Uwe Wystup

Mixed Local Volatility Model Boosts Distribution of Exotics

Speed, flexible calibration, and market acceptance make MLVs a necessary tool in the FX quant's armory.

As a market maker for FX derivatives, and especially for exotic products on a single desk platform, one needs the best tradeoff between precision and speed in one's exotics model. A Mixed Local Volatility (MLV) model is a simplified, yet powerful version of the fully Stochastic Local Volatility (SLV) model. It ignores the correlation between spot and volatility (which is common in FX) [illegible]

Figure 1: Volatility processes in SLV and MLV

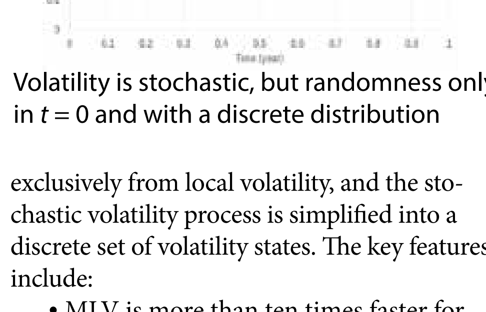

Volatility driven by a diffusive process (e.g., Heston) with continuous distribution

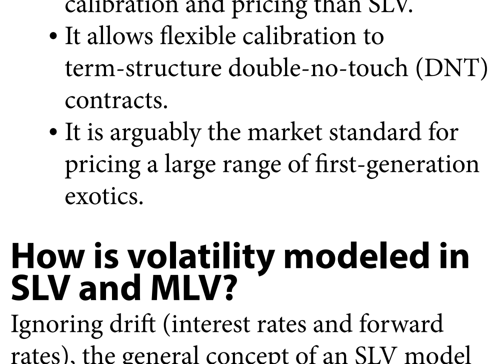

Volatility is stochastic, but randomness only in $t = 0$ and with a discrete distribution

exclusively from local volatility, and the stochastic volatility process and volatility into a discrete set of volatility states. The key features include:

- MLV is more than ten times faster for calibration and pricing than SLV.
- It allows flexible calibration to [illegible] double-no-touch (DNT) contracts.
- It is arguably the market standard for pricing a large range of first-generation exotics.

How is volatility modeled in SLV and MLV?

Ignoring drift (interest rates and forward rates), the general dynamics of an SLV model is using a product of a local and a stochastic volatility model, stochastic for the model dynamics and to reflect nondeterministic volatility, local to fit the model to a given vanilla options volatility surface:

[illegible]

The choice for the stochastic volatility is typically a diffusion process in SLV models, indicated by infinitely many paths on the left-hand side in Figure 1. For MLV, one can think of tossing a coin and generating a few volatility states [illegible], see right-hand side in Figure 1.

How do traders mark an MLV model?

As a trader, one either buys or distributes FX volatility market data in the form of At-The-Money (ATM) volatilities, Risk Reversals (RR), and Butterflies (BF) for 25 delta and 10 delta volatilities for the usual tenors, as listed in Table 1.

A trader's job is then to mark the term structure of mixing factors empirically – for example, the 1M mixing factor of 45 percent means that [illegible]

Table 1: Volatility market data and MLV mixing factors (MIX)

	ATM	RR25	RR10	BF25	BF10	MIX
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
3M	9.2%	−0.80%	0.40%	−1.52%	1.44%	50.00%
6M	9.5%	−0.80%	0.40%	−1.52%	1.44%	50.00%
9M	10.0%	−0.80%	0.40%	−1.52%	1.44%	55.00%
1Y	11.0%	−0.80%	0.40%	−1.52%	1.44%	55.00%
18M	11.5%	−0.80%	0.40%	−1.52%	1.44%	55.00%
2Y	12.0%	−0.80%	0.40%	−1.52%	1.44%	55.00%

45 percent of BF25 is generated by stochastic volatility (or mixture)
55 percent of BF25 is generated by the local volatility

Traders mark MIX to match a set of one-touch (OT) or DNT contracts. Figure 2 shows the difference of MLV-based OT prices and their corresponding theoretical value (TV) as a function of the TV, an illustration generally referred to as the OT-moustache. Different mixing factors (MIX) of 0 percent, 25 percent, 50 percent, 75 percent, and 100 percent generate different price differences. The trader chooses the MIX that would represent the market prices of OT contracts. Notice that the interbank bid-offer spread is about 2 percent in USDJPY, which is a one-unit block on the y-axis; thus, the choice of MIX is crucially important when running an automatic trading platform.

