the model with parameters (4) fall systematically within market bid-ask spreads. The overall shape of the shorter-dated SPX smiles shown in figure 3 are reproduced accurately.

Obviously, fits can be made even better by reducing the range of strikes of interest or by fine-tuning the calibration, notably through improving the $\theta_0(\cdot)$ function. We are currently working on a fast calibration approach, inspired by recent works on neural networks. ■

[4] *Market data is from OptionMetrics via Wharton Data Research Services (WRDS).*

[5] *Note that we can always take $\eta = 1$ up to a rescaling of the other parameters.*

Jim Gatheral is presidential professor of mathematics at Baruch College, CUNY in New York; Paul Jusselin is a PhD student in financial mathematics at École Polytechnique in Paris; and Mathieu Rosenbaum is a full professor at École Polytechnique in Paris, where he holds the chair in Analytics and Models for Regulation and is co-head of the quantitative finance (El Karoui) master's program. They thank Julien Guyon for numerous inspiring discussions, and Stefano De Marco and the two reviewers for relevant comments. Paul Jusselin and Mathieu Rosenbaum gratefully acknowledge the financial support of the ERC (grant 679836: Staqamof) and of the chair Analytics and Models for Regulation.
Email: jim.gatheral@baruch.cuny.edu, paul.jusselin@polytechnique.edu, mathieu.rosenbaum@polytechnique.edu.

REFERENCES

Baldeaux J and A Badran, 2014
Consistent modelling of VIX and equity derivatives using a 3/2 plus jumps model
Applied Mathematical Finance 21(4), pages 299–312

Bayer C, P Friz and J Gatheral, 2016
Pricing under rough volatility
Quantitative Finance 16(6), pages 887–904

Blanc P, J Donier and J-P Bouchaud, 2017
Quadratic Hawkes processes for financial prices
Quantitative Finance 17(2), pages 171–188

CBOE, 2019
Vix: CBOE volatility index
White Paper, Chicago Board Options Exchange, available at www.cboe.com/micro/vix/vixwhite.pdf

Cont R and T Kokholm, 2013
A consistent pricing model for index options and volatility derivatives
Mathematical Finance 23(2), pages 248–274

Dandapani A, P Jusselin and M Rosenbaum, 2019
From quadratic Hawkes processes to super-Heston rough volatility models with Zumbach effect
Preprint (arXiv:1907.06151)

De Marco S and P Henry-Labordere, 2015
Linking vanillas and VIX options: a constrained martingale optimal transport problem
SIAM Journal on Financial Mathematics 6(1), pages 1171–1194

El Euch O and M Rosenbaum, 2018
Perfect hedging in rough Heston models
Annals of Applied Probability 28(6), pages 3813–3856

El Euch O and M Rosenbaum, 2019
The characteristic function of rough Heston models
Mathematical Finance 29(1), pages 3–38

El Euch O, J Gatheral, R Radoičić and M Rosenbaum, 2019a
The Zumbach effect under rough Heston
Quantitative Finance 20(2), pages 235–241

El Euch O, J Gatheral and M Rosenbaum, 2019b
Roughening Heston
Risk May, pages 84–89

Gatheral J, T Jaisson and M Rosenbaum, 2018
Volatility is rough
Quantitative Finance 18(6), pages 933–949

Goutte S, A Ismail and H Pham, 2017
Regime-switching stochastic volatility model: estimation and calibration to VIX options
Applied Mathematical Finance 24(1), pages 38–75

Guyon J, 2014
Path-dependent volatility
Working Paper, SSRN, available at https://bit.ly/2w8bwDz

Guyon J, 2018
On the joint calibration of SPX and VIX options
Presentation, available at https://bit.ly/2UzL4MG

Guyon J, 2019a
Inversion of convex ordering in the VIX market
Working Paper, SSRN, available at https://bit.ly/3dLMOK9

Guyon J, 2019b
The joint S&P 500/VIX smile calibration puzzle solved
Working Paper, SSRN, available at https://bit.ly/3bIiSNe

Jusselin P and M Rosenbaum, 2018
No-arbitrage implies power-law market impact and rough volatility
Preprint (arXiv:1805.07134)

Kokholm T and M Stisen, 2015
Joint pricing of VIX and SPX options with stochastic volatility and jump models
Journal of Risk Finance 16(1), pages 27–48

Pacati C, G Pompa and R Renò, 2018
Smiling twice: the Heston++ model
Journal of Banking & Finance 96, pages 185–206

Papanicolaou A and R Sircar, 2014
A regime-switching Heston model for VIX and S&P 500 implied volatilities
Quantitative Finance 14(10), pages 1811–1827

Zumbach G, 2009
Time reversal invariance in finance
Quantitative Finance 9(5), pages 505–515

Zumbach G, 2010
Volatility conditional on price trends
Quantitative Finance 10(4), pages 431–442