

Mindless Fitting?

8th Winter School on Mathematical Finance

Lunteren, 15 January 2009

Piotr Karasinski

email: piotr.g.karasinski@gmail.com

Introduction

We are required to mark-to-market non-plain, exotic, products consistently with the *market-observed* prices of liquid vanilla products.

Thus for each exotic we must have a one-to-one mapping between vanilla prices and the exotic's price. Such mapping is called the mark-to-market model as it produces mark-to-market price and risk exposure. Risk management policies (risk limits, desire to minimise volatility of the mark-to-market P&L) typically compel traders to hedge exotics with vanillas such that the combined risk exposure, measured by the mark-to-market model, is close to zero.

In the traditional approach we set the price of an exotic equal to its' value given by a traditional derivatives valuation model that assumes a certain stochastic evolution of the relevant risk factors. To fit vanilla prices practitioners often use (are forced to use?) over-parametrised models in which risk factor dynamics can be counter-intuitive. Does this produce a good model, i.e., does hedging to such model's risk exposure result in realised replication costs that is close to the initial exotic's price the model produces? How can we find an answer to this question?

What are the alternatives? Can we start with a price of an exotic produced by a standard derivatives valuation model, with risk factors' dynamics that makes sense (who is to judge?), and somehow, externally, adjust the price to reflect the difference between market and model prices of relevant vanilla options? Would the resulting mapping produce a hedging model that is better than the one based on the traditional approach?

Modelling Challenge:

**Pricing Long-Dated Bermudan
Cross-Currency Products**

FX Implied Volatilities on 16th Nov 06

<i>Expiry</i>	<i>FwdFx</i>	<i>ATM</i>	<i>25 RR</i>	<i>25 ST</i>	<i>10 RR</i>	<i>10 ST</i>
ON	118.0	8.8%	-0.7%	0.2%	-1.1%	0.5%
1W	117.9	6.3%	-0.6%	0.2%	-1.0%	0.5%
1M	117.5	6.5%	-0.6%	0.2%	-1.0%	0.5%
3M	116.5	6.8%	-0.7%	0.2%	-1.3%	0.6%
6M	115.2	7.0%	-0.9%	0.2%	-1.5%	0.8%
1Y	112.8	7.3%	-1.0%	0.2%	-1.9%	1.0%
2Y	108.7	7.3%	-1.5%	0.2%	-2.8%	1.3%
3Y	105.1	7.3%	-2.0%	0.2%	-3.8%	1.3%
5Y	98.6	7.6%	-2.9%	0.1%	-5.3%	1.6%
7Y	92.7	8.1%	-3.6%	0.0%	-6.5%	1.9%
10Y	85.8	9.9%	-4.6%	-0.1%	-8.0%	2.0%
15Y	75.9	12.5%	-4.8%	0.3%	-8.1%	2.6%
20Y	68.0	15.5%	-4.9%	0.1%	-8.1%	2.2%
30Y	55.4	18.6%	-4.9%	0.0%	-8.1%	1.9%

25RR (10RR) refer to the level of 25% (10%) delta risk reversals, defined as the difference between implied volatilities of 25% (10%) delta calls and puts

25ST (10ST) refer to 25% (10%) strangles, defined as the average sum of 25% (10%) delta calls and puts implied volatilities less ATM volatility.