If prices of first-generation exotics are not readily available, a statistical estimate of MIX can also be obtained by looking at the historical correlation between spot and RR25.

Figure 2: One-touch moustache in the mixed local volatility model with various mixing factors

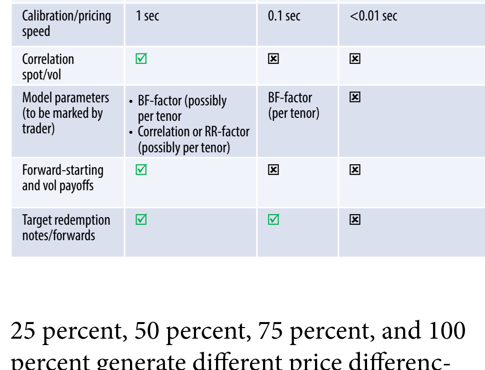

Table 2: Features of SLV, MLV, and Vanna–Volga models

	SLV	MLV	Vanna–Volga (Skew)
Accurate vanilla calibration	☑	☑	⚠Perfect calibration only at 3 strikes
Barrier/touch options, 1st generation exotics	☑ (calibration at short maturity)	☑	☑ (⚠ little flexibility in adjusting exotic prices)
Calibration/pricing speed	1 sec	0.1 sec	<0.01 sec
Correlation spot/vol	☑	☒	☒
Model parameters (to be marked by trader)	• BF-factor (possibly per tenor) • Correlation or RR-factor (possibly per tenor)	BF-factor (per tenor)	☒
Forward-starting and vol payoffs	☑	☒	☒
Target redemption notes/forwards	☑	☑	☒

Figure 3: Marking SLV and MLV to market prices of OT contracts

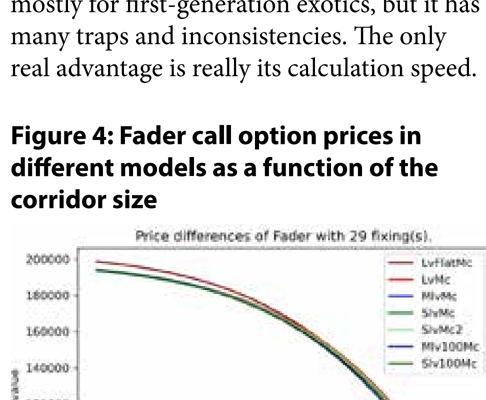

Comparison of SLV/MLV/Vanna–Volga

I summarize the key features of the three most common industry exotics models in Table 2. The Vanna–Volga approach, which was popular until the first decade, works mostly for first-generation exotics, but it has many traps and inconsistencies. The only real advantage is really its calculation speed.

Figure 4: Fader call option prices in different models as a function of the corridor size

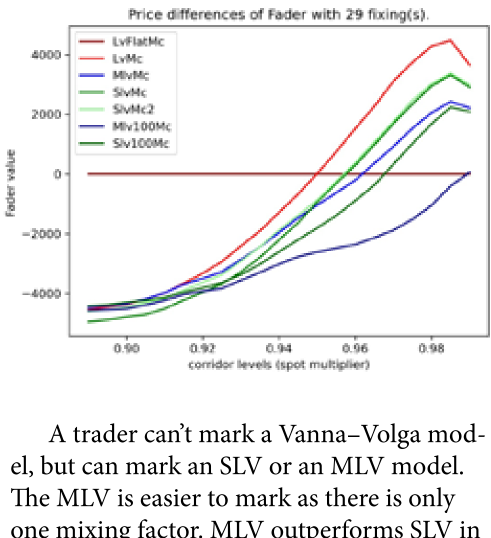

Figure 5: Fader call option price differences in different models as a function of the corridor size

