Joint Dynamics for the Underlying Asset and its Implied Volatility Surface: A New Methodology for Option Risk Management

Pascal François
In collaboration with: Rémi Galarneau-Vincent, Geneviève Gauthier, Frédéric Godin

Seminar at Bayes Business School - Financial Engineering Workshop

May 2025

Derivative markets and the implied volatility surface

- On derivative markets, the IV surface is the price snapshot of all options currently traded.
- It is the cornerstone for option trading, pricing and hedging.
- Option pricing model performance is typically tested on its ability to fit the IV surface.
- Consistent pricing of OTC derivatives using calibration on IV surface.
- The IV surface is also used for the calculation of volatility indices such as the VIX.

Derivative markets and the implied volatility surface

- In the standard option pricing literature, the IV surface is the model *output*.
- We contribute to a nascent literature in which the IV surface is the model *input*.
- A tight fit with current market data is guaranteed.
- Forecasting option prices takes full advantage of all market information.
- More in line with market practice.

Our contribution

- Incorporate a model for underlying asset returns into a parametric and dynamic IV surface. No-arbitrage condition imposed on the dynamics of underlying asset returns.
- One of the variance factors is anchored to the 1-month, at-the-money IV level.
- A Gaussian copula captures the dependence structure between the underlying asset returns and the IV factors $\rightarrow$ JIVR model (Joint Implied Volatility & Return).
- Estimate the JIVR on S&P 500 index return and option data for the 1996-2020 period.
- Challenge the JIVR for producing reliable Value-at-Risk estimates for straddles and accurate forecasts of the VIX distribution.

Related literature

- Direct modelling of the IV surface (Zhu and Avellaneda, 1998, Schönbucher, 1999, Fengler, 2006, Daglish et al., 2007). Diffusion processes for implied volatilities. Hard to derive constraints on the risk-neutral dynamics to prevent arbitrage.
- Carr and Wu (2016) restrict the modelling of the IV surface to near-term dynamics. They suggest using their framework with a parametric specification for the underlying asset return.
- Carr and Wu (2020) limits the diffusion modelling of the IV to the management of the instantaneous P&L of an option position.
- Extract IV surface explanatory factors in a non-parametric fashion (Cont and Da Fonseca, 2002, Israelov and Kelly, 2017, Cont and Vuletic, 2023). Only applies to dense regions of the IV surface, ignoring some peripheral options that are very informative about higher-order moments.

Outline

Model

The implied volatility $\sigma(M,\tau)$ observed on a given day for an option with moneyness M and time to maturity τ is modeled as

$$\sigma(M,\tau) = \underbrace{\beta_1}_{\text{Long-term ATM IV}} + \beta_2 \underbrace{f_2(\tau, T_{\text{conv}})}_{\text{Time to maturity slope}} + \beta_3 \underbrace{f_3(M)}_{\text{Moneyness slope}} + \beta_4 \underbrace{f_4(M^2, \tau, T_{\max})}_{\text{Smile attenuation}} + \beta_5 \underbrace{f_5(M^3, \tau, T_{\max})}_{\text{Smirk}}.$$

$T_{\max} = 5$ years is the maximal maturity represented by the model.
$T_{\text{conv}} = 0.25$ represents the location of a fast convexity change in the IV term structure with respect to time to maturity.
$f_j, (j = 2, ..., 5)$ are well-behaved, differentiable functions.

Factors

Long-term ATM IV

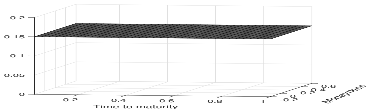

Time to maturity slope

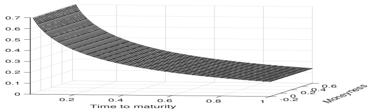

Moneyness slope

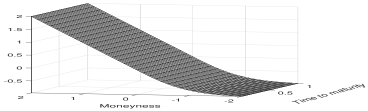

Smile attenuation

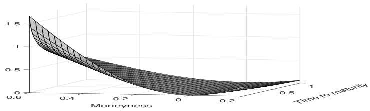

Smirk

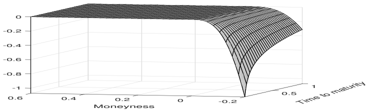

Outline

Estimation and Performance

Our model is a linear function of the factors, which suggests that the ordinary least square (OLS) estimation approach is straightforward.

However, there can be several sets of parameters which produce very similar surfaces.

