Singular exotic perturbation

Florian Monciaud and Adil Reghai combine singular perturbation techniques with a price-adjustment argument to analyse the impact of the smile dynamics, ie, the price difference between local stochastic volatility and local volatility on exotic products. They obtain an elegant formula that is exact for vanilla options and they propose a set of well-chosen scenarios to compute the impact efficiently

The main driver when selecting a model for pricing and risk management derivatives products is its capacity to explain the profit and loss (PnL) evolution (Bergomi 2018; Reghai & Kettani 2020; Reghai 2015). Choosing the right model to successfully price and hedge financial instruments is based on a careful study of the financial structure to be considered and the market in which it evolves. The quantitative finance literature initially promoted the local volatility (LV) and then the pure stochastic volatility (SV) models as a means of explaining the observed market smile. However, when we consider the dynamic hedging of exotic products such as autocalls, we rapidly conclude that matching the smile is not enough; we also need to control the way in which the latter evolves when the spot moves. Neither the LV model nor the SV model can describe the smile and its evolution properly. However, a fine-tuned mix between the two gives the flexibility to fit both the vanilla options and the way they evolve when the spot moves. These kinds of models are known in the literature as the local stochastic volatility (LSV) models.

The literature on this topic is vast and covers a diversity of approaches to the definition of the models or the way to calibrate them. Lipton (2002) introduces a universal diffusion model presentation with applications to foreign exchange derivatives. Lipton *et al* (2014) survey LSV models applied to a variety of first-generation exotics. Many papers cover the different approaches to calibration, as this is one of the most important building blocks of the computation. These include Monte Carlo based approaches (Henry-Labordère 2009) and McKean's particle method (Guyon & Henry-Labordère 2011). Fouque *et al* (2011) take a perturbation approach to SV models with one or two factors. However, the calibration of the vanilla is not considered therein. Reghai *et al* (2012) introduce a mixing weight to control the correlation and the volatility of volatility of the process. The LSV impact is computed for exotics. However, this method does not span all possible stochastic volatility parameters and only works for mild parameters. Hagan *et al* (2018) perform a singular perturbation analysis on a term-structure stochastic alpha-beta-rho model with fast-varying parameters. However, this is a pure stochastic volatility model.

The objective of this paper is to apply a singular perturbation approach in the case of the LSV model. For that purpose, we use a singular perturbation approach without focusing on the vanilla calibration as described in Fouque *et al* (2011). We then recover the vanilla fit using a price adjustment, as described in Hagan (2005). The obtained formulas are then computed effectively using a well-chosen scenario, as is done in the computation of the Exotic Theta by Bergomi (2018).

This is the main motivation of this work. We indeed propose an extremely fast algorithm that prices the LSV impact at a much lower computational cost than traditional LSV implementations. It is based on LV prices calculated on a well-chosen volatility scenario. This is not only a game changer for real-time risk management but also a powerful way to infer the stochastic volatility parameters in the presence of exotic prices. One last property of the proposed technique is that it reprices vanilla options perfectly, removing the known burden of LSV calibration. Ultimately, this formula offers a rapid, robust and easy implementation of an essential model in real-time management. Finally, this technique is quite general and could open the door for other industrial applications, which will make it possible to enhance all those existing perturbation formulas developed over the years that did not find industrial applications due to the lack of their equivalent in the presence of the smile.

Problem formulation and main result

In this section, for clarity we first formulate the problem, and then go on to state the main result of the paper.

Formulation of the problem. Assume that the dynamics of the stock are given by the following LSV process:

$$\frac{\mathrm{d}S_t}{S_t} = \frac{\sigma_{\mathrm{D}}(t, S_t)}{f_\varepsilon(t, S_t)} h(Y_t^\varepsilon)\,\mathrm{d}W_t \tag{1}$$

where $h(x) = \mathrm{e}^x$ and Y_t satisfies an Orstein-Uhlenbeck process:

$$\mathrm{d}Y_t^\varepsilon = -\frac{\kappa}{\varepsilon} Y_t^\varepsilon\,\mathrm{d}t + \frac{\nu}{\sqrt{\varepsilon}}\,\mathrm{d}B_t \tag{2}$$

with $\langle \mathrm{d}B_t, \mathrm{d}W_t \rangle = \rho\,\mathrm{d}t$.

In order to preserve the vanilla calibration, we make an adjustment, f, of the Dupire local volatility $\sigma_{\mathrm{D}}(t, S_t)$, which we define as (Henry-Labordère 2009):

$$f_\varepsilon^2(t, S) = \mathbb{E}(h^2(Y_t^\varepsilon) \mid S_t = S) \tag{3}$$

The particular choice to represent the mean reversion κ and the volatility of volatility of the dynamic ν as $-\kappa/\varepsilon$ and $\nu/\sqrt{\varepsilon}$, respectively, highlights the fact that these parameters are large in practice, in order to fit the anticipated breakeven values. At this stage we note that κ^2 is homogeneous to ν. This choice is dictated by the fact that κ has an inverse time dimension, whereas ν has an inverse of the square root of time dimension.

