

Volatility and Hedging Errors

Jim Gatheral

September, 25 1999

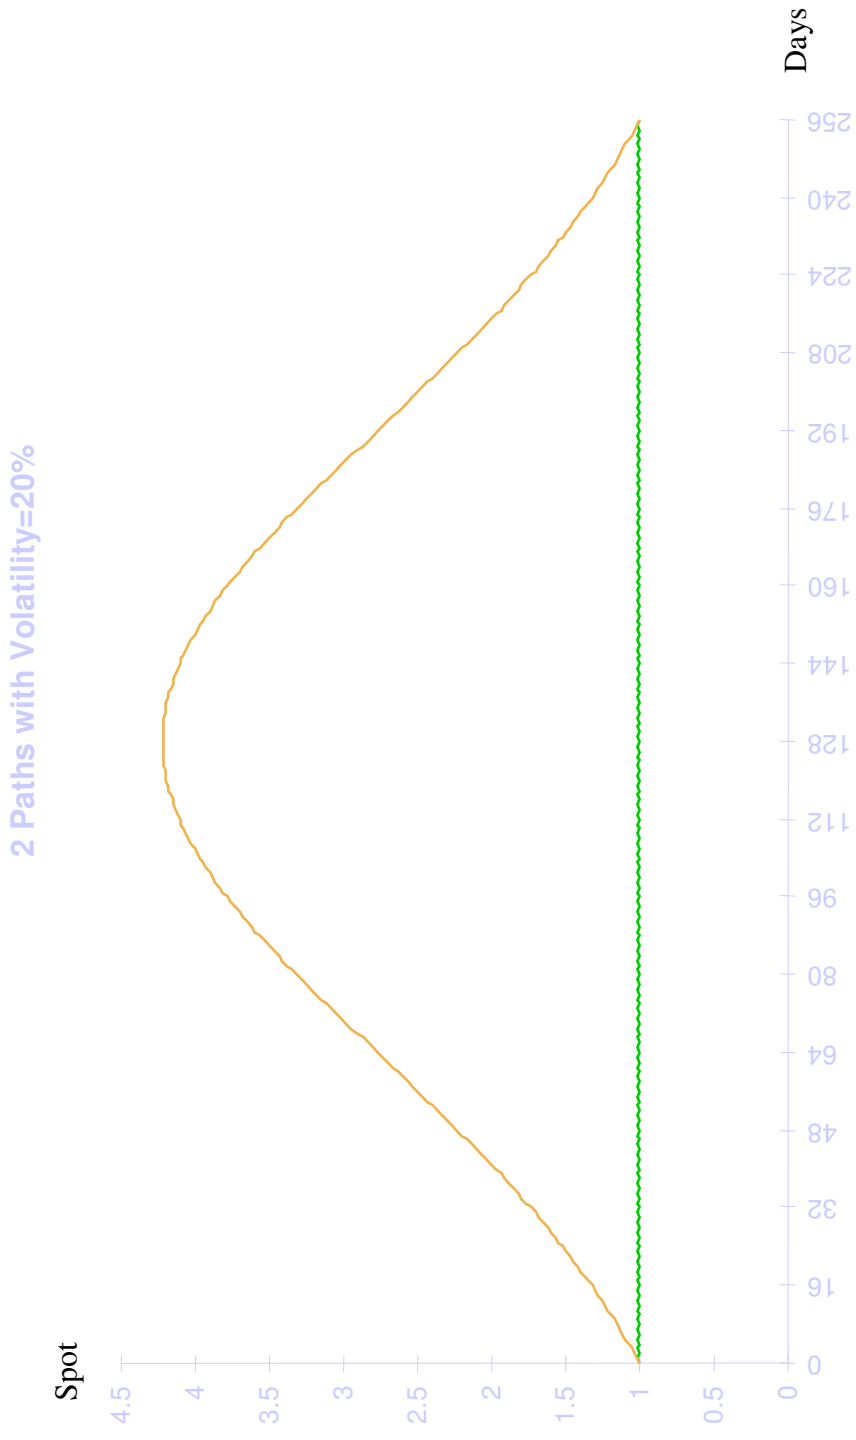
Background

- Derivative portfolio bookrunners often complain that
 - hedging at market-implied volatilities is sub-optimal relative to hedging at their best guess of future volatility but
 - they are forced into hedging at market implied volatilities to minimise mark-to-market P&L volatility.
- All practitioners recognise that the assumptions behind fixed or swimming delta choices are wrong in some sense. Nevertheless, the magnitude of the impact of this delta choice may be surprising.
- Given the P&L impact of these choices, it would be nice to be able to avoid figuring out how to delta hedge. Is there a way of avoiding the problem?