Introduction to FX/IR Models

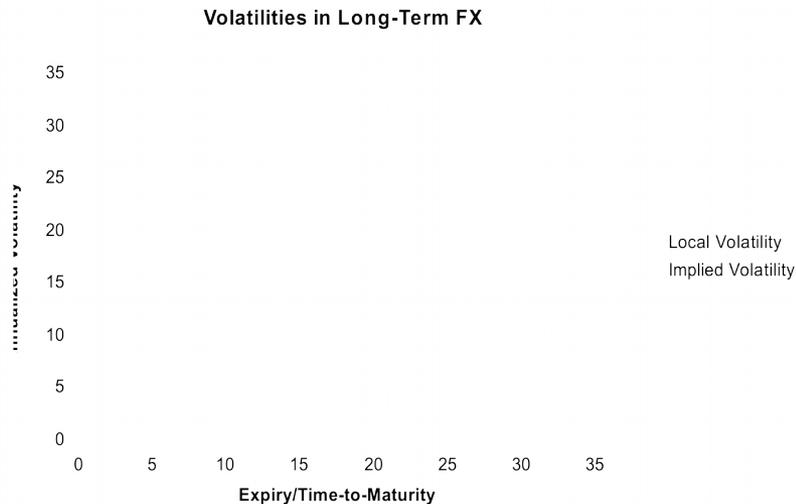
- To price long-dated FX and FX/IR products on a given currency pair we need to jointly model evolution of:
 - spot FX rate
 - interest rates for each of the two currencies.
- Why do we need to model rates jointly with spot FX for long-dated FX products?
 - forward FX rates levels and their volatilities are determined by spot FX and level of rates in the two currencies and their respective volatilities
- As a result an FX/IR model must have at least 3 stochastic factors

The Most Simple FX/IR Model

- The most simple and at the same time the most analytically tractable model we can build is a 3-factor model in which:
 - spot FX follows the standard one-factor lognormal model
 - rates for each of the two currencies follow one-factor normal (Gaussian) mean-reverting model.
- In this model forward FX rate follows lognormal process with local volatility function that is determined by:
 - spot FX volatility
 - basis point volatility of the spread between rates in the two countries
 - correlation between changes in spot FX and changes in spread between the two rates (“domestic” minus “foreign”).
- As the forward FX rate is lognormal, the price of a standard European put or call is given by the Black-Scholes formula with volatility input that can be easily computed from the underlying three-factor model inputs.

FX Volatility in the Simple Model: Example I

- The graph below shows implied Black volatility (bottom curve) as a function of time to expiry and local volatility (upper curve) of a forward FX contract as a function of time to maturity when:
 - spot FX volatility is 10%
 - intra-curve rate spread volatility is 100 basis points and
 - correlation between changes in spot FX and changes in rate spread is zero.



The forward FX local volatility decreases with the decrease in time_ to_ maturity following the upper curve.

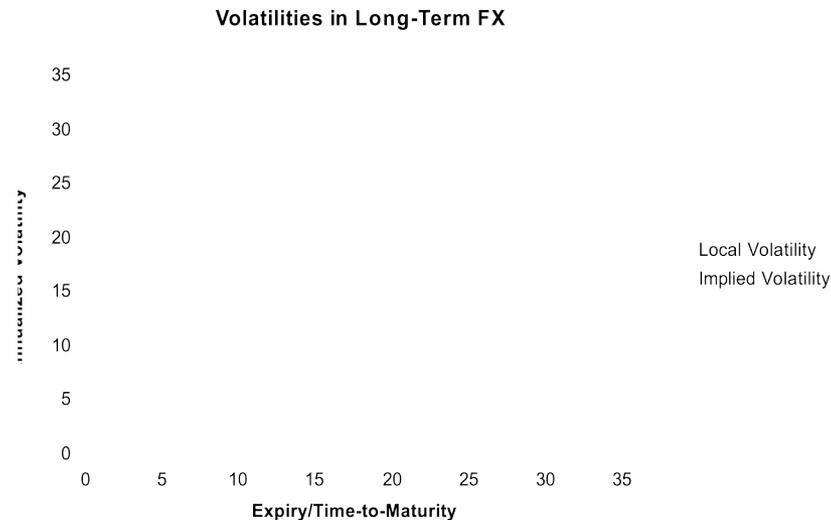
The implied volatility for a given expiry is an average value (strictly speaking mean-square-root average) of the level of forward FX volatilities for maturities between 0 and expiry.

At the 30 year point: the implied FX volatility - at 20% - is twice, and the forward FX volatility - at 31.6% - is three times, the size of local spot FX volatility.

Volatility of intra-rate spread is responsible for this effect.