Statement of the main result. The main contributions of the paper are, first, to provide a methodology that combines perturbation techniques with calibration and, second, to suggest a computation strategy based on exotic Greeks that performs well both theoretically and numerically.

The paper shows a detailed application to the stochastic volatility model. To efficiently compute the LSV impact $\pi_{\mathrm{LSV}} - \pi_{\mathrm{LV}}$, we derive the following formula:

$$\pi_{\mathrm{LSV}} \approx \pi_{\mathrm{LV}} + \tfrac{1}{2}\sigma_y^2 \frac{\partial_{\mathrm{E}}^2 \pi_{\mathrm{LV}}^\beta}{\partial \beta^2}\bigg|_{\beta=0} + \frac{\rho\nu}{\kappa}\frac{\partial_{\mathrm{E}}^2 \pi_{\mathrm{LV}}}{\partial \ln S_0 \partial\sigma} \tag{4}$$

where:

- π_{LSV} is the price under the LSV process (1), which can be decomposed as follows:

$$\left.\begin{aligned} \frac{\mathrm{d}S_t}{S_t} &= \sigma_{\mathrm{D}}(t, S_t)\frac{\mathrm{e}^{Y_t^\varepsilon}}{\sqrt{\mathbb{E}(\mathrm{e}^{2Y_t^\varepsilon} \mid S_t)}}\,\mathrm{d}B_t \\ \mathrm{d}Y_t^\varepsilon &= -\frac{\kappa}{\varepsilon}Y_t^\varepsilon\,\mathrm{d}t + \frac{\nu}{\sqrt{\varepsilon}}\,\mathrm{d}B_t \end{aligned}\right\} \tag{5}$$

- $\sigma_y^2 = \nu^2/2\kappa$ represents the variance of the invariant distribution of Y_t^ε when $\varepsilon \to 0$;
- π_{LV}^β is defined as pricing under the LV process, which depends on the initial conditions β:

$$\frac{\mathrm{d}S_t^\beta}{S_t^\beta} = \sigma_{\mathrm{D}}(t, S_t^\beta)\mathrm{e}^{\beta-\sigma_y^2}\,\mathrm{d}W_t \tag{6}$$

- $\partial_{\mathrm{E}}\cdot$ denotes a particular exotic Greek.

We recall that any exotic Greek $\partial_{\mathrm{E}}\cdot$ of a payoff π is formed from the standard Greek adjusted from the vanilla contribution. More precisely, if we define:

$$q_{KT} = \frac{\partial_{\sigma(K,T)}\pi}{\partial_{\sigma(K,T)}C_{KT}} \tag{7}$$

q_{KT} represents the quantity of vanilla C_{KT} to be retained in order to hedge volatility surface movement:

$$\partial_{\mathrm{E}}(\pi) = \partial(\pi) - \int_0^T \int_0^\infty q_{KT}\,\partial(C_{KT})\,\mathrm{d}T\,\mathrm{d}K$$

We also show below how to compute the exotic Greek through a well-chosen scenario at a very low computational cost. We conclude by showing numerical results on the most-traded instruments in equity derivatives, ie, autocalls.

PnL explain

The most important feature of a model is its ability to explain the PnL evolution on a daily basis. For this exercise, the model is based on the following three pillars:

- **Option:** P, which denotes the exposure.
- **Market:** $(\mathrm{d}S/S)^2$, $\mathrm{d}\Sigma\,\mathrm{d}S$, which denote the breakeven values of the stock volatility.
- **Model:** $\sigma_{\mathrm{D}}, \alpha, \kappa, \nu$, etc, which detail the intrinsic property of the model.

A delta hedged position under the Black-Scholes model gives the following PnL explanatory formula:

$$\delta\mathrm{PL} = \tfrac{1}{2}S^2\frac{\mathrm{d}^2P}{\mathrm{d}S^2}\left[\left(\frac{\delta S^2}{S^2}\right)^{\mathrm{rlzd}} - \sigma^2\delta t\right] \tag{Gamma}$$

The fair price is obtained by putting the model parameter σ to its corresponding realised value $(\delta S^2/S^2)^{\mathrm{rlzd}}$, or at least its anticipated level.

1 Autocall Greeks

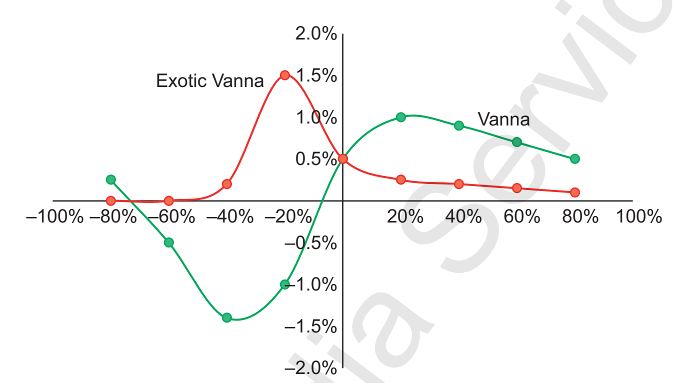

Likewise, if we use an advanced model such as an LV or LSV model and perform only a delta hedge strategy, we obtain the following PnL explanatory formula:

$$\delta\mathrm{PL} = \frac{\mathrm{d}P}{\mathrm{d}\sigma_{KT}}[(\delta\sigma_{KT})^{\mathrm{rlzd}} - (\delta\sigma_{KT})^{\mathrm{model}}] \tag{Vega}$$

$$+ \tfrac{1}{2}S^2\frac{\mathrm{d}^2P}{\mathrm{d}S^2}\left[\left(\frac{\delta S^2}{S^2}\right)^{\mathrm{rlzd}} - \sigma^2\delta t\right] \tag{Gamma}$$

$$+ S\sigma_{KT}\frac{\mathrm{d}^2P}{\mathrm{d}S\,\mathrm{d}\sigma_{KT}} \times \left[\left(\frac{\delta S\delta\sigma_{KT}}{S\sigma_{KT}}\right)^{\mathrm{rlzd}} - \left(\frac{\delta S\delta\sigma_{KT}}{S\sigma_{KT}}\right)^{\mathrm{model}}\right] \tag{Vanna}$$

$$+ \tfrac{1}{2}\sigma_{KT}^2\frac{\mathrm{d}^2P}{\mathrm{d}\sigma_{KT}^2}\left[\left(\frac{\delta\sigma_{KT}^2}{\sigma_{KT}^2}\right)^{\mathrm{rlzd}} - \left(\frac{\delta\sigma_{KT}^2}{\sigma_{KT}^2}\right)^{\mathrm{model}}\right] \tag{Volga}$$

Before tackling the fair value pricing through matching the breakeven value, we need to cancel the Vega risk, which is of first-order magnitude and is in practice 20 times more important than second-order terms. By doing so, the PnL explanatory formula changes not only with the first-order term, which disappears, but also with exposures, which are now adjusted by Vega hedging. The new exposures are called exotic exposures as they are nil for vanilla options or any product that is replicable with vanillas.

More precisely, to cancel Vega risk, the trader needs to sell q_{KT} vanilla options C_{KT} (see (7)).

The trader's new PnL equation is then given by:

$$P^{\mathrm{H}} = P - q_{KT}C_{KT}$$

$$\delta\mathrm{PL}^{\mathrm{H}} = \tfrac{1}{2}S^2\frac{\mathrm{d}^2P^{\mathrm{H}}}{\mathrm{d}S^2}\left[\left(\frac{\delta S^2}{S^2}\right)^{\mathrm{rlzd}} - \sigma^2\delta t\right] \tag{Exotic Gamma}$$

$$+ S\sigma_{KT}\frac{\mathrm{d}^2P^{\mathrm{H}}}{\mathrm{d}S\,\mathrm{d}\sigma_{KT}} \times \left[\left(\frac{\delta S\delta\sigma_{KT}}{S\sigma_{KT}}\right)^{\mathrm{rlzd}} - \left(\frac{\delta S\delta\sigma_{KT}}{S\sigma_{KT}}\right)^{\mathrm{model}}\right] \tag{Exotic Vanna}$$

$$+ \tfrac{1}{2}\sigma_{KT}^2\frac{\mathrm{d}^2P^{\mathrm{H}}}{\mathrm{d}\sigma_{KT}^2} \times \left[\left(\frac{\delta\sigma_{KT}^2}{\sigma_{KT}^2}\right)^{\mathrm{rlzd}} - \left(\frac{\delta\sigma_{KT}^2}{\sigma_{KT}^2}\right)^{\mathrm{model}}\right] \tag{Exotic Volga}$$

As an illustration, we can compute the exotic Greeks of an autocall product (figure 1).

2 Market breakeven versus LSV1F breakeven

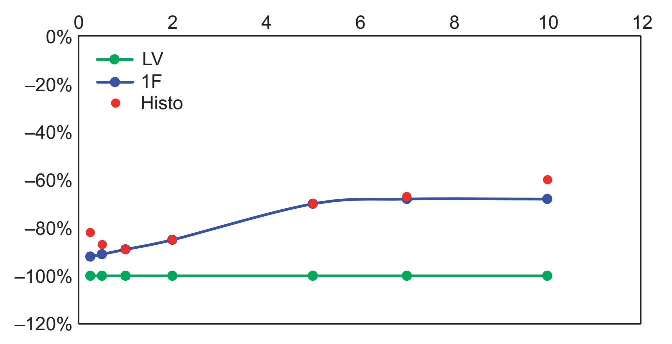

The Exotic Vanna for the autocall product (long position) remains positive regardless of the spot scenarios. This has important consequences for the hedging with the LV model, for which the model breakeven point is exactly -1, and therefore generates a negative carry for the seller of the autocall.

At this stage, we can conclude that hedging an autocall with the LV model will generate a systematic loss on the Exotic Vanna term. A way to compensate for this and retain a fair pricing is to move to a local stochastic volatility one-factor (LSV1F) model that fits the correlation breakeven value (figure 2). This is done via extreme values of stochastic volatility parameters: $\kappa, \nu \gg 1$ and ρ close to -1.

Price-adjustment technique

We start with the following modified Newton lemma.

