

Path-dependent volatility

Combining benefits from local volatility and stochastic volatility

Julien Guyon

Bloomberg L.P.
Quantitative Research

Columbia University
IEOR

Financial Engineering Practitioners Seminar
New York, September 17, 2018

jguyon2@bloomberg.net

Risk magazine, October 2014



Cutting edge: Derivatives pricing

Path-dependent volatility

So far, path-dependent volatility models have drawn little attention compared with local volatility and stochastic volatility models. In this article, Julien Guyon shows they combine benefits from both and can also capture prominent historical patterns of volatility

Three main volatility models have been used so far in the finance industry: constant volatility, local volatility (LV) and stochastic volatility (SV). The first two models are complete: since the asset price is driven by a single Brownian motion, every payoff admits a unique self-financing replicating portfolio consisting of cash and the underlying asset. Therefore, its price is uniquely defined as the initial value of the replicating portfolio, independent of utilities or preferences. Unlike the constant volatility models, the LV model is flexible enough to fit any arbitrage-free surface of implied volatilities (henceforth, 'smile'), but then no more flexibility is left. Calibrating to the market smile is useful when one sells an exotic option whose risk is well mitigated by trading vanilla options – then the model correctly prices the hedging instruments at inception.

For their part, SV models are incomplete: the volatility is driven by one of several extra Brownian motions, and as a result perfect replication and price uniqueness are lost. Modifying the drift of the SV leaves the model arbitrage-free, but changes option prices.

Using SV models allows us to gain control of key risk factors such as volatility of volatility (vol-of-vol), forward skew and spot-vol correlation. SV models generate joint dynamics of the asset and its implied

price uniqueness and parsimony: it is remarkable that so many popular properties of SLV models can be captured using a single Brownian motion. Although perfect delta-hedging is unrealistic, incorporating the path-dependency of volatility into the delta is likely to improve the delta-hedge. Not only that, we will see that, thanks to their huge flexibility, PDV models can generate spot-vol dynamics that are not attainable using SLV models.

Below, we first introduce the class of PDV models and then explain how we calibrate them to the market smile. Subsequently, we investigate how to pick a particular PDV.


Path-dependent volatility models

PDV models are those models where the instantaneous volatility σ_t depends on the path followed by the asset price so far:


$$\frac{dS_t}{S_t} = \sigma(t, (S_u, u \leq t)) dW_t$$

where, for simplicity, we have taken zero interest rates, repo and dividends. In practice, the volatility $\sigma_t = \sigma(t, S_t, X_t)$ will often be





VIX 2012-18.png



Why study and use path-dependent volatility?

- Path-dependent volatility (PDV) models have drawn little attention compared with local volatility (LV) and stochastic volatility (SV) models
- This is **unfair**: PDV models combine benefits from both LV and SV, and even go beyond
- Like LV: **complete** and can **fit exactly the market smile**
- Like SV: produce a **wide variety of joint spot-vol dynamics**
- Not only that:
 - 1 Can generate spot-vol dynamics that are **not attainable using SV models**
 - 2 Can also **capture prominent historical patterns of volatility**

Outline

- A short recap on volatility modeling
- PDV models: a (necessarily) brief history, and what we want from them
- Smile calibration of PDV models
- Choose a particular PDV to generate desired spot-vol dynamics
- Choose a particular PDV to capture historical patterns of volatility
- Concluding remarks
- Discussion

A short recap on volatility modeling

- **Constant volatility** (Bachelier, 1905; Black and Scholes, 1973) and **LV** (Dupire, 1994) are **complete** models: every payoff admits a unique self-financing replicating portfolio consisting of cash and the underlying asset \implies **unique price**
- **LV** flexible enough to **fit exactly any arbitrage-free smile**—but no more flexibility is left
- **SV** models are **incomplete** \implies **no unique price**. But they give control on key risk factors such as vol of vol, forward skew, and spot-vol correlation. Unlike LV, they generate **rich joint dynamics of the asset and its implied volatilities**
- To allow SV models to perfectly calibrate to the market smile, one can use **SLV** models + **particle method**. Modifies spot-vol dynamics, but only slightly (except maybe for small times t): usually the LV component (**leverage function**) flattens as t grows

Can we combine benefits from both LV and SV?

- Can we build **complete** models that have all the nice properties of SLV models, namely, **rich spot-vol dynamics**, and **calibration to market smile**?
- For instance, can we build a complete model that is calibrated to a flat smile, and yet produces very negative short term forward skews?
- Tempting but wrong to quickly answer 'no', by arguing that the only complete model calibrated to the smile is the LV model.
- This is **not true**: we will show that PDV models, which are complete, can produce rich spot-vol dynamics and, on top of that, can be perfectly calibrated to the market smile
- Benefits of model completeness: price uniqueness and parsimony. All properties of SLV models can be captured **using a single Brownian motion**. Although perfect delta-hedging is unrealistic, incorporating the path-dependency of volatility into the delta is likely to improve the delta-hedge.
- PDV models actually **go beyond SLV models**: they can generate spot-vol dynamics that are not attainable using SLV models.

PDV models

- PDV models are those models where the instantaneous volatility σ_t depends on the **path** followed by the asset price so far:

$$\frac{dS_t}{S_t} = \sigma(t, (S_u, u \leq t)) dW_t$$

- In practice, $\sigma_t \equiv \sigma(t, S_t, X_t)$ where $X_t =$ **finite set of path-dependent variables**: running or moving averages, maximums or minimums, realized variances, etc.
- Most famous examples: ARCH/GARCH models. Discrete-time and hardly used in the derivatives industry

The Hobson-Rogers model (1998)

- $\sigma_t = \sigma(X_t)$; $X_t = (X_t^1, \dots, X_t^n)$ where the X_t^m are exponentially weighted moments of all the past log increments of the asset price:

$$X_t^m = \int_{-\infty}^t \lambda e^{-\lambda(t-u)} \left(\ln \frac{S_t}{S_u} \right)^m du$$