An Idealised Model


- To get intuition about hedging options at the wrong volatility, we consider two particular sample paths for the stock price, both of which have realised volatility = 20%
 - a whipsaw path where the stock price moves up and down by 1.25% every day
 - a sine curve designed to mimic a trending market

Whipsaw and Sine Curve Scenarios



Whipsaw vs Sine Curve: Results

P&L vs Hedge Vol.



Conclusions from this Experiment

- If you knew the realised volatility in advance, you would definitely hedge at that volatility because the hedging error at that volatility would be zero.
- In practice of course, you don't know what the realised volatility will be. The performance of your hedge depends not only on whether the realised volatility is higher or lower than your estimate but also on whether the market is range bound or trending.

Analysis of the P&L Graph

- If the market is range bound, hedging a short option position at a lower vol. hurts because you are getting continuously whipsawed. On the other hand, if you hedge at very high vol., and market is range bound, your gamma is very low and your hedging losses are minimised.
- If the market is trending, you are hurt if you hedge at a higher vol. because your hedge reacts too slowly to the trend. If you hedge at low vol., the hedge ratio gets higher faster as you go in the money minimising hedging losses.


Another Simple Hedging Experiment

- In order to study the effect of changing hedge volatility, we consider the following simple portfolio:
 - short \$1bn notional of 1 year ATM European calls
 - long a one year volatility swap to cancel the vega of the calls at inception.
- This is (almost) equivalent to having sold a one year option whose price is determined *ex-post* based on the actual volatility realised over the hedging period.
- Any P&L generated by this hedging strategy is pure hedging error. That is, we eliminate any P&L due to volatility movements.

Historical Sample Paths


- In order to preempt criticism that our sample paths are too unrealistic, we take real historical FTSE data from two distinct historical periods: one where the market was locked in a trend and one where the market was range bound.
 - For the range bound scenario, we consider the period from April 1991 to April 1992
 - For the trending scenario, we consider the period from October 1996 to October 1997
 - In both scenarios, the realised volatility was around 12%