Lemma 1 (Modified Newton) *Suppose we have a function f of a variable x. Suppose that $x^\star$ is set such that $g(x^\star) = 0$, ie, g satisfies a given constraint. The value of f on $x^\star$ is given by the following formula:*

$$f(x^\star) = f(x) - \frac{\partial_x f(x)}{\partial_x g(x)} g(x) + O(x^\star - x)^2$$

Typically, this calibration can be denoted by $g(x) := \pi_{\mathrm{Model}}(x) - \pi_{\mathrm{Market}}$. In many dimensions, the adjustment takes the following form:

$$f(\pi, \boldsymbol{\beta}^\star) = f(\pi, \boldsymbol{\beta}) - \nabla_{\boldsymbol{\beta}} f_{n\times 1} \cdot [\partial_i g_j]^{-1}_{n\times n} \cdot [g_i(\boldsymbol{\beta})]_{n\times 1} + O(\|\boldsymbol{\beta}^\star - \boldsymbol{\beta}\|^2) \qquad (8)$$

Proof of lemma 1 First we follow the lines of Newton's approach by searching for $x^\star$ as a perturbation of x, ie, $x^\star = x + \varepsilon$. As $g(x^\star) = 0$ we can expand as follows:

$$g(x + \varepsilon) = g(x) + \varepsilon \partial_x g(x) = 0$$

Then, $\varepsilon = -g(x)/\partial_x g(x)$. Now, we expand $f(x^\star)$ to obtain the final result. □

This result is the basis of price adjustment in order to fit a given set of constraints. Fitting is another word for calibration. Indeed, the constraint $g(x) = 0$ is usually written in finance as follows:

$$g(x) := \pi_{\mathrm{Model}}(x) - \pi_{\mathrm{Market}}$$

where x plays the role of model parameters.

We can adjust our LV model price, in order to match vanilla prices exactly in more classical financial notation using the following vector formula:

$$\pi_{\mathrm{LV}} \leftarrow \pi_{\mathrm{LV}} - \int_0^T \int_0^\infty q_{KT} (C^{K,T}_{\mathrm{LV}} - C^{K,T}_{\mathrm{Market}}) \,\mathrm{d}T \,\mathrm{d}K$$

where q_{KT} is defined in (7).

The objective of this paper is to combine this idea with the design of a well-chosen scenario that will permit the precise computation of the LSV impact at a very low computational cost, by interpreting the adjustment as an exotic Greek.

Singular exotic perturbation

In this section, we present the method for a singular exotic perturbation. The objective is to solve the pricing dynamic as a function of ε and see how it converges when $\varepsilon \to 0$ for (1).

We denote by f the conditional expectation in such a way that we fit the vanilla. It satisfies (3).

Making the expansion in the presence of f_ε is hard. Instead, we use the following method, which we name the 'singular exotic perturbation':

- We perform the singular perturbation without calibrating the vanilla, ie, $\lim_{\varepsilon\to 0} f^2_\varepsilon(t, S) = \mathbb{E}(h^2(Y^{\varepsilon\to 0}_t))$. The zeroth order gives the Dupire local volatility model. We identify the higher orders as Volga and Vanna contributions.
- We apply the modified Newton lemma in order to adjust the expansion and recover an exact calibration.
- We explicitly compute zeroth-, first- and second-order order adjustments due to the singular perturbation and correct them in order to maintain the vanilla fit.
- Finally, we design well-chosen scenarios in order to simplify the above computations.

Singular perturbation on the non-calibrated process. The non-calibrated dynamics have the following form:

$$\frac{\mathrm{d}S_t}{S_t} = \sigma_{\mathrm{D}}(t, S_t) \mathrm{e}^{Y^\varepsilon_t - \sigma^2_y} \,\mathrm{d}W_t \qquad (9)$$

Let $u(t, x, y)$ be the price of the derivative. It satisfies the following partial differential equation (PDE):

$$u_t + \tfrac{1}{2}\sigma^2_{\mathrm{D}}(t, x) x^2 \mathrm{e}^{2(y-\sigma^2_y)} u_{xx} + \frac{1}{\varepsilon}\mathcal{L}_y u + \frac{1}{\sqrt{\varepsilon}} \rho x \sigma_{\mathrm{D}}(t, x) \mathrm{e}^{y-\sigma^2_y} \nu u_{xy} = 0$$

where $\mathcal{L}_y = -\kappa u_y + \frac{1}{2}\nu^2 u_{yy}$. We search for $u = u_0 + \sqrt{\varepsilon} u_1 + \varepsilon u_2$ as follows:

$$O\Big(\frac{1}{\varepsilon}\Big): \quad \mathcal{L}_y u_0 = 0, \quad u_0(t, x, \cancel{y})$$

$$O\Big(\frac{1}{\sqrt{\varepsilon}}\Big): \quad \mathcal{L}_y u_1 + \rho\nu x \sigma_{\mathrm{D}}(t, x) \mathrm{e}^{y-\sigma^2_y} \cancel{(u_0)_{xy}} = 0, \quad u_1(t, x, \cancel{y})$$

$$O(1): \quad \mathcal{L}_y u_2 + \rho\nu x \sigma_{\mathrm{D}}(t, x) \mathrm{e}^{y-\sigma^2_y} \cancel{(u_1)_{xy}} + (\partial_t + \tfrac{1}{2}\sigma^2_{\mathrm{D}}(t, x) x^2 \mathrm{e}^{2(y-\sigma^2_y)} \partial_{xx}) u_0 = 0$$

$$O(\sqrt{\varepsilon}): \quad \mathcal{L}_y u_3 + \rho\nu x \sigma_{\mathrm{D}}(t, x) \mathrm{e}^{y-\sigma^2_y} (u_2)_{xy} + (\partial_t + \tfrac{1}{2}\sigma^2_{\mathrm{D}}(t, x) x^2 \mathrm{e}^{2(y-\sigma^2_y)} \partial_{xx}) u_1 = 0$$