- $n = 1$: σ_t depends only on $X_t^1 = \ln S_t - \int_{-\infty}^t \lambda e^{-\lambda(t-u)} \ln S_u du =$ the difference between current log price and a weighted average of past log prices \implies **vol determined by local trend of the asset price** over a period of order $1/\lambda$ years (e.g., 1 month if $\lambda = 12$)
- Supported by empirical studies (see later)
- Choice of an infinite time window and exponential weights only guided by **computational convenience**: ensures that (S_t, X_t) is a Markov process \implies price of a vanilla option reads $u(t, S_t, X_t)$ where u is the solution to a second order parabolic PDE
- Implied vols at time 0 in the model depend not only on the strike, maturity, and S_0 , but also on **all** the past asset prices through X_0

Four natural and important questions arise

- 1 Can we specify $\sigma(\cdot)$ and λ so that the model **fits exactly the market smile**?
Platania & Rogers, Figà-Talamanca & Guerra only gave approximate calibration results
- 2 Does the calibrated model have **desired dynamics of implied volatility**, such as large negative short term forward skew for instance?
- 3 In the definition of X_t , can we use **general weights and a finite time window** $[t - \Delta, t]$ instead of $(-\infty, t]$, so that the vol truly depends only a limited portion of the past? The generalization in Foschi & Pascucci is partial as it requires positive weights on $[0, t]$.
- 4 Much more importantly: how do we **generalize to other choices of X_t** ?
The generalization in Hubalek *et al.*, where the vol depends on a particular modified version of the offset X_t^1 , is also very partial

Our approach will solve these four questions all at once

- First we choose **any** set of path-dependent variables X_t and **any** function $\sigma(t, S, X)$ so that the PDV model with $\sigma_t = \sigma(t, S_t, X_t)$ has desired spot-vol dynamics and/or captures historical patterns of volatility
- Then we define a new model by **multiplying $\sigma(t, S_t, X_t)$ by a leverage function $l(t, S_t)$** and we perfectly calibrate l to the market smile of S using the **particle method** (Guyon and Henry-Labordère, 2011)
- Usually, multiplying $\sigma(t, S_t, X_t)$ by the calibrated leverage function **distorts only slightly** the spot-vol dynamics
- This way we mimic SLV models, with the **'pure'** PDV $\sigma(t, S_t, X_t)$ playing the role of SV, but we stay in the world of complete models
- Not only that: thanks to their huge flexibility, PDV models can generate spot-vol dynamics that are **not attainable using SLV models**
- Same program can be run by choosing two functions $a(t, S, X)$ and $b(t, S, X)$ instead of only one function $\sigma(t, S, X)$, and then defining $\sigma_t^2 = a(t, S_t, X_t) + b(t, S_t, X_t)l(t, S_t)$
 $b \equiv 1$: complete analogue of incomplete *additive* SLV models

Smile calibration of PDV models: Particle method

- Given a PDV $\sigma(t, S, X)$, we can uniquely build the leverage function $l(t, S)$ such that the PDV model

$$\frac{dS_t}{S_t} = \sigma(t, S_t, X_t) l(t, S_t) dW_t \quad (1)$$

fits exactly the market smile of S

- From Itô-Tanaka's formula, Model (1) is exactly calibrated to the market smile of S if and only if

$$\mathbb{E}^{\mathbb{Q}} [\sigma(t, S_t, X_t)^2 | S_t] l(t, S_t)^2 = \sigma_{\text{Dup}}^2(t, S_t)$$

where \mathbb{Q} denotes the unique risk-neutral measure and σ_{Dup} the Dupire LV
 \implies **calibrated model satisfies the nonlinear McKean stochastic differential equation**

$$\frac{dS_t}{S_t} = \frac{\sigma(t, S_t, X_t)}{\sqrt{\mathbb{E}^{\mathbb{Q}}[\sigma(t, S_t, X_t)^2 | S_t]}} \sigma_{\text{Dup}}(t, S_t) dW_t$$

- The **particle method** (Guyon & Henry-Labordère, 2011) computes the above conditional expectation, hence the leverage function $l(t, S) = \sigma_{\text{Dup}}(t, S) / \sqrt{\mathbb{E}^{\mathbb{Q}}[\sigma(t, S_t, X_t)^2 | S_t = S]}$, on the fly while simulating the paths

Smile calibration of PDV models

- Brunick and Shreve (2013): Given a general Itô process $dS_t = \sigma_t S_t dW_t$ and a special type of path-dependent variable X , there exists a PDV $\sigma(t, S_t, X_t)$ such that, for each t , the **joint** distribution of (S_t, X_t) is the same in both models:

$$\sigma(t, S_t, X_t)^2 = \mathbb{E}^{\mathbb{Q}}[\sigma_t^2 | S_t, X_t]$$

- Only X 's satisfying a type of Markov property are admissible though: running averages are admissible, but moving averages are not; instead, one must pick $X_t = (S_u, t - \Delta \leq u \leq t)$
- Take $X_t = (S_u, 0 \leq u \leq t)$: Brunick-Shreve \implies the price process produced by any SV/SLV model has the same distribution, **as a process**, as a PDV model (**not only the marginal distributions**) \implies There always exists a PDV model that produces exactly the same prices of, **not only vanilla options, but all options**, including path-dependent, exotic options \implies No surprise that PDV models can reproduce popular SLV spot-vol dynamics (see below)

How to choose a particular path-dependent volatility?

Now the crucial question is:

How to choose a particular PDV?