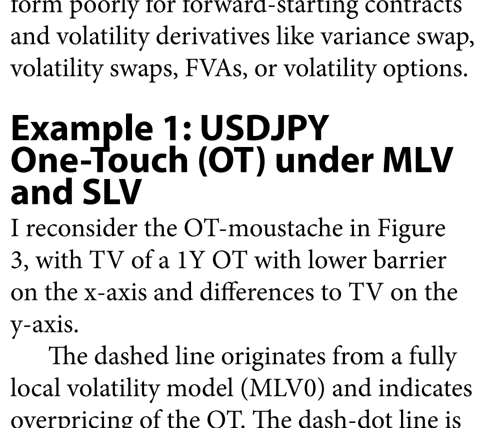

A trader can't mark a Vanna–Volga model, but can mark an SLV or an MLV model. The MLV is easier to mark as there is only one mixing factor. MLV outperforms SLV in calibration and pricing speed, but would perform poorly for forward-starting contracts and volatility derivatives like variance swap, volatility swaps, FVAs, or volatility options.

Example 1: USDJPY One-Touch (OT) under MLV and SLV

I reconsider the OT-moustache in Figure 3, with TV of a 1Y OT with lower barrier on the x-axis and differences to TV on the y-axis.

The dashed line originates from a fully local volatility model (MLV0) and indicates overpricing of the OT. The dash-dot line is a fully stochastic volatility model (MLV100) and indicates underpricing of the OT. The black curve represents market prices. A suitable fit of an SLV model in red can be obtained by an SLV mixing factor of 70 percent (SLV_OT). An equally suitable fit of the MLV model can be obtained by a mixing factor of 49 percent (MLV49.0). Note that the mixing factors generally differ, depending on the model choice. Vanna–Volga (VV), in blue, has a similar pattern but still leads to mispricing of the OT contracts, as it is often outside the market bid-offer.

Example 2: Fader call

Next, I consider a fader call option with a notional proportional to the number of times the FX rate fixes in a predefined corridor, whose lower end is set by the spot multiplier and upper end is calculated symmetrically – e.g., 1.1 for a spot multiplier of 0.9. The prices generated by various SLV and MLV models are exhibited in Figure 4.

Table 3: Model/product matrix

		Vanillas & European
Models	BS + volatility interpolation	✓
	Vanna–Volga	✓
	Dupire	✓
	Local-vol mixture	✓
	Local-vol mixture with transition	✓
	SLV	✓

Since the prices overlap each other, it is better to consider the differences of the prices to TV in Figure 5. We observe that for wide corridors, all the models under consideration yield similar results but for small corridors, the differences are much larger. I conclude that model risk is not only a function of the product, but also of the contractual parameters of the product. Faders with small fade-in corridors will be priced differently in different SLV and MLV models. The way to go forward is to make the fader price at least consistent with first-generation exotics by calibrating an MLV model with the appropriate MIX.

And, of course, the risk managers and model governance officers should be aware of the model risk.

A parsimonious choice of model

An overview of which model I consider suitable for which product class is presented in Table 3. Red colors refer to a too simple model and mispricing, blue indicates a too complex model and overkill, and green the right level of complexity. Yellow indicates borderline cases, e.g., a Vanna–Volga approach does not even work precisely for vanilla options because the volatility is only correct for three strikes, but usually wing volatilities are too low. A pure Heston model, while studied intensively in the academic literature, does not really perform satisfactorily for any product.

In summary: MLV is a simplified version of SLV

- The stochastic volatility process is discretized into a small number of volatility states.
- Zero correlation between spot and volatility is assumed (common assumption in FX).
- MLV contains the necessary flexibility to calibrate to a large range of first-generation exotics and most of the target forwards family.
- Calibration and pricing are very fast, suitable for structuring and a flow trading environment.

Acknowledgments

I thank my colleagues Andreas Weber and Frédéric Bossens at MathFinance for their contributions to this column.

FURTHER READING

There are not many papers out there on mixed local volatility models, but some worth reading include the following:

Dupire, B. 1994. Pricing with a smile. *Risk*, 7(7), 18–20.
Ren, Y., Madan, D., Qian Madani, M. 2007. Calibrating and pricing with embedded local volatility models. *Risk Magazine*, September, 138–143.
Clark, I.J. 2011. *Foreign Exchange Option Pricing – A Practitioner's Guide*, Wiley Finance.

About the Author

Uwe Wystup (uwe.wystup@mathfinance.com) is the founder and Managing Director of MathFinance AG, an independent consulting and software company that specializes in FX derivatives pricing.