We then apply the Poisson formula in u_2 (centring condition):

$$(\partial_t + \tfrac{1}{2}\sigma_{\mathrm{D}}^2(t,x)x^2\langle \mathrm{e}^{2(y-\sigma_y^2)}\rangle\partial_{xx})u_0 = 0$$

where $\langle\cdot\rangle$ denotes the integration over the invariant distribution of Y.

This becomes:

$$\mathcal{L}_{\mathrm{D}}u_0 = 0$$

where L_{D} is the Dupire operator:

$$\mathcal{L}_{\mathrm{D}} = (\partial_t + \tfrac{1}{2}\sigma_{\mathrm{D}}^2(t,x)x^2\partial_{xx})$$

We apply the Poisson condition on u_3:

$$\mathcal{L}_{\mathrm{D}}u_1 + \rho\nu x\sigma_{\mathrm{D}}(t,x)\langle \mathrm{e}^{y-\sigma_y^2}\partial_{xy}u_2\rangle = 0$$

But:

$$\mathcal{L}_y u_2 = -\tfrac{1}{2}x^2\sigma_{\mathrm{D}}^2(t,x)(\mathrm{e}^{2y-\sigma_y^2} - \langle \mathrm{e}^{2(y-\sigma_y)^2}\rangle)\partial_{xx}u_0$$

with $\mathcal{L}_y\phi = \mathrm{e}^{2(y-\sigma_y^2)}$. This implies:

$$\begin{aligned} u_2 &= -\tfrac{1}{2}x^2\sigma_{\mathrm{D}}^2(t,x)\phi(y)\partial_{xx}u_0 + C(t,x,y) \\ \mathcal{L}_{\mathrm{D}}u_1 &+ \rho\nu x\sigma_{\mathrm{D}}(t,x)\partial_x(-\tfrac{1}{2}x^2\sigma_{\mathrm{D}}(t,x)\partial_{xx}u_0) = 0 \end{aligned}$$

with $u_1(T,x) = 0$, $V_3 = \langle \mathrm{e}^{(y-\sigma_y)}\phi'(y)\rangle(-\tfrac{1}{2}(\rho\nu/\kappa))$ (cf. Fouque *et al* 2011).

Therefore, by applying the Feynman-Kac formula, the first-order equation above can be computed using only the local volatility model and its derivatives:

$$\begin{aligned} u_1 &= V_3\mathbb{E}\int_0^T S_t\sigma_{\mathrm{D}}(t,S_t)\partial_x(S_t^2\sigma_{\mathrm{D}}^2(t,S_t)\partial_{xx}u_0)\,\mathrm{d}t \\ &= -2V_3\mathbb{E}\int_0^T S_t\sigma_{\mathrm{D}}(t,S_t)\partial_x(\partial_{\sigma_{\mathrm{D}}(t,x)}u_0)\,\mathrm{d}t \\ &= -2V_3\mathbb{E}\int_0^T\int_0^T\int_0^\infty S_t\sigma(t,S_t) \\ &\qquad \times\partial_x(\partial_{\sigma_{KT}}u_0\partial_{\sigma_{\mathrm{D}}(t,x)}\sigma(K,T))\,\mathrm{d}K\,\mathrm{d}T\,\mathrm{d}t \\ &= -2V_3\mathbb{E}\int_0^T\int_0^T\int_0^\infty S_t\sigma_{\mathrm{D}}(t,S_t)\partial_{\sigma_{\mathrm{D}}(t,x)} \\ &\qquad \times\sigma(K,T)\partial_x\partial_{\sigma_{KT}}u_0\,\mathrm{d}K\,\mathrm{d}T\,\mathrm{d}t \qquad \text{(I)} \\ &\quad - 2V_3\mathbb{E}\int_0^T\int_0^T\int_0^\infty S_t\sigma_{\mathrm{D}}(t,S_t)\partial_{\sigma_{KT}}u_0 \\ &\qquad \times\partial_x(\partial_{\sigma_{\mathrm{D}}(t,x)}\sigma(K,T))\,\mathrm{d}K\,\mathrm{d}T\,\mathrm{d}t \qquad \text{(II)} \end{aligned}$$