FX Volatility in the Simple Model: Example II

- It instructive to see in the graph below what happens to both implied and local volatilities of FX forwards from **Example I** when:
 - we set the spot FX volatility to zero
 - while keeping the intra-curve spread volatility unchanged at 100 basis point:



At 30 year time horizon reducing spot FX volatility from 10% to 0%:

- lowers implied FX volatility from 20% to 17.3%
- lowers local FX forward rate volatility from 31.6% to 30%

Cognitive Bias in Model Building

Value of a Bermudan on a Call Date = $\max(\text{HoldValue}, \text{ExerValue})$

$$= \text{ExerValue} + \max(\text{HoldValue} - \text{ExerValue}, 0)$$

People are obsessed with building a model that matches Exercise Value!
Why? Because this value is typically known assuming liquid vanilla markets.

In fact we need to predict the difference between HoldingValue and Exercise Value and not Exercise Value alone.

A model that matches vanillas today is not necessarily good at predicting the cost of dynamically replicating the value of Bermudan callability = $\max(\text{HoldValue} - \text{ExerValue}, 0)$

Is it possible that a model that fits vanillas less well than the perfect model can actually do better job at predicting $\text{HoldValue} - \text{ExerValue}$?

Fischer Black Believed that Simple Models are Better!

Quote from Perry Mehrling's book entitled "FischerBlack and the Revolutionary Idea of Finance" page 14:

“In a world where nothing is constant, complex models are inherently fragile, and prone to break down when you lean on them too hard. Simple models are potentially more robust , and easier to adapt as the world changes. Fisher [Black] embraced simple models as his anchor in the flux because he thought they were more likely to survive Darwinian selection as the system changes.”

Let's Think Out-of-the-Box!

Natural Selection and Derivatives Pricing

Derivatives Modelling and Darwin's "Natural Selection"

Food-for-thought on derivatives modelling provided by a quote from Steve Jones's book

Almost Like a Whale: The Origin of Species Updated (*), Chapter IV Natural Selection,

"I once worked for a year or so, for what seemed good reasons at the time, as a fitter's mate in a soap factory on the Wirral Peninsula, Liverpool's Left Bank. It was a formative episode, and was also, by chance, my first exposure to the theory of evolution.

To make soap powder, a liquid is blown through a nozzle. As it streams out, the pressure drops and a cloud of particles forms. These fall into a tank and after some clandestine coloration and perfumery are packaged and sold. In my day, thirty years ago, the spray came through a simple pipe that narrowed from one end to the other. It did its job quite well, but had problems with changes in the size of the grains, liquid spilling through or - worst of all - blockages in the tube.

Those problems have been solved. The success is in the nozzle. What used to be a simple pipe has become an intricate duct, longer than before, with many constrictions and chambers. The liquid follows a complex path before it sprays from the hole. Each type of powder has its own nozzle design, which does the job with great efficiency.

(continued)

(*). Published 1999 by Doubleday

Derivatives Modelling and Darwin's "Natural Selection"

(continued)

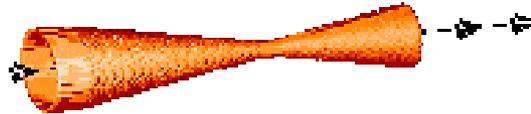
What caused such progress? Soap companies hire plenty of scientists, who have long studied what happens when a liquid sprays out to become a powder. The problem is too hard to allow even the finest engineers to do what enjoy the most, to explore the question with mathematics and design the best solution. Because that failed, they tried another approach. It was the key to evolution, design without a designer: the preservation of favourable variations and the rejection of those injurious. It was, in other words, natural selection.

The engineers used the idea that moulds life itself: descent with modification. Take a nozzle that works quite well and make copies, each changed at random. Test them for how well they make powder. Then, impose a struggle for existence by insisting that not all can survive. Many of the altered devices are no better (or worse) than the parental form. They are discarded, but the few able to do a superior job are allowed to reproduce and are copied - but again not perfectly. As generations pass there emerges, as if by magic, a new and efficient pipe of complex and unexpected shape.

Natural selection is a machine that makes almost impossible things.”