To preserve the economic interpretation of each factor, the least squares method is coupled with a Bayesian approach for regularization purposes, thereby avoiding avoid erratic behavior in parameter time series.

Daily estimates

Panel A: β_1 (Level) and β_3 (Moneyness slope)

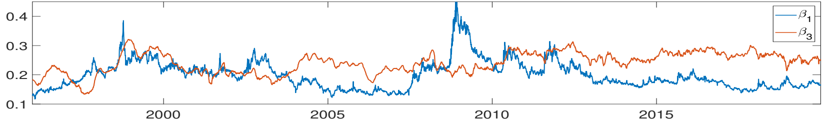

Panel B: β_2 (Time-to-maturity slope)

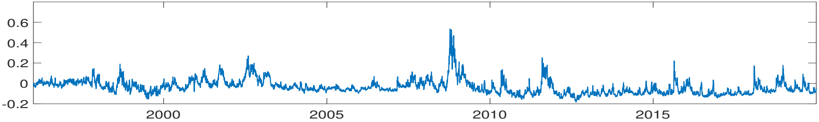

Panel C: β_4 (Smile attenuation) and β_5 (Smirk)

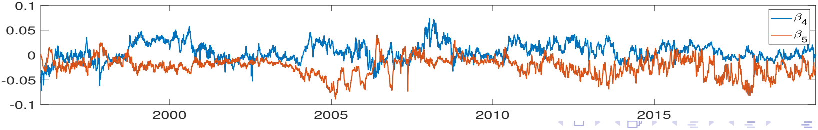

Performance against some benchmarks

Root mean squared IV error

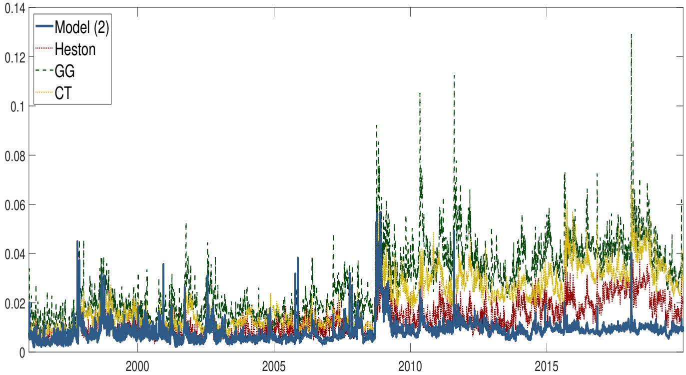

GG: Goncalves and Guidolin (2006); CT: Chalamandaris and Tsekrekos (2011)

Outline

Log-returns dynamics

- The S&P500 price evolution is $S_{t+1} = S_t \exp(Y_{t+1})$.
- Daily sample, $\Delta = \frac{1}{252}$.
- The annualized variance process, h, is anchored around the 1-month ATM IV thereby incorporating the forward-looking information.
- $\{\epsilon_{t,Y}\}$ iid standardized NIG random variables.

$$
\begin{aligned}
\underbrace{Y_{t+1}}_{\text{Log-returns}} &= r_t - q_t + \underbrace{\lambda h_{t+1,Y}\Delta}_{\text{equity risk premium}} - \underbrace{\psi(\sqrt{h_{t+1,Y}\Delta}}_{\text{convexity correction}}) + \sqrt{h_{t+1,Y}\Delta}\,\epsilon_{t+1,Y}\,, \\
h_{t+1,Y} &= V_t + \kappa_Y\left(h_{t,Y} - V_t\right) + a_Y h_{t,Y}\left(\epsilon^2_{t,Y} - 1 - \underbrace{2\gamma_Y\epsilon_{t,Y}}_{\text{leverage effect}}\right), \\
\sqrt{V_t} &= \omega_Y\,\sigma\left(M = 0, \tau = \frac{1}{12}, \beta_t\right).
\end{aligned}
$$

Long-term ATM IV level and other factors

- The process is also anchored around the 1-month ATM IV.
- It is similar to Carr and Wu (2016) where the volvol is proportional to the volatility level.
- $\{\epsilon_{t,1}\}_{t=1}^{T}$ iid standardized NIG random variables.

$$\beta_{t+1,1} = \alpha_1 + \sum_{j=1}^{5} \theta_{1,j}\beta_{t,j} + \sqrt{h_{t+1,1}\Delta}\,\epsilon_{t+1,1},$$

$$h_{t+1,1} = U_t + \kappa_1 \left(h_{t,1} - U_t\right) + a_1 h_{t,1}(\epsilon_{t,1}^2 - 1 - 2\epsilon_{t,1}\gamma_1),$$

$$\sqrt{U_t} = \omega_1\ \sigma\left(0, \frac{1}{12}, \beta_t\right).$$

- The four other daily coefficients $\beta_{t,2}, \ldots, \beta_{t,5}$ are represented by a VAR-NGARCH process with NIG innovations.

Estimation procedure

- JIVR model parameters (excluding the Copula parameters) are estimated via maximum likelihood. The parameters of the marginal processes are estimated in a univariate fashion.
- The Gaussian copula parameters are subsequently estimated over the residuals.

Goodness of fit

	p-value
S&P 500 log-returns	29.4%
Long-term level ($\beta_{t,1}$)	65.3%
TmT slope ($\beta_{t,2}$)	85.2%
Moneyness slope ($\beta_{t,3}$)	58.1%
Smile attenuation ($\beta_{t,4}$)	65.6%
Smirk ($\beta_{t,5}$)	30.8%

Table: Cramér-von Mises goodness-of-fit tests

p-values of the Cramér-von Mises test applied to residuals over the daily sample ranging from January 4, 1996, to December 31, 2020.

Nested sub-models

	BS	Gaussian	Gaussian	NIG	JIVR Model	
	AR(1)	GARCH	NGARCH	NGARCH	Indep.	Copula
Log-likelihood						
S&P 500 log-returns	10,124	11,365	11,466	11,600	11,709	
Long-term level	13,788	15,617	15,620	15,785	15,861	
TmT slope	7,115	8,681	8,843	9,304	9,319	
Moneyness slope	15,392	15,674	15,672	15,709	15,709	
Smile attenuation	14,922	15,251	15,250	15,352	15,352	
Smirk	14,441	14,861	14,860	14,995	14,995	
Joint	75,785	81,452	81,713	82,748	82,947	85,873

Table: Out-of-sample model performance

Outline

Forecasting the option portfolio return distribution

1. Parameter estimation (over the training sample).
   For each iteration N, where $N = 2007, \ldots, 2020$, the model is first estimated over the training sample, which covers years 1996 to $N - 1$.
2. Joint simulation of the log-returns and the IV surface for each day t in year N over a horizon of d days.
3. Option portfolio evaluation
   (using the surface-implied, risk-neutral density).
4. Return empirical distribution.
5. Risk measures.

VaR coverage test; 1-day and 5-day horizon

	TmT	Straddles 1	Straddles 3	Straddles 6	Strangles 1	Strangles 3	Strangles 6
1%	% breaches	1.11%	0.79%	0.77%	1.08%	0.60%	0.57%
	p-value	(53.25%)	(20.30%)	(14.52%)	(64.57%)	(0.91%)	(0.50%)
5%	% breaches	3.43%	2.95%	3.04%	4.77%	3.43%	3.15%
	p-value	(0%)	(0%)	(0%)	(52.06%)	(0%)	(0%)
95%	% breaches	5.84%	4.74%	4.79%	6.13%	5.28%	4.94%
	p-value	(2.50%)	(47.10%)	(57.28%)	(0.30%)	(45.50%)	(86.17%)
99%	% breaches	1.70%	1.16%	1.36%	1.79%	1.50%	1.45%
	p-value	(0.01%)	(34.27%)	(4.08%)	(0%)	(0.52%)	(1.24%)
1%	% breaches	0.71%	0.74%	1.02%	0.79%	0.91%	1.05%
	p-value	(6.72%)	(10.05%)	(89.93%)	(20.30%)	(57.63%)	(76.89%)
5%	% breaches	2.72%	4.11%	4.34%	3.29%	4.28%	4.68%
	p-value	(0%)	(1.28%)	(6.63%)	(0%)	(4.57%)	(37.97%)
95%	% breaches	5.19%	4.40%	4.88%	4.99%	4.60%	5.13%
	p-value	(60.41%)	(9.39%)	(74.16%)	(98.46%)	(26.45%)	(71.47%)
99%	% breaches	1.36%	1.28%	1.65%	1.36%	1.30%	1.62%
	p-value	(4.08%)	(11.35%)	(0.04%)	(4.08%)	(8.22%)	(0.07%)

Outline

The VIX

- The Chicago Board Options Exchange (CBOE) computes the VIX index from a portfolio of available OTM put and call options as explained in detail in CBOE (2014). For a given time-to-maturity τ,

$$\mathrm{VIX}_{t,\tau} = 100\sqrt{\frac{2}{\tau}\sum_{i=1}^{N_{t,\tau}} \frac{\Delta K_{i,\tau}}{K_{i,\tau}^2} e^{r_\tau \tau} P_{t,\tau}(K_{i,\tau}) - \frac{1}{\tau}\left(\frac{F_{t,\tau}}{K_{i,\tau}} - 1\right)^2}$$

- The squared index is a linear interpolation between $\mathrm{VIX}_{t,\tau_1}^2$ and $\mathrm{VIX}_{t,\tau_2}^2$ where $\tau_1 < T < \tau_2$ are the two nearest available time-to-maturities surrounding T:

$$\mathrm{VIX}_{t,T}^{\text{index}} = \sqrt{\frac{\tau_2 - T}{\tau_2 - \tau_1}\mathrm{VIX}_{t,\tau_1}^2 + \frac{T - \tau_1}{\tau_2 - \tau_1}\mathrm{VIX}_{t,\tau_2}^2}$$

The JIVR forecast

- We use the JIVR model to generate forecasts for the VIX variation

$$\Delta \mathrm{VIX}_{t,t+d}^{\mathrm{JIVR}} = \mathrm{VIX}_{t+d}^{\mathrm{JIVR}} - \mathrm{VIX}_{t}^{\mathrm{JIVR}}$$

  with d the prediction horizon (1 or 5 days), $\mathrm{VIX}_{t}^{\mathrm{JIVR}}$ obtained by replacing the OTM option quoted prices by the ones obtained from the fitted IV surface that day, and $\mathrm{VIX}_{t+d}^{\mathrm{JIVR}}$ obtained from Monte Carlo simulations of the IV surface and the underlying asset log-return.

- The forecast is finally given by

$$\mathrm{VIX}_{t+d,T}^{\mathrm{index}} = \mathrm{VIX}_{t,T}^{\mathrm{index}} + \Delta \mathrm{VIX}_{t,t+d}^{\mathrm{JIVR}}$$

VIX distribution forecasting

	JIVR	Direct model	p-values
2014	**-346.58**	-352.33	22.9%
2015	-434.79	**-430.55**	75.8%
2016	**-372.33**	-379.62	10.0%
2017	**-251.82**	-269.93	2.8%
2018	**-436.39**	-445.06	16.5%
2019	**-359.79**	-362.01	24.5%
2020	**-519.20**	-537.27	1.8%
Total	**-2,720.90**	-2,776.77	0.4%

Table: Out-of-sample log-likelihood for VIX distribution forecasting

Direct modeling of the VIX index refers to:

$$\mathrm{VIX}_t = \alpha_{\mathrm{VIX}} + \beta_{\mathrm{VIX}}\mathrm{VIX}_{t-1} + \sqrt{h_{t,\mathrm{VIX}}\Delta}\,\epsilon_{t,\mathrm{VIX}}$$
$$h_{t+1}^{\mathrm{VIX}} = \sigma_{\mathrm{VIX}}^2 + \kappa\left(h_t^{\mathrm{VIX}} - \sigma_{\mathrm{VIX}}^2\right) + a h_t^{\mathrm{VIX}}(\epsilon_{t,\mathrm{VIX}}^2 - 1 - 2\epsilon_{t,\mathrm{VIX}}\gamma)$$

VIX point forecasting

	RG	EG	G	HN	HNvd	JIVR
2007 – 2018	2.870	3.259	3.278	4.298	3.903	1.894
2007 – 2012	3.383	3.972	3.944	5.497	4.866	2.247
2013 – 2018	2.242	2.339	2.435	2.591	2.603	1.452

Table: One-day ahead VIX forecast RMSE

The first five columns are extracted from table 5 of Hansen et al. (2024). RG: Realized GARCH model, EG: EGARCH, G: GARCH, HN: Heston-Nandi GARCH, HNvd: Heston-Nandi GARCH with the variance-dependent SDF

Conclusion

- This study develops a parametric model for implied volatility surfaces.
  - The factors have a financial interpretation (long-term level, time to maturity slope, moneyness slope, smile attenuation and smirk).
  - The convenient IV surface asymptotic behavior allows for extrapolation beyond levels observed in the data.
  - Daily estimate via Bayesian-OLS approach.
- Joint dynamics of underlying asset returns and IV surface coefficients.
  - It incorporates forward volatility measure in the S&P500 conditional variance.
  - Efficient way to assess market risk of option portfolios.
  - Our approach significantly alleviates what has long been considered as a major obstacle to the development of IV market models.