The first term in the last equality, (I), shows the full Vanna of the product, summing up all contributions for a comovement of $S_t\sigma_{\mathrm{D}}(t,S_t)$. The second term, (II), has a contribution coming only from vanilla options weighted by the Vega of the product. Therefore, when the product is Vega KT hedged, the first term, which is a Vanna, becomes an Exotic Vanna, and the second term, which is a combination of European contributions, just vanishes. More precisely:[1]

$$\begin{aligned} \text{(I)} &= \mathbb{E}\int_0^T\int_0^T\int_0^\infty \sigma_{\mathrm{D}}(t,S_t)\partial_{\sigma_{\mathrm{D}}(t,x)} \\ &\qquad \times\sigma(K,T)\partial_{\ln x}\partial_{\sigma_{KT}}u_0\,\mathrm{d}K\,\mathrm{d}T\,\mathrm{d}t \\ &= \int_0^T\int_0^\infty \partial_{\ln x}\partial_{\sigma_{KT}}u_0\,\mathrm{d}K\,\mathrm{d}T \\ &= \partial^2_{\ln x,\sigma}u_0 \\ \text{(II)} &= \mathbb{E}\int_0^T\int_0^T\int_0^\infty \sigma_{\mathrm{D}}(t,S_t)\partial_{\sigma_{KT}}u_0 \\ &\qquad \times\partial_{\ln x}(\partial_{\sigma_{\mathrm{D}}(t,x)}\sigma(K,T))\,\mathrm{d}K\,\mathrm{d}T\,\mathrm{d}t \\ &= \int_0^T\int_0^\infty q_{KT}\partial_{\sigma_{KT}}C_{KT}\partial_{\ln x}\sigma(K,T)\,\mathrm{d}K\,\mathrm{d}T \end{aligned}$$

Probabilistic interpretation. In the computation of the second-order impact, the stochastic volatility is important. We start as if there is no recalibration:

$$\frac{\mathrm{d}S_t}{S_t} = \sigma_{\mathrm{D}}(t,S_t)\mathrm{e}^{Y-\sigma_y^2}\,\mathrm{d}W_t \tag{10}$$

where Y is the invariant distribution of the previous process. Let π_{LV}^y be as defined in (9). Note that:[2]

$$\pi_{\mathrm{LV}}^{\sigma_y^2} = \pi_{\mathrm{LV}}$$

Then:

$$\begin{aligned} u_0 + u_2 &\approx \pi_{\mathrm{LV}}^{\sigma_y^2} + \tfrac{1}{2}\operatorname{Var}(Y)\frac{\pi_{\mathrm{LV}}^{+\beta+\sigma_y^2} + \pi_{\mathrm{LV}}^{-\beta+\sigma_y^2} - 2\pi_{\mathrm{LV}}^{+\sigma_y^2}}{\beta^2} \\ &= u_0 + \tfrac{1}{2}\sigma_y^2\frac{\partial^2\pi_{\mathrm{LV}}^\beta}{\partial\beta^2}\bigg|_{\beta=+\sigma_y^2} \end{aligned}$$

Adjusting the prices to recover the vanilla fit. At this stage, we can apply the modified Newton lemma in order to compensate for the non-calibration. Let us denote by P^{NC} the price obtained using $u_0 + u_2$ which is non-calibrated. Let us denote by P the calibrated (to the vanilla) price. We see that this pricing does not match vanilla options due to the extra term:

$$u_2 = \tfrac{1}{2}\sigma_y^2\frac{\partial^2\pi_{\mathrm{LV}}^\beta}{\partial\beta^2}\bigg|_{\beta=+\sigma_y^2}$$

We can construct the adjusted price P by compensating at first order:

$$\begin{aligned} P &= P^{\mathrm{NC}} - \int_0^T\int_0^\infty q_{KT}(P^{\mathrm{NC}}(C_{KT}) - C_{\mathrm{Market}}^{K,T})\,\mathrm{d}T\,\mathrm{d}K \\ &= \pi_{\mathrm{LV}} + \tfrac{1}{2}\sigma_y^2\frac{\partial^2\pi_{\mathrm{LV}}^\beta}{\partial\beta^2} \\ &\quad - \int_0^T\int_0^\infty q_{KT}\bigg(C_{\mathrm{LV}}^{K,T} + \tfrac{1}{2}\sigma_y^2\frac{\partial^2 C_{\mathrm{LV}}^{\beta,K,T}}{\partial\beta^2} - C_{\mathrm{Market}}^{K,T}\bigg)\,\mathrm{d}T\,\mathrm{d}K \\ &= \pi_{\mathrm{LV}} + \tfrac{1}{2}\sigma_y^2\bigg(\frac{\partial^2\pi_{\mathrm{LV}}^\beta}{\partial\beta^2} - \int_0^T\int_0^\infty q_{KT}\frac{\partial^2 C_{\mathrm{LV}}^{\beta,K,T}}{\partial\beta^2}\,\mathrm{d}T\,\mathrm{d}K\bigg) \end{aligned}$$

where we have used $C_{\mathrm{LV}}^{K,T} = \pi_{\mathrm{Market}}^{K,T}$.

[1] *We use the following equation:* $\mathbb{E}\int_0^T \sigma_{\mathrm{D}}(t,S_t)\partial_{\sigma_{\mathrm{D}}(t,x)}\sigma(K,T)\,\mathrm{d}t = 1.0$.