Darwin on the Shop Floor: Evolution of a Nozzle (*)

- Before mutation and selection



- After 20 generations of mutation & selection



Outline

- **What Makes a Good Model?**
- **Foundations of Derivatives Pricing**
- **Marking-to-Market Non-Plain Products: Traditional Approach**
- **New Approach to Marking-to-Market Non-Plain Products**
- **Maths of the New Approach**
- **Advantages of the New Approach?**

What Makes a Good Model?

What is a Good Mark-to-Market Model?

Mark-to-Market model has two components

- Term-structure IR model
- The way the IR model is used: model calibration, etc.

Some of the criteria used in judging mark-to-market models

- Matching prices of reference plain liquid products
- Matching market prices of non-plain products (produced by other banks?)
- Accepted market practice, market standard
- Easy to explain and can be disclosed to wider audiences
 - Marketers, Model Validation, Risk Management, Auditors, Regulators, etc.
- Conservatism
- Ease of implementation and cost of running

What is a good hedging model?

Hedging models has two components

- Term-structure IR model
- The way the IR model is used: model calibration, risk measures against which we hedge, trade-off between local risk and transactions costs, etc.

Criteria used in judging hedging models

- Doing the best possible job replicating non-plain products
- Capturing value of a non-plain trade
- Trade-off between local risk hedging and transactions costs
- Minimising uncertainty in the realised replication cost
 - How can we measure such uncertainty?

What Model Should We Use for Hedging?

- Is hedging to mark-to-market model the best way to replicate non-plain products with plain ones?
- If not we should consider using:
 - Alternative term-structure models
 - Increasing number of factors, different skew properties, etc.
 - Alternative method of calibration, additional risk measures
- Am I taking model risk if I hedge to a model is not the same as the model I use for marking-to-market?
 - My mark-to-market model is likely to report non-zero risk even though my portfolio has zero risk according to the hedging model

Non-Plain Product's Model Risk

Model-Dependence of Prices

- Prices depend on the terms-structure model used
 - Number of curve factors included and rate-level dependence of volatility
 - Assumptions about jumps and stochastic volatility
- Prices depend on how a particular model is used
 - What we calibrate model parameters to and how
 - Mark-to-market methodology based on the chosen model

Risk in Capturing P&L

- Realised replication cost is uncertain
 - depends on what model we use and how
 - the market doesn't follow any particular model, can have structural changes, etc.
 - transactions costs

Model Risk in Capturing P&L

- How do we know that our models, and the way we use them, will allow us to capture the economic value (lock-in P&L) of a non-plain portfolio or a single trade over the trade's/portfolio's lifetime which sometime may extend to 30 and even 40 years?
- How can we convince ourselves, and importantly significant others, that Warren Buffet's prophecy (from 2000 Berkshire Hathaway's annual report) which says "In extreme cases, mark-to-model degenerates into what I would call mark-to-myth" will not be fulfilled?
- In order to build tractable pricing/hedging models we make simplifying assumptions about the random behaviour of risk factors that determine the size of our non-plain products' cashflows and cost of replicating these cashflows with plain products.
- Our cumulative experience in building and using models, as well as intuition derived from this experience, guide us in choosing models, and the way we use them, and provides a level of confidence.

Where Do Prices of Liquid Plain Products Come From?

Prices of liquid plain products are impacted by:

- Supply-and-demand (end-users' hedging needs and risk-aversion, risk limits and risk appetite of providers, ...), overall market liquidity
- Beliefs about future swap curve behaviour (volatilities and correlations, etc) as these beliefs can be translated into a term-structure model which produces prices of plain products that can be compared with currently observed market prices of those non-plain products.
- Our beliefs about future swap curve behaviour are shaped:
 - by swap curve's past behaviour
 - “economic theory”
- Some of the relative value players are fitting parameters in their term-structure models to swaption market data and measure relative value by
 - Comparing current levels of such implied parameters against past levels: past highs/lows etc.
 - Looking at individual swaptions' current and historical levels of fit residuals.
 - Such relative value measures impact trading decisions and thus feed-back to market mid prices.
- Overall market liquidity and risk aversion is another factor. During the LTCM crisis the implied vols of long-dated swaptions collapsed.

Foundations of Derivatives Pricing

Pricing and Hedging of Non-Plain Products

How much should we charge for an interest rate exotic?