Two main possible goals:

- 1 Generate desired spot-vol dynamics
- 2 Capture historical features of volatility

These two goals are not mutually exclusive: it might very well happen, and it is desirable, that a given choice of a PDV fulfills both objectives at a time

Choose a particular PDV
to generate desired spot-vol dynamics

Choose a particular PDV to generate desired spot-vol dynamics

- Can we choose a PDV $\sigma(t, S, X)$ that, for instance, generates **large negative short term forward skews**, even when it is **calibrated to a flat smile**?
- LV/SLV analogy \implies We need $\sigma(t, S_t, X_t)$ to be negatively correlated with S_t
- May be achieved by picking a decreasing function σ of S alone, but smile calibration would bring us back to pure LV model:

$$\frac{dS_t}{S_t} = \frac{\sigma(t, S_t, X_t)}{\sqrt{\mathbb{E}^{\mathbb{Q}}[\sigma(t, S_t, X_t)^2 | S_t]}} \sigma_{\text{Dup}}(t, S_t) dW_t \quad (2)$$

- What we actually need is $\sqrt{\eta(t, S_t, X_t)}$ to be negatively correlated with S_t , where

$$\eta(t, S, X) \equiv \frac{\sigma(t, S, X)^2}{\mathbb{E}^{\mathbb{Q}}[\sigma(t, S_t, X_t)^2 | S_t = S]}$$

- $\eta(t, S, X) = \text{PDLVR} = \text{'path-dependent to local variance ratio'}$
- The PDLVR or alternatively $D(t, S) = \text{Var}(\eta(t, S_t, X_t) | S_t = S) = \mathbb{E}[(\eta(t, S_t, X_t) - 1)^2 | S_t = S]$ measures deviation from LV: $\text{LV} \iff \eta \equiv 1 \iff D \equiv 0$

Choose a particular PDV to generate desired spot-vol dynamics

- Recall that we want

$$\sqrt{\eta(t, S, X)} \equiv \frac{\sigma(t, S, X)}{\sqrt{\mathbb{E}^{\mathbb{Q}}[\sigma(t, S_t, X_t)^2 | S_t = S]}}$$

to tend to be large when S is small, and conversely

- $\implies \sigma(t, S, X)$ must be negatively linked to S , but **not perfectly**: target correl of the **levels** of spot and vol is more around, say, -50% than around -1% or -99% . **Moderate correlation** property
- Usual SLV models: Mean reversion in the SV \implies Moderate correlation, even if **increments** of spot and vol are extremely negatively correlated

Example 1

Ex.	X_t	$\sigma(S, X)$ producing large forward skew
1	$S_{t-\Delta}$	$\bar{\sigma}1_{\{\frac{S}{X} \leq 1\}} + \underline{\sigma}1_{\{\frac{S}{X} > 1\}}$

Comparing with SV models:

- $\bar{\sigma} - \underline{\sigma} \longleftrightarrow$ **vol of vol**: we need it to be large enough to generate large negative short term forward skew
- $\Delta \longleftrightarrow$ **spot-vol correlation**:
 - S_t small \implies more likely that S_t be smaller than $S_{t-\Delta} \implies$ more likely that σ_t be large
 - The larger Δ , the larger the correlation
- $\Delta \longleftrightarrow$ **mean reversion** too: the smaller Δ , the more **ergodic** the volatility, hence the flatter the forward smile (cf. Fouque-Papanicolaou-Sircar, 2000)

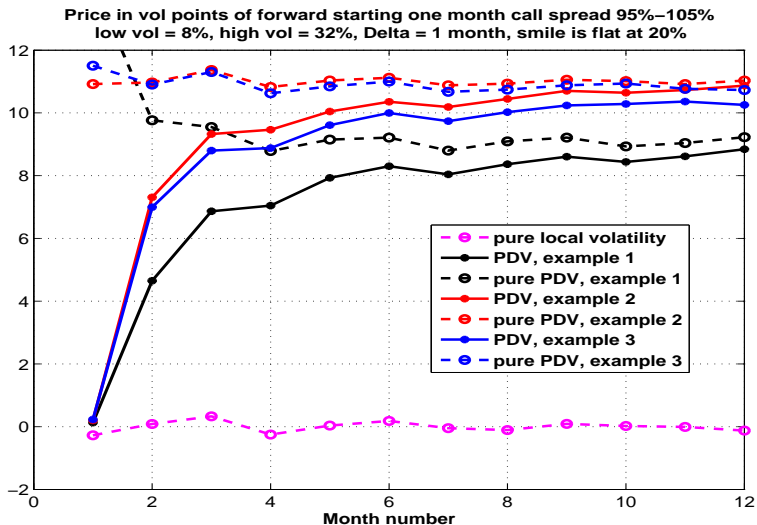
Examples 2 and 3

Ex.	X_t	$\sigma(S, X)$ producing large forward skew
1	$S_{t-\Delta}$	$\bar{\sigma}1_{\{\frac{S}{X} \leq 1\}} + \underline{\sigma}1_{\{\frac{S}{X} > 1\}}$
2	\bar{S}_t^Δ	as above
3	(m_t^Δ, M_t^Δ)	$\bar{\sigma}1_{\{\frac{S-m}{M-m} \leq \frac{1}{2}\}} + \underline{\sigma}1_{\{\frac{S-m}{M-m} > \frac{1}{2}\}}$

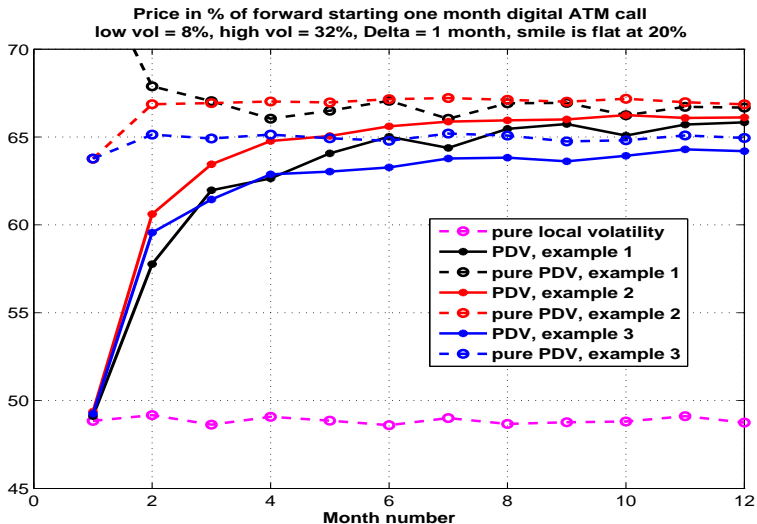
$$\bar{S}_t^\Delta = \frac{\int_0^\Delta w_\tau S_{t-\tau} d\tau}{\int_0^\Delta w_\tau d\tau}, \quad m_t^\Delta = \inf_{t-\Delta \leq u \leq t} S_u, \quad \text{and} \quad M_t^\Delta = \sup_{t-\Delta \leq u \leq t} S_u$$