[2] $\mathbb{E}(f(Y)) \approx f(\mathbb{E}(Y)) + \tfrac{1}{2}\operatorname{Var}(Y)f''(\mathbb{E}(Y))$.

3 Annualised Totem prices

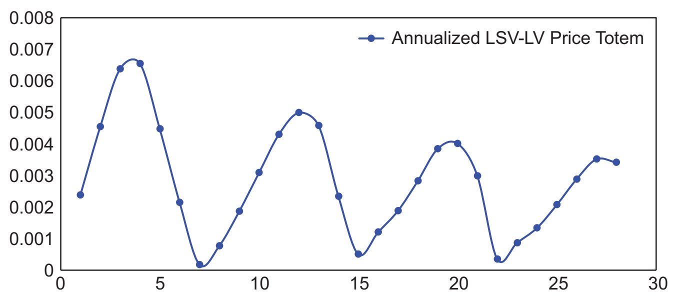

We introduce the 'Exotic Volga' Greek:

$$\frac{\partial^2_{\mathrm{E}}\pi_{\mathrm{LV}}}{\partial\beta^2} = \left(\frac{\partial^2\pi^{\beta}_{\mathrm{LV}}}{\partial\beta^2} - \int_0^T\int_0^\infty q_{KT}\frac{\partial^2 C^{\beta,K,T}_{\mathrm{LV}}}{\partial\beta^2}\,\mathrm{d}T\,\mathrm{d}K\right)$$

The Exotic Volga appears naturally for perturbation Greeks once we adjust for the calibration of the vanilla.

Similarly, we adjust the term I by introducing the effect of the calibration on the vanilla options and calculate I_{C}:

$$\begin{aligned} I_{\mathrm{C}} &= I - \int_0^T\int_0^\infty q_{KT}\partial^2_{\ln x,\sigma_{KT}} C_{KT}\,\mathrm{d}K\,\mathrm{d}T \\ &= \frac{\partial^2_{\mathrm{E}}\pi_{\mathrm{LV}}}{\partial\ln S_0\partial\sigma} \end{aligned}$$

We obtain the final result described in (4).

We have obtained the desired result. In particular we have identified the functions that intervene in the expansion, and their exotic nature yields an important property of the formula: it ensures that the vanilla has exactly zero impact on the LSV. This formula of great interest for understanding the effect of stochastic volatility on top of the local volatility. However, a brute force implementation will be needed to provide all the q_{KT}. We can, for example, use an offline computation, as proposed in Hagan (2005).

We could also use algorithmic automatic differentiation (AAD) to compute all the q_{KT} at a cost that does not exceed four price computations. We instead propose the design of particular scenarios in order to compute these exotic Greeks based on two building blocks: the implied Black-Scholes calculator and local volatility pricer. We approach the problem as a computation of the Exotic Theta as presented in Bergomi (2018).

Designing exotic scenarios

In order to compute an exotic version of a sensitivity we:

- compute the sensitivity using a classical scenarios approach; and
- adjust the scenarios in such a way that the effect on vanilla options disappears.

We use the perturbed local volatility (9).

We therefore obtain an exact definition of the asymmetric bump that makes the vanilla Volga disappear:

$$\sigma_{KT} + x_{KT} = C^{-1}_{KT}\{C^{K,T}_{\mathrm{LV}} + (C^{K,T}_{\mathrm{LV}} - C^{-\beta,K,T}_{\mathrm{LV}})\} \quad (11)$$

4 Exotic Vanna (eight-year tenor) for the different barrier levels

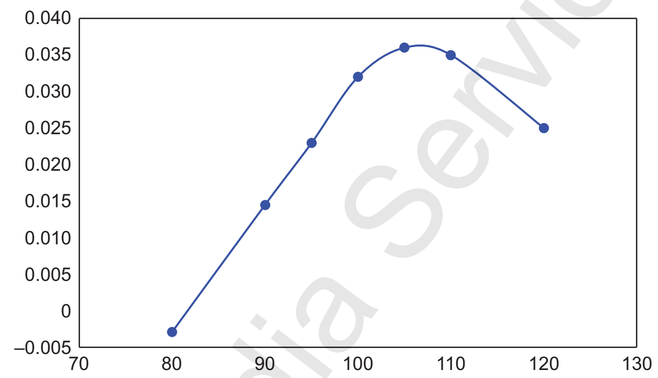

The 'Exotic Volga' Greek is then computed as follows:[3]

$$\frac{\partial^2_{\mathrm{E}}\pi_{\mathrm{LV}}}{\partial\beta^2} \overset{\beta\to 0}{=} \frac{\pi_{\mathrm{LV}}(\sigma + x) - 2\pi_{\mathrm{LV}} + \pi^{-\beta}_{\mathrm{LV}}}{\beta^2}$$

This Greek is the result of three local volatility prices. One of them has already been computed as it is the central price with no deformation of the volatility surface. The two others are computed by first generating a scenario of volatility deformation using a bump with a value of β and then creating an implied volatility bump $x(K,T)$ constructed point by point by inferring the volatility from (11). This construction guarantees that the exotic term is mechanically zero on vanilla options. This property gives it its name: exotic.