How should we hedge it?

- These two questions are obviously inter-related:
 - The price we should charge for an exotic should equal the present value of future hedging costs plus a reasonable profit margin.
- So, how do we go about designing a hedging strategy?
- Given a hedging strategy, how do we find the present value of realised hedging costs?
- Will the present value of realised hedging costs depend on the future evolution of market prices/rates for our yield curve and volatility hedges?
If so, what is the distribution of potential outcomes?

Fundamental Theorem of Derivatives Pricing

When:

All relevant market risk factors can be hedged, i.e., we can hedge yield curve, volatility and correlation risk with traded instruments.

We can hedge in continuous time — we can re-hedge as often as we need — without incurring transactions cost.

Market risk factors follow a *diffusion process*. The instantaneous covariance matrix between changes in risk factors' levels is a known function of time and risk factors' levels.

A derivative product has a unique price.

In plain English this means that there exists a unique hedging strategy for which the realised cost of replication is path-independent.

Marking-to-Market Non-Plain Products: Traditional Approach

Marking Non-Plain Products Relative to Plain Ones

- We are required to mark-to-market non-plain products in a way that is consistent with mark-to-market prices of plain products
- Thus for each non-plain product we need to design a functional form

$$V_{non-plain}^{(mkt)}(t, y, \sigma_V)$$

– for the dependence of the product's mark-to-market price on:

- valuation date
- t swap curve
- y parameters that enter into our mark-to-market model for plain products
- σ_V

Traditional Approach

- We set the non-plain products mark-to-market value to its' value given by a term-structure model

$$V_{non-plain}^{(mkt)}(t, y, \sigma_V) = V_{non-plain}^{(model)}(t, y, p(y, \sigma_V))$$

- We calibrate model parameters to a product-dependent set of market prices of plain caps/floors and swaptions
 - Typical set of model parameters for a one-factor model: mean-reversion speed, parameters entering into the local volatility function, skew parameter.

Fitting Model Parameters for Bermudans

- Exact fit to a small set of “properly” chosen benchmarks
 - Number of fitted parameters equal to the number of benchmarks
 - Example: fix mean-reversion speed, fit the local volatility function so that we match “diagonal” swaptions with “properly” chosen strikes:
 - at-the-money
 - equal to the underlying’s swap rate (for Bermudans on plain swaps)
 - etc.
- Least-square fit (could be vega weighted)
 - Number of benchmarks $>$ Number of model parameters
 - Issues with multiple local minima

Potential Problems with the Traditional Approach

- Overparametrization
- Risk factors' dynamics that “does not make sense”
- Excessive variability of model parameters
- Proliferation of model calibrations
- Inconsistent risk for plain products we calibrate to and/or hedge with:
 - Risk produced by our term structure model could differ from risk produced by the plain products' mark-to-market model
- Hedging to the “mark-to-model” model may not be the best way of capturing economic value.

New Approach to Marking-to-Market

Non-Plain Products

New Framework for Marking Non-Plain Products

- We construct a portfolio that consists of a non-plain product and so called “mark-to-market” hedge portfolio of plain products that:
 - hedges away sensitivities of the non-plain product’s model price to model parameters
 - satisfies additional criteria: “stability” conditions, etc.
- We set the non-plain product’s mark-to-market price equal to:
 - the combined model price of the non-plain product and “mark-to-market” hedge portfolio
 - less the mark-to-market price of the “mark-to-market” hedge portfolio
- The above framework satisfies, by construction, the requirement that we mark non-plain products consistently with market prices of plain products.

Calibrating Model Parameters to Plain Products

- We may prefer to have a small number of model parameters and model parameterizations that make economic sense.
- We need to decide which model parameters to keep fixed. For example, we may keep correlations between factor shocks constant.
- We may want to impose priors on the model parameters we calibrate based on historically fitted parameter levels and other criteria.
- We need to decide which specific caps/floors and swaptions to include in our calibration set: we need to pick expiries, underlyers' tenors, and strikes.
- While we may prefer not to have trade-specific calibration sets, one model calibration may not work for all non-plain products. We hope that it will suffice to have a small number of product-group specific model calibrations.