FTSE 100 since 1985



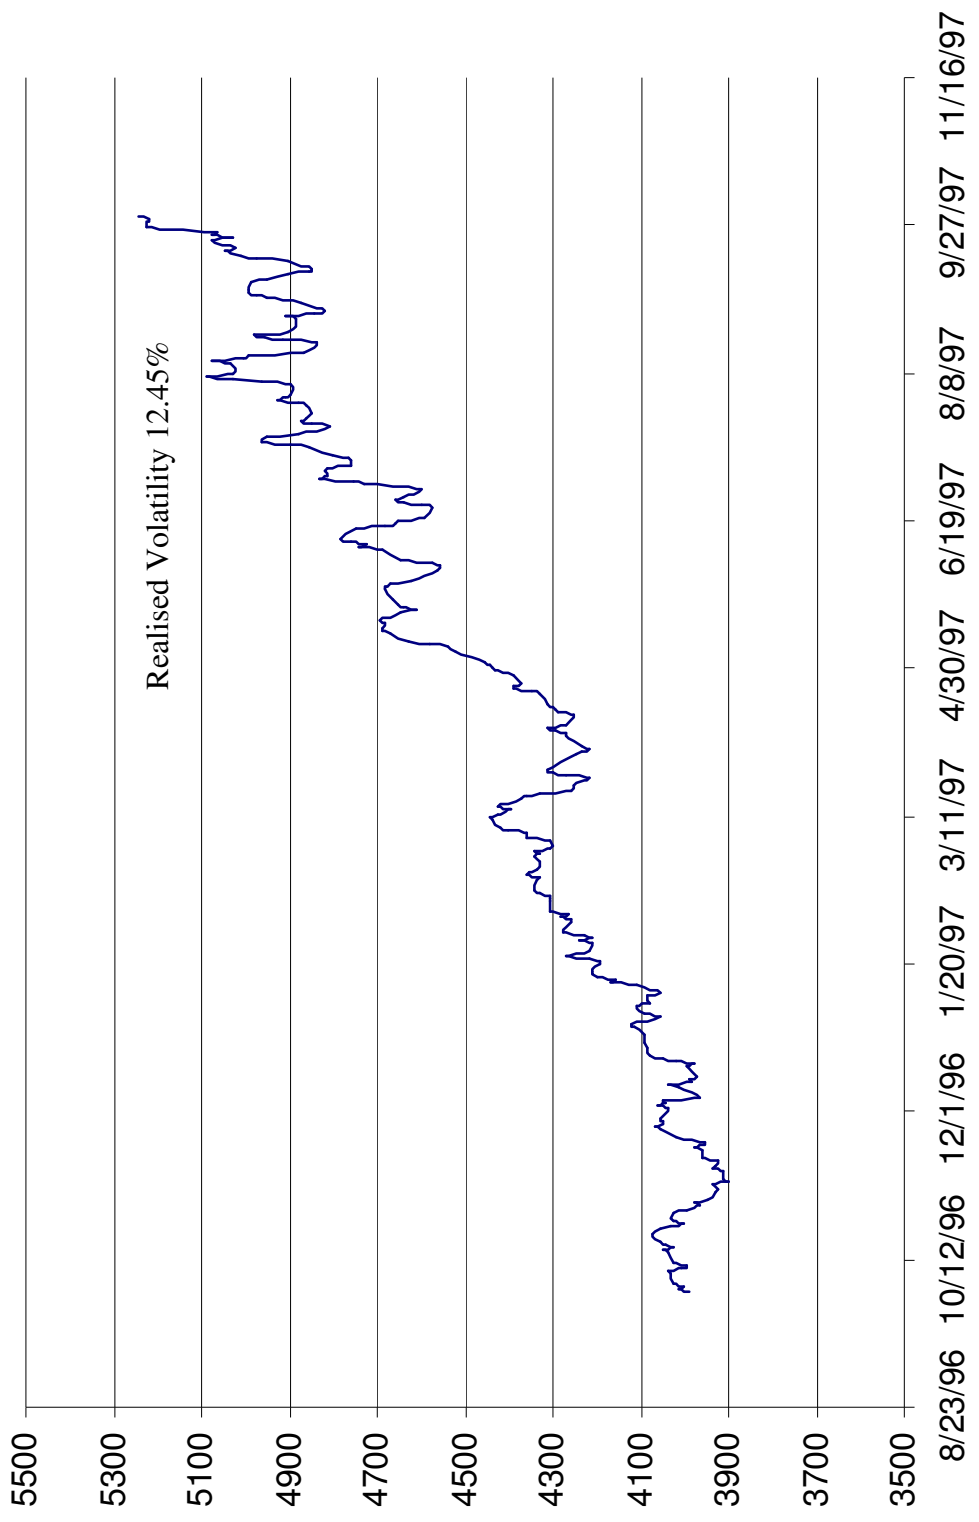
Range Scenario

FTSE from 4/1/91 to 3/31/92

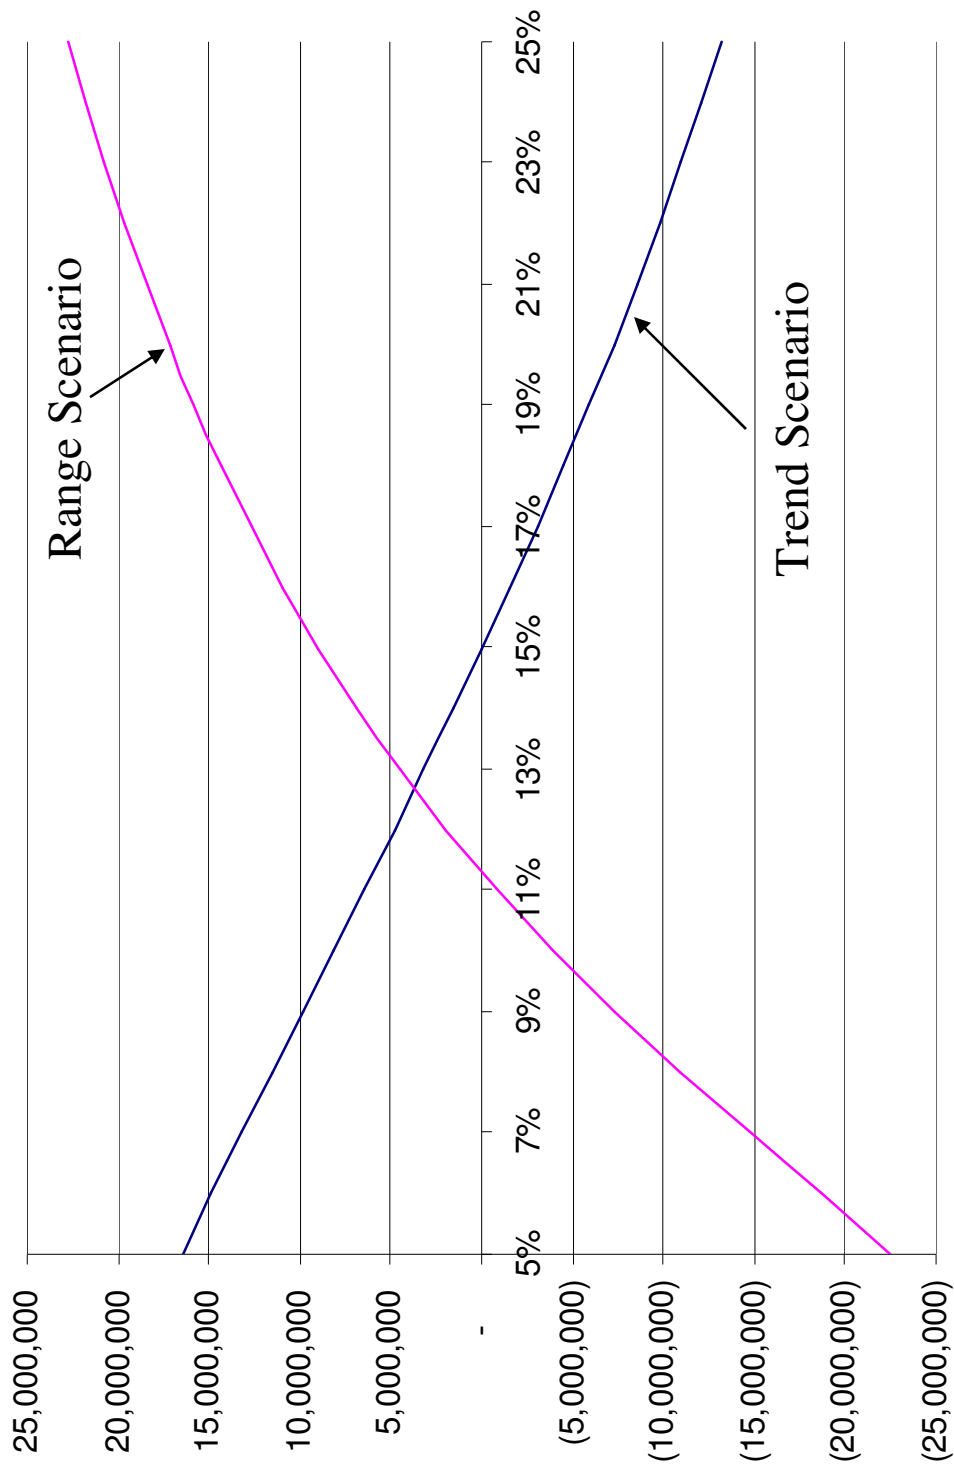


Trend Scenario

FTSE from 11/1/96 to 10/31/97



P&L vs Hedge Volatility



Discussion of P&L Sensitivities

- The sensitivity of the P&L to hedge volatility did depend on the scenario just as we would have expected from the idealised experiment.
 - In the range scenario, the lower the hedge volatility, the lower the P&L consistent with the whipsaw case.
 - In the trend scenario, the lower the hedge volatility, the higher the P&L consistent with the sine curve case.
- In each scenario, the sensitivity of the P&L to hedging at a volatility which was wrong by 10 volatility points was around \$20mm for a \$1bn position.


Questions?

- Suppose you sell an option at a volatility higher than 12% and hedge at some other volatility. If realised volatility is 12%, do you make money?
 - Not necessarily. It is easy to find scenarios where you lose money.
- Suppose you sell an option at some implied volatility and hedge at the same volatility. If realised volatility is 12%, when do you make money?
 - In the two scenarios analysed, if the option is sold and hedged at a volatility greater than the realised volatility, the trade makes money. This conforms to traders' intuition.
 - Later, we will show that even this is not always true.

Sale/ Hedge Volatility Combinations

	<i>Sale Vol.</i>	<i>Volatility P&L</i>	<i>Hedge Vol.</i>	<i>Hedge P&L</i>	<i>Total</i>
Range Scenario	13%	+\$3.30mm	10%	-\$3.85mm	-\$0.55mm
Trend Scenario	13%	+\$2.20mm	17%	-\$3.07mm	-\$0.87mm


P&L from Selling and Hedging at the Same Volatility



Delta Sensitivities