- Ex. 2: $S_{t-\Delta}$ replaced by moving average \bar{S}_t^Δ . Makes more financial sense: why put all the weight w_τ on $\tau = \Delta$?
- Ex. 3 uses that $\frac{S_t - m_t^\Delta}{M_t^\Delta - m_t^\Delta}$ is positively correlated with S_t . The larger Δ , the larger the correlation


Forward starting 1M call spread 95%-105%



Forward starting 1M ATM digital call




Smiles of pure PDV models



Example 3: Leverage function $l(t, S)$

Leverage function $l(t, S)$: Example 3
low vol = 8%, high vol = 32%, Delta = 1 month, smile is flat at 20%



Short term forward skew: comparison with SLV models

- Consider the SLV model where the SV = exponential O-U process
- To reproduce Example 2 prices, a **huge** positive **return** in volatility is needed when asset returns are negative \implies one must use a very large vol of vol, **beyond 550%**, together with an extreme spot-vol correlation (say, -99%)
- \implies A very large mean reversion above 20 is then somehow **artificially needed** to keep volatility within a reasonable range
- By comparison PDV models, which can directly relate the asset returns to the volatility **levels**, look much more handy and can more naturally generate large forward skews

- Large **positive** short term forward skew: exchange $\bar{\sigma}$ and $\underline{\sigma}$
- One may use **smoothed versions** of the PDV by replacing the Heaviside function by $\frac{1}{2}(1 + \tanh(\lambda x))$ for instance

U-shaped short term forward smile

- What if we want a PDV model **calibrated to a flat smile** and yet that generates a **pronounced U-shaped short term ($\tau = 1M$) forward smile**?
- \implies We need

$$\sqrt{\eta(t, S, X)} \equiv \frac{\sigma(t, S, X)}{\sqrt{\mathbb{E}^{\mathbb{Q}}[\sigma(t, S_t, X_t)^2 | S_t = S]}}$$

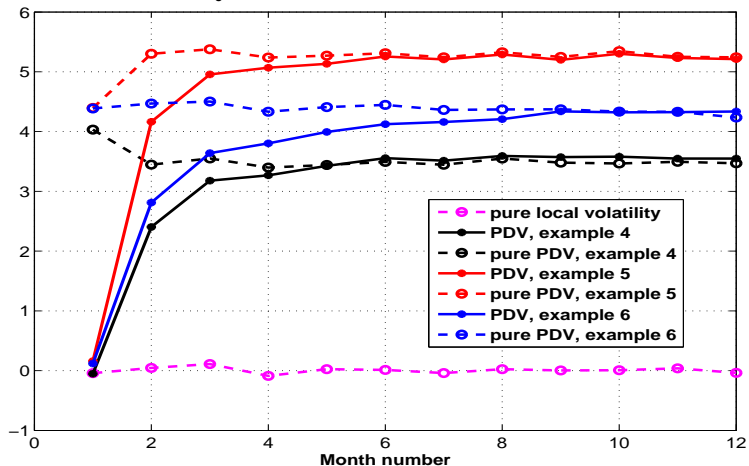
to be highly volatile and uncorrelated with S

- Examples 1–3 cannot capture this:
 - $\Delta \ll \tau \implies$ ergodic vol \implies flat forward smile
 - $\Delta \approx \tau \implies \sqrt{\eta(t, S, X)}$ correlated with S
 - $\Delta \gg \tau \implies \sqrt{\eta(t, S, X)}$ almost constant
- Examples 4–6 are natural candidates: **vol is large if and only if recent asset returns (up or down) are as well**. Produce vanishing ATM forward skew


Ex.	X_t	$\sigma(S, X)$ producing U-shaped forward smile
4	$S_{t-\Delta}$	$\bar{\sigma} 1_{\left\{ \left \frac{S}{X} - 1 \right > \kappa \sigma_0 \sqrt{\Delta} \right\}} + \underline{\sigma} 1_{\left\{ \left \frac{S}{X} - 1 \right \leq \kappa \sigma_0 \sqrt{\Delta} \right\}}$
5	\bar{S}_t^Δ	as above
6	(m_t^Δ, M_t^Δ)	$\bar{\sigma} 1_{\left\{ \frac{M}{m} - 1 > \kappa \sigma_0 \sqrt{\Delta} \right\}} + \underline{\sigma} 1_{\left\{ \frac{M}{m} - 1 \leq \kappa \sigma_0 \sqrt{\Delta} \right\}}$

Forward starting 1M butterfly spread 95%-100%-105%

Price in vol points of forward starting one month butterfly spread 95%-100%-105%
 low vol = 10/9/8%, high vol = 50/45/40%, kappa = 1/0.35/1.2,
 $\sigma_{\text{a}_0} = 20\%$, Delta = 1 month, smile is flat at 20%




Smiles of pure PDV models



Example 6: Leverage function $l(t, S)$

Leverage function $l(t, S)$: Example 6, low vol = 8%, high vol = 40%
Delta = 1 month, sigma0 = 20%, kappa = 1.2, smile is flat at 20%



Spot-vol dynamics beyond what SLV models can attain

- PDV models have so many degrees of freedom—the path-dependent variables X , and the function $\sigma(t, S, X)$ —that they can generate **spot-vol dynamics that are not attainable using SLV models**
- Example: imagine that a sophisticated client asks a quote on the conditional variance swap with payoff

$$H_T = \sum_{i=1}^{n-1} r_{i+1}^2 1_{\{r_i \leq 0\}} \approx \int_0^T \sigma_t^2 1_{\left\{\frac{S_t}{S_{t-\Delta}} \leq 1\right\}} dt, \quad r_i = \frac{S_{t_i} - S_{t_{i-1}}}{S_{t_{i-1}}}, \quad \Delta = t_i - t_{i-1} = 1 \text{ day}$$