Choosing Plain Product Hedges

- We need to decide up-front which plain products to include in the mark-to-market hedge portfolio for a given non-plain product.
- The number of plain products we choose could be larger than the number of model parameters.
- We fix swaptions' strike levels, expiry dates and underlier tenors, cap strikes and start/end dates.
- For Bermudan swaptions we would include both on-diagonal and off-diagonal swaptions in the hedge portfolio.

Finding Hedge Portfolio Weights

We find the hedge portfolio weights by requiring that:

- We hedge away sensitivities of the non-plain product's model price to model parameters while at the same time we:
 - minimize the total portfolio gamma
 - satisfy hedge stability condition, i.e., that the hedged non-plain product portfolio's delta with respect to model parameters is insensitive to small curve and model parameter shifts
 - satisfy “other” constraints

Maths of the New Approach

Maths of Mark-to-Market: Notation

σ_V	vector of vol - type parameters in plain products' mark - to - market model
p	vector of model parameters
$V(t, p), V^{(mkt)}(t, \sigma_V)$	time t non - plain model and mark - to - market value
$O_k(t, p), O_k^{(mkt)}(t, \sigma_V)$	time t model and market value of k - th plain hedge instrument
N_O	number of plain instruments which can be included in hedge portfolio
t_i	i - th discrete time point at which we mark - to - market non - plain product
$p^{(i)}$	model parameter vector calibrated at time t_i
$\omega_1^{(i)}, \dots, \omega_{N_O}^{(i)}$	time t_i amounts of plain instruments in hedge portfolio
$P(t, p)$	model value, $V + \sum_{k=1}^{N_O} \omega_k O_k$, of the "non - plain + hedges" portfolio
$P^{(mkt)}(t, \sigma_V)$	market value, $V^{(mkt)} + \sum_{k=1}^{N_O} \omega_k O_k^{(mkt)}$, of the "non - plain + hedges" portfolio

Maths of Mark-to-Market

- We set the combined time t_i mark-to-market value of the non-plain product and the hedge portfolio equal to the respective model value:

$$P_i^{(mkt)} = P_i = V(t_i, p^{(i)}) + \sum_{k=1}^{N_o} \omega_k^{(i)} O_k(t_i, p^{(i)})$$

- The mark-to-market value of the non-plain product equals the above value less the mark-to-market value of the hedge portfolio:

$$V^{(mkt)}(t_i, \sigma^{(i)}) = V(t_i, p^{(i)}) + \sum_{k=1}^{N_o} \omega_k^{(i)} (O_k(t_i, p^{(i)}) - O_k^{(mkt)}(t_i, \sigma^{(i)}))$$

non-plain product
mark-to-market
non-plain product
model price
adjustment due to differences between model
and market prices of plain hedge instruments

Decomposing Changes in Mark-to-Market

- Profit-and-Loss attribution plays an important role in managing a portfolio of derivative products. Here is how it works for non-plain products marked to market using our prescription.
- We write the change in mark-to-market value of a non-plain product as a sum of three terms:

$$\begin{aligned}
 V^{(mkt)}(t_{i+1}, \sigma_V^{(i+1)}) - V^{(mkt)}(t_i, \sigma_V^{(i)}) &= V(t_{i+1}, p^{(i+1)}) - V(t_i, p^{(i)}) \\
 &+ \sum_{k=1}^{N_O} \omega_k^{(i)} \left[\left(O_k(t_{i+1}, p^{(i+1)}) - O_k^{(mkt)}(t_{i+1}, \sigma_V^{(i+1)}) \right) - \left(O_k(t_i, p^{(i)}) - O_k^{(mkt)}(t_i, \sigma_V^{(i)}) \right) \right] \\
 &+ \sum_{k=1}^{N_O} \left(\omega_k^{(i+1)} - \omega_k^{(i)} \right) \left(O_k(t_{i+1}, p^{(i+1)}) - O_k^{(mkt)}(t_{i+1}, \sigma_V^{(i+1)}) \right)
 \end{aligned}$$