It is non-zero if and only if the product is a non-vanilla option, ie, not replicable using a vanilla option.

At this stage we have shown an analytic formula for the LSV impact and given details of the low-complexity scenarios that permit the computation of the exotic in order to implement the formula.

In the next section we shall apply the previous formula to autocalls and show how precise it is for these types of product.

Numerical examples

Twenty-eight different structures are exchanged via the Totem service from IHS Markit:

- Seven barrier levels: 80%, 90%, 95%, 100%, 105%, 110%, 120%.
- Four maturities: $T_1 = 1$ years, $T_2 = 3$ years, $T_3 = 5$ years, $T_4 = 8$ years.
- Quarterly coupon equal to 1.25%.

Totem provides the running cost of these $28 = 4 \times 7$ structures (figure 3):

$$\pi_{\mathrm{Totem}}(B,T) = \frac{\pi_{\mathrm{LSV}}(B,T) - \pi_{\mathrm{LV}}(B,T)}{T}$$

We compute the exotic Greeks for different payoffs and then simulate hundreds of random stochastic volatility model parameters (ρ, ν, κ). We compare

[3] *Note that $\pi^{-\beta,K,T}_{\mathrm{LV}}$ is the price of vanilla options obtained with the local vol process. This is done using a forward PDE sweep.*

5 Comparing a full LSV with the approximation

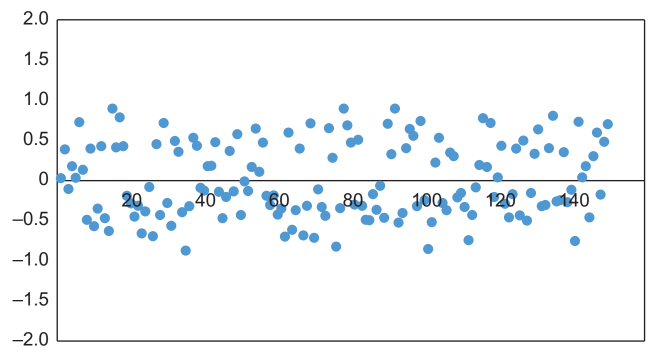

the full LSV impact computed using a full implementation with the formula for the different barrier levels. We show the results in figures 4 and 5.

These exotic Greeks depend only on the products and the volatility surface. They do not depend on the stochastic volatility parameters.

We then use our formula for randomly simulated stochastic volatility parameters and compare the full LSV with the proxy formula. Needless to say, the formula is instantaneous, whereas the full LSV takes a non-negligible time to run. We note that the results are extremely good with no bias (average error 0 basis points) and a standard deviation of around 2bp to be compared with the 200bp we ought to match.

Conclusion

In this paper we have introduced a new methodology: singular exotic perturbation. This is an efficient approach to compute the impact of smile dynamics without running a costly LSV model. It instead builds on a well-chosen scenario priced completely under the simpler LV model. Our proposed formula ensures zero impact on vanilla and performs very well on more complex products such as autocalls. The methodology proposed can be used in different contexts, as it combines three different building blocks: singular perturbation, first-order price adjustment and computability through the introduction of exotic Greeks. We suggest that this methodology be used in other cases such as those involving stochastic rates, multiple assets or correlation skew. Further work is needed to detail these practical applications. ■

Adil Reghai is head of the quantitative research team for equities, commodities and hybrid business at Natixis Investment Bank's Groupe BPCE. Florian Monciaud is a quantitative analyst in the same team. Both are based in Paris.
Email: florian.monciaud@gmail.com, adil.reghai@polytechnique.org.

REFERENCES

Bergomi L, 2018
Exotic theta
Conference presentation at `Quant minds'

Bergomi L, 2016
Stochastic Volatility Modelling
Wiley

Lipton A, 2002
The vol smile problem
Risk February, pages 81–85, www.risk.net/1530435

Lipton A, A Gal and A Lasis, 2014
Pricing of vanilla and first-generation exotic options in the local stochastic volatility framework: survey and new results
Quantitative Finance 14(11), pages 1,899–1,922

Fouque JP, G Papanicolaou, R Sircar and K Sølna, 2011
Multiscale Stochastic Volatility for Equity, Interest Rate, and Credit Derivatives
Cambridge University Press

Hagan PS, 2005
Adjusters: turning good prices into great prices
Best of Wilmott , page 297

Hagan PS, A Lesniewski and D Woodward, 2018
Managing vol surfaces
Wilmott 2018(93), pages 24–43

Henry-Labordère P, 2009
Calibration of local stochastic volatility models to market smiles: a Monte-Carlo approach
Risk September, pages 112–117

Guyon J and P Henry-Labordère, 2011
The smile calibration problem solved
SSRN preprint, https://ssrn.com/abstract=1885032

Reghai A and O Kettani, 2020
Financial Models in Production
Springer

Reghai A, O Kettani, V Klaeyle and A Boukhaffa, 2012
LSV models with a mixing weight
SSRN preprint, https://ssrn.com/abstract=2008207

Reghai A, 2015
Quantitative Finance: Back to Basics
Palgrave MacMillan