- SLV model: for a given risk-neutral probability \mathbb{Q} ,

$$\mathbb{E}^{\mathbb{Q}} \left[\sigma_t^2 1_{\left\{\frac{S_t}{S_{t-\Delta}} \leq 1\right\}} \middle| S_t \right] \approx \mathbb{E}^{\mathbb{Q}} [\sigma_t^2 | S_t] \mathbb{E}^{\mathbb{Q}} \left[1_{\left\{\frac{S_t}{S_{t-\Delta}} \leq 1\right\}} \middle| S_t \right] \approx \frac{1}{2} \sigma_{\text{Dup}}^2(t, S_t)$$

\implies Both the SLV price and the LV price are very close to the variance swap price halved:

$$\text{SLV price} \approx \text{LV price} \approx \frac{1}{2} \int_0^T \mathbb{E}^{\mathbb{Q}} [\sigma_{\text{Dup}}^2(t, S_t)] dt = \frac{1}{2} \text{var swap price}$$

whatever the choice of \mathbb{Q}

Spot-vol dynamics beyond what SLV models can attain

- However, if client requests such quote, it is probably because they observed that for this asset r_{i+1}^2 is large when $r_i \leq 0$, and small otherwise, expect this to continue, and try to statistically arbitrage a counterparty
- **All the models** commonly used in the industry today would fail to capture this risk but the PDV model of Ex. 1, with $\Delta = t_i - t_{i-1} = 1$ day, grasps it very well
- Δ small $\implies \mathbb{E}^{\mathbb{Q}} [\sigma(t_i, S_{t_i}, S_{t_{i-1}})^2 | S_{t_i}] \approx \frac{\bar{\sigma}^2 + \underline{\sigma}^2}{2}$ is almost constant \implies

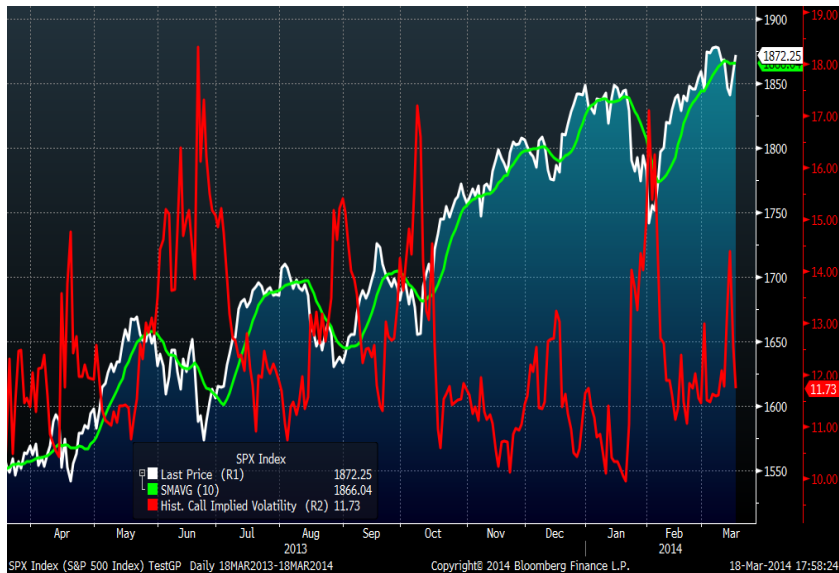
$$\text{PDV price} \approx \int_0^T \mathbb{E}^{\mathbb{Q}} \left[\frac{\bar{\sigma}^2}{\frac{\bar{\sigma}^2 + \underline{\sigma}^2}{2}} \sigma_{\text{Dup}}^2(t, S_t) 1_{\left\{ \frac{S_t}{S_{t-\Delta}} \leq 1 \right\}} \right] dt \approx \frac{\bar{\sigma}^2}{\bar{\sigma}^2 + \underline{\sigma}^2} \text{var swap price}$$

- For reasonable values of $(\underline{\sigma}, \bar{\sigma})$, e.g., (10%, 40%), this is close to the (unconditional) var swap price = **twice the SLV price and twice the LV price**. In such a case, an investment bank equipped with PDV may avoid a large mispricing

Choose a particular PDV
to capture historical patterns of volatility

Choose a particular PDV to capture historical patterns of volatility

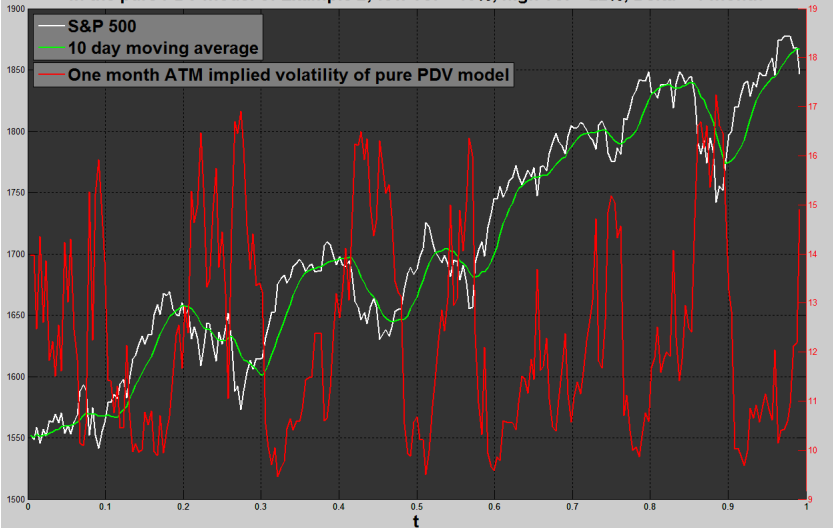
- Like the local correlation models presented in Guyon (2013), PDV models are flexible enough to **reconcile implied calibration** (e.g., calibration to the market smile) **with historical calibration** (calibration from historical time series of asset prices):
 - 1 one chooses a PDV $\sigma(t, S, X)$ from the observation of the time series, e.g., the short term ATM volatility is a certain function of S_t/\bar{S}_t^Δ
 - 2 one then multiplies it by a leverage function l and eventually calibrates l to the market smile using the particle method
- By construction, PDV models are flexible enough to capture **any** path-dependency of the volatility. For a given choice of PDV, what remains to be numerically checked is
 - 1 how much and how long the smile calibration distorts the link between past prices and current instantaneous volatility
 - 2 whether the model produces suitable dynamics of implied volatility



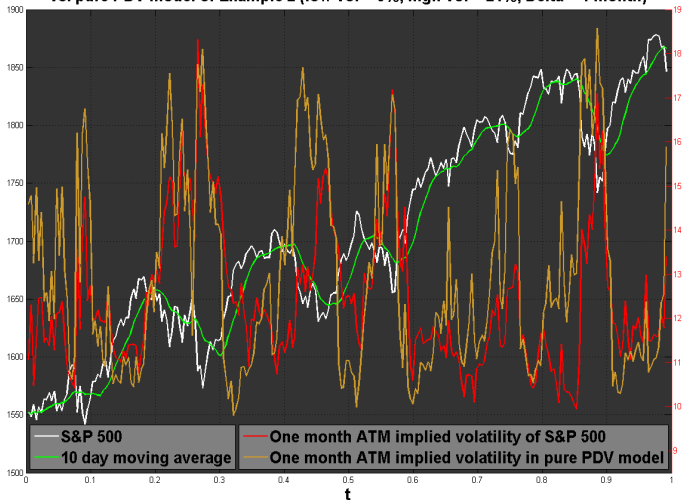
Choose a particular PDV to capture historical patterns of volatility

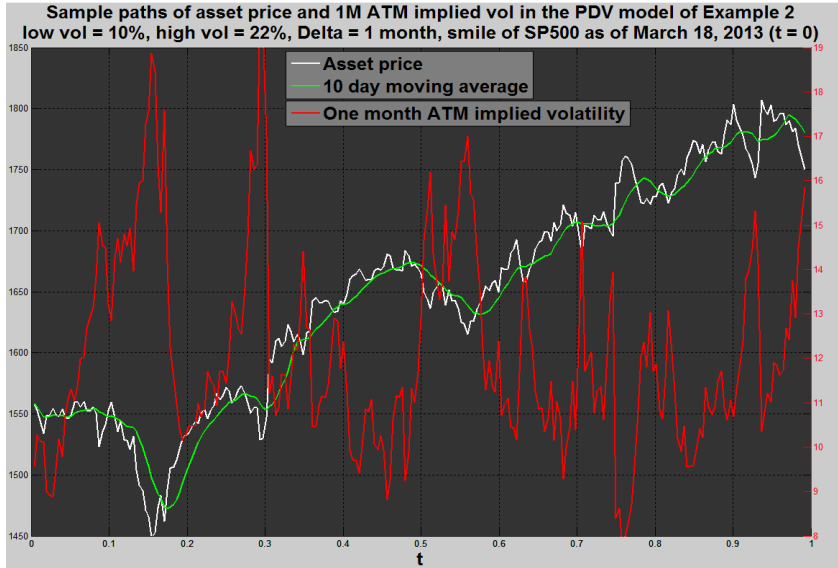
- For the S&P 500, the volatility level is **not** determined by the asset price level, but by the recent **changes** in the asset price
- Examples 1–3, which relate volatility **levels** to recent asset price **returns**, easily capture this
- Actually, the two basic quantities that possess a natural scale are the volatility **levels** and the asset **returns** so we believe that a good model should relate these two quantities
- LV model links the volatility **level** to the asset **level**, does not make much financial sense: well designed PDV models need not be recalibrated as often as the LV model
- SV models connect the **change** in volatility to the **relative change** in the asset price. Has limitations:
 - Only unreasonable levels of vol of vol allow large movements (e.g., a 70% return in 2 weeks) of instantaneous volatility
 - Therefore a large mean-reversion needs to be artificially added to keep volatility within its natural range
- By contrast PDV models can easily capture such large changes in volatility

S&P500 from March 18, 2013 to March 18, 2014 and corresponding path of the 1M ATM implied vol in the pure PDV model of Example 2; low vol = 10%, high vol = 22%, Delta = 1 month




S&P 500 from March 18, 2013 to March 18, 2014; 1M ATM implied vol: actual vs. pure PDV model of Example 2 (low vol = 8%, high vol = 21%, Delta = 1 month)

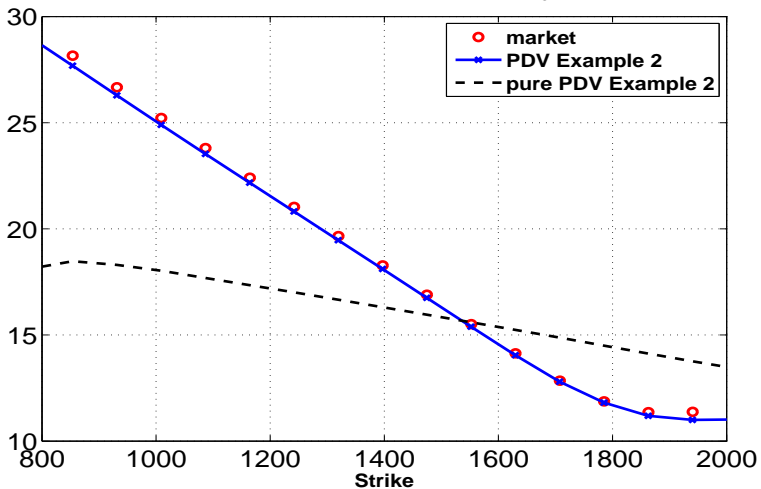




Leverage function $l(t,S)$: Example 2
low vol = 10%, high vol = 22%, Delta = 1 month
smile of SP500 as of March 18, 2013, $S_0 = 1552$



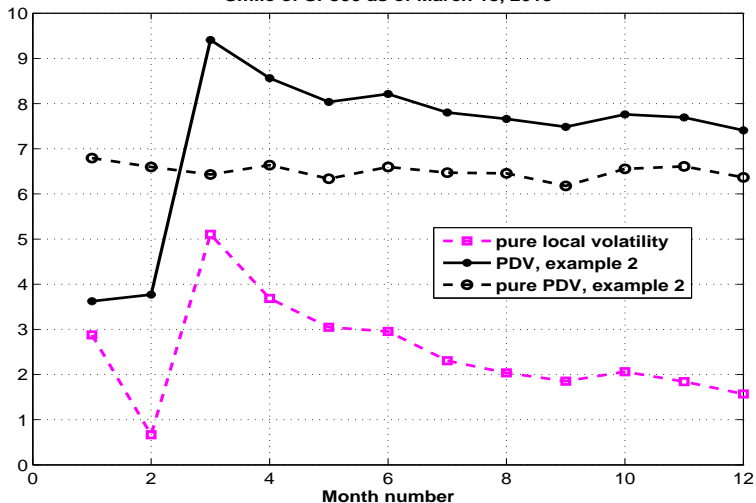
Implied volatility at maturity ($T = 1$): Example 2
low vol = 10%, high vol = 22%, Delta = 1 month
Smile of SP500 as of March 18, 2013, $S_0 = 1552$



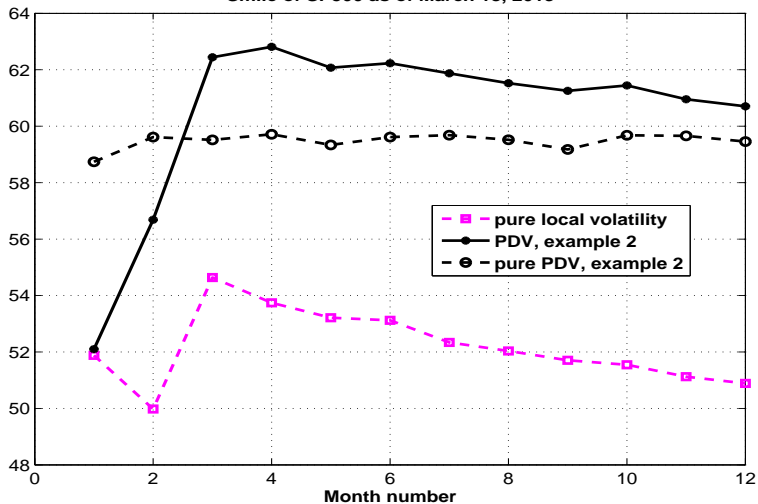
Example 2: $\underline{\sigma} = 10\%$, $\bar{\sigma} = 22\%$, $\Delta = 1$ month

- The implied volatility varies continuously (with spikes when the market is locally bearish) \implies **Modeling instantaneous volatility as a pure jump process is not problematic**: no one has ever seen such quantity—it may actually not exist
- With such parameter values, the PDV model of Example 2 captures what we believe is a **major pattern of the historical joint behaviour of the S&P 500 and its short term implied volatilities**
- What about pricing? As the volatility interval $[\underline{\sigma}, \bar{\sigma}]$ is not as wide as $[8\%, 32\%]$, the forward skew is not as expensive. However, still sizeably larger than the LV forward skew

Price in vol points of forward starting one month call spread 95%–105%
 Example 2: low vol = 10%, high vol = 22%, Delta = 1 month
 Smile of SP500 as of March 18, 2013



Price in % of forward starting one month digital ATM call
Example 2: low vol = 10%, high vol = 22%, Delta = 1 month
Smile of SP500 as of March 18, 2013



Generalized local ARCH/GARCH models

- That volatility depends on recent asset returns was also supported by other statistical analyses (Platania-Rogers, 2003; Foschi-Pascucci, 2007)
- Some empirical studies show that vol may depend on recent realized volatility. So far, only the ARCH (Engle, 1982) and GARCH (Bollerslev, 1986) models, and their descendants, could capture this
- ARCH/GARCH capture tail heaviness, volatility clustering and dependence without correlation, like Examples 1–6 above
- Our approach generalizes them by defining **local ARCH models**, in which the ARCH volatility is multiplied by a leverage function in order to fit a smile and the function $\sigma(X)$ is arbitrary:

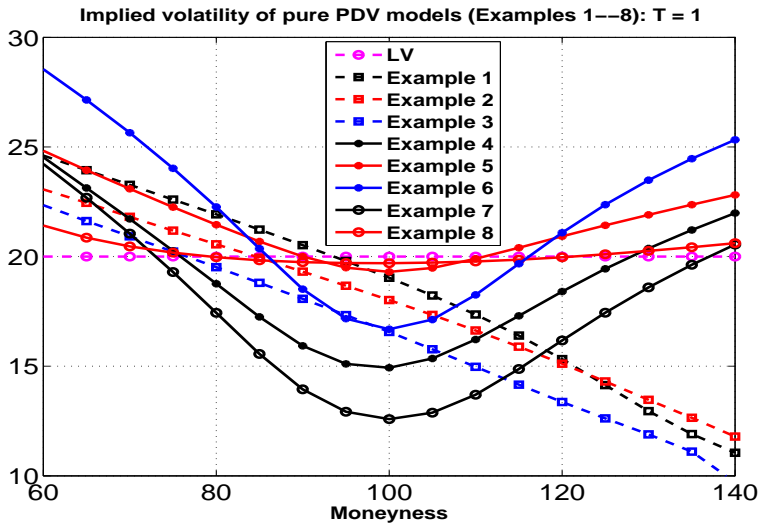
$$\frac{dS_t}{S_t} = \sigma(X_t)l(t, S_t) dW_t, \quad X_t = \sum_{t-\Delta < t_i \leq t} r_i^2, \quad r_i = \frac{S_{t_i} - S_{t_{i-1}}}{S_{t_{i-1}}}$$

Generalized local ARCH/GARCH models

$$\frac{dS_t}{S_t} = \sigma(X_t)l(t, S_t) dW_t, \quad X_t = \sum_{t-\Delta < t_i \leq t} r_i^2, \quad r_i = \frac{S_{t_i} - S_{t_{i-1}}}{S_{t_{i-1}}}$$


- Ex. 7: $\sigma(t, X) = \underline{\sigma}$ if $X \leq \sigma_0$ and $\bar{\sigma}$ otherwise. Vanishing ATM forward skew. Forward starting butterfly spreads cost around 2.4 points of volatility
- Ex. 8 (to mimic ARCH models): $\sigma(X)^2 = \alpha + \beta X$ with $\alpha > 0$, $\beta < 1$. Much flatter pure PDV smile and a much flatter leverage function l . Vanishing ATM forward skew. Forward starting butterfly spreads around 0.7 point of volatility

Smiles of pure PDV models




Example 7: Leverage function $l(t, S)$

Leverage function $l(t, S)$: Example 7, low vol = 10%, high vol = 40%
Delta = 1 week, $\sigma_0 = 20\%$, smile is flat at 20%



Example 8: Leverage function $l(t, S)$

Leverage function $l(t, S)$: Example 8
Delta = 1 week, alpha = 0.008, beta = 0.8, smile is flat at 20%



Conclusion

- PDV models are excellent candidates to **challenge the duopoly of LV and SV** which has dominated option pricing for 20 years
- Like the LV model: complete and can be calibrated to the market smile
⇒ **all derivatives have a unique price which is consistent with today's prices of vanilla options**
- Like SV models: can produce **rich spot-vol dynamics**, such as large negative short term forward skews or large forward smile curvatures
- Huge flexibility: one can choose any set of path-dependent variables X and any PDV $\sigma(t, S, X) \implies$ PDV models actually span a **much broader range of spot-vol dynamics** than SV models, and can also **capture important historical features of asset returns**, such as volatility *levels* depending on recent asset *returns*, tail heaviness, volatility clustering, and dependence without correlation

Conclusion

- All SV/LSV models can be mimicked by PDV models: There exists a PDV model that has the same distribution **as a process** – **ALL options** (not only vanilla options) **have the same price in both models!**
- In practice, the **particle method** is so **simple and efficient** that the smile calibration is not a problem \implies **Efforts can be concentrated on the choice of a convenient PDV**, depending on the market and derivative under consideration

Risk magazine, October 2014



Cutting edge: Derivatives pricing

Path-dependent volatility

So far, path-dependent volatility models have drawn little attention compared with local volatility and stochastic volatility models. In this article, Julien Guyon shows they combine benefits from both and can also capture prominent historical patterns of volatility

Three main volatility models have been used so far in the finance industry: constant volatility, local volatility (LV) and stochastic volatility (SV). The first two models are complete: since the asset price is driven by a single Brownian motion, every payoff admits a unique self-financing replicating portfolio consisting of cash and the underlying asset. Therefore, its price is uniquely defined as the initial value of the replicating portfolio, independent of utilities or preferences. Unlike the constant volatility models, the LV model is flexible enough to fit any arbitrage-free surface of implied volatilities (henceforth, 'smile'), but then no more flexibility is left. Calibrating to the market smile is useful when one sells an exotic option whose risk is well mitigated by trading vanilla options – then the model correctly prices the hedging instruments at inception.

For their part, SV models are incomplete: the volatility is driven by one of several extra Brownian motions, and as a result perfect replication and price uniqueness are lost. Modifying the drift of the SV leaves the model arbitrage-free, but changes option prices.

Using SV models allows us to gain control of key risk factors such as volatility of volatility (vol-of-vol), forward skew and spot-vol correlation. SV models generate joint dynamics of the asset and its implied

price uniqueness and parsimony: it is remarkable that so many popular properties of SLV models can be captured using a single Brownian motion. Although perfect delta-hedging is unrealistic, incorporating the path-dependency of volatility into the delta is likely to improve the delta-hedge. Not only that, we will see that, thanks to their huge flexibility, PDV models can generate spot-vol dynamics that are not attainable using SLV models.

Below, we first introduce the class of PDV models and then explain how we calibrate them to the market smile. Subsequently, we investigate how to pick a particular PDV.









Path-dependent volatility models

PDV models are those models where the instantaneous volatility σ_t depends on the path followed by the asset price so far:








$$\frac{dS_t}{S_t} = \sigma(t, (S_u, u \leq t)) dW_t$$

where, for simplicity, we have taken zero interest rates, repo and dividends. In practice, the volatility $\sigma_t = \sigma(t, S_t, X_t)$ will often be

A few selected references

-  Bergomi, L.: *Smile dynamics II*, Risk, October, 2005.
-  Bollerslev, T.: *Generalized Autoregressive Conditional Heteroskedasticity*, Journal of Econometrics, 31:307-327, 1986.
-  Brunick, G. and Shreve, S.: *Mimicking an Itô process by a solution of a stochastic differential equation*, Ann. Appl. Prob., 23(4):1584–1628, 2013.
-  Dupire, B.: *Pricing with a smile*, Risk, January, 1994.
-  Engle, R.: *Autoregressive Conditional Heteroscedasticity with Estimates of Variance of United Kingdom Inflation*, Econometrica 50:987-1008, 1982.
-  Figà-Talamanca, G. and Guerra, M.L.: *Fitting prices with a complete model*, J. Bank. Finance 30(1), 247?258, 2006.
-  Foschi, P, and Pascucci, A.: *Path Dependent Volatility*, Decisions Econ. Finan., 2007.
-  Fouque, J.-P., Papanicolaou, G. and Sircar, R.: *Derivatives in Financial Markets with Stochastic Volatility*, Cambridge University Press, 2000.

A few selected references

-  Guyon, J. and Henry-Labordère, P.: *Being Particular About Calibration*, Risk, January, 2012.
-  Guyon, J. and Henry-Labordère, P.: *Nonlinear Option Pricing*, Chapman & Hall/CRC Financial Mathematics Series, 2013.
-  Guyon, J.: *Local correlation families*, Risk, February, 2014.
-  Hobson, D. G. and Rogers, L. C. G.: *Complete models with stochastic volatility*, Mathematical Finance 8 (1), 27-48, 1998.
-  Hubalek, F., Teichmann, J. and Tompkins, R.: *Flexible complete models with stochastic volatility generalising Hobson-Rogers*, working paper, 2004.
-  Klüppelberg, C., Lindner, A. and Maller, R.: *A continuous time GARCH process driven by a Lévy process: stationarity and second order behaviour*, J. Appl. Probab., 41(3):601-622, 2004.
-  Platania, A. and Rogers, L.C.G.: *Putting the Hobson-Rogers model to the test*, working paper, 2003