Terms in the P&L Decomposition

- Change in the model value of a non-plain product due to yield curve shift, change in model calibration, and passage of time:

$$V(t_{i+1}, p^{(i+1)}) - V(t_i, p^{(i)})$$

- Change in the price adjustment due to change in discrepancy between model and market values of the time t_i hedge portfolio (unchanged hedge portfolio weights):

$$\sum_{k=1}^{N_O} \omega_k^{(i)} \left[\left(O_k(t_{i+1}, p^{(i+1)}) - O_k^{(mkt)}(t_{i+1}, \sigma_V^{(i+1)}) \right) - \left(O_k(t_i, p^{(i)}) - O_k^{(mkt)}(t_i, \sigma_V^{(i)}) \right) \right]$$

- Adjustment to mark-to-market price due to discrepancy between model and market values of the hedge added at time t_{i+1} :

$$\sum_{k=1}^{N_O} (\omega_k^{(i+1)} - \omega_k^{(i)}) \left(O_k(t_{i+1}, p^{(i+1)}) - O_k^{(mkt)}(t_{i+1}, \sigma_V^{(i+1)}) \right)$$

Non-Plain Product's Delta Risk

Hedge portfolio based on Parameter Hedging

- Non-plain product's vega risk is equal to the total derivative of the product's mark-to-market value with respect to volatility parameter vector σ_V

$$\begin{aligned}
 \frac{\partial V^{(mkt)}(t, y, \sigma_V)}{\partial \sigma_V} &= \frac{\partial V(t, y, p)}{\partial \sigma_V} + \sum_{N_o} \omega_k \times \left[\frac{\partial O_k(t, y, p)}{\partial \sigma_V} - \frac{\partial O_k^{(mkt)}(t, y, \sigma_V)}{\partial \sigma_V} \right] \\
 &\quad + \sum_{N_o} \left[\frac{\partial \omega_k(t, y, p)}{\partial \sigma_V} + \frac{\partial \omega_k(t, y, p)}{\partial p} \frac{\partial p(t, y, \sigma_V)}{\partial \sigma_V} \right] \times (O_k - O_k^{(mkt)})
 \end{aligned}$$

mark-to-market curve delta model curve delta difference between model and mark-to-market curve deltas of the hedge portfolio derivative of a portfolio weight w.r.t. yield curve bumps difference between model and market values for a plain product additional curve delta adjustment

Non-Plain Product's Vega Risk

Hedge portfolio based on Parameter Hedging

- What is vega risk?
 - Risk due to small shifts in parameters entering mark-to-market model for plain caps/floors and swaptions
 - ATM vols or other vol-type variables, skew/smile parameters, etc.
- What is the vega risk of a non-plain product equal to?
 - When the “mark-to-market” hedge portfolio is chosen to hedge away deltas of the non-plain product model value with respect to model calibration
 - The vega risk of a non-plain product's mark-to-market value equals to:
 - negative of vega risk of the “mark-to-market” hedge portfolio plus a small adjustment

Non-Plain Product's Vega Risk Formula

Hedge portfolio based on Parameter Hedging

Non-plain product's vega is equal to the partial derivative of the product's mark-to-market value with respect to volatility parameter vector σ_V

$$\frac{\partial V^{(mkt)}(t, y, \sigma_V)}{\partial \sigma_V} = - \sum_{k=1}^{N_O} \omega_k \times \frac{\partial O_k^{(mkt)}(t, y, \sigma_V)}{\partial \sigma_V} + \frac{\partial p(t, y, \sigma_V)}{\partial \sigma_V} \times \sum_{k=1}^{N_O} \frac{\partial \omega_k(t, y, p)}{\partial p} \times (O_k - O_k^{(mkt)})$$

vega of "mark-to-market"
 plain product's vega risk
 model calibration's delta to market vol params
 portfolio weight's delta to model calibration
 plain instrument's model mis-marking

Advantages of the New Approach

Advantages of the New Approach

- Reduced number of model parameters
- We can have risk factors' dynamics that makes economic sense
- Reduced variability of model parameters
- Small number of model calibrations
- Better aggregation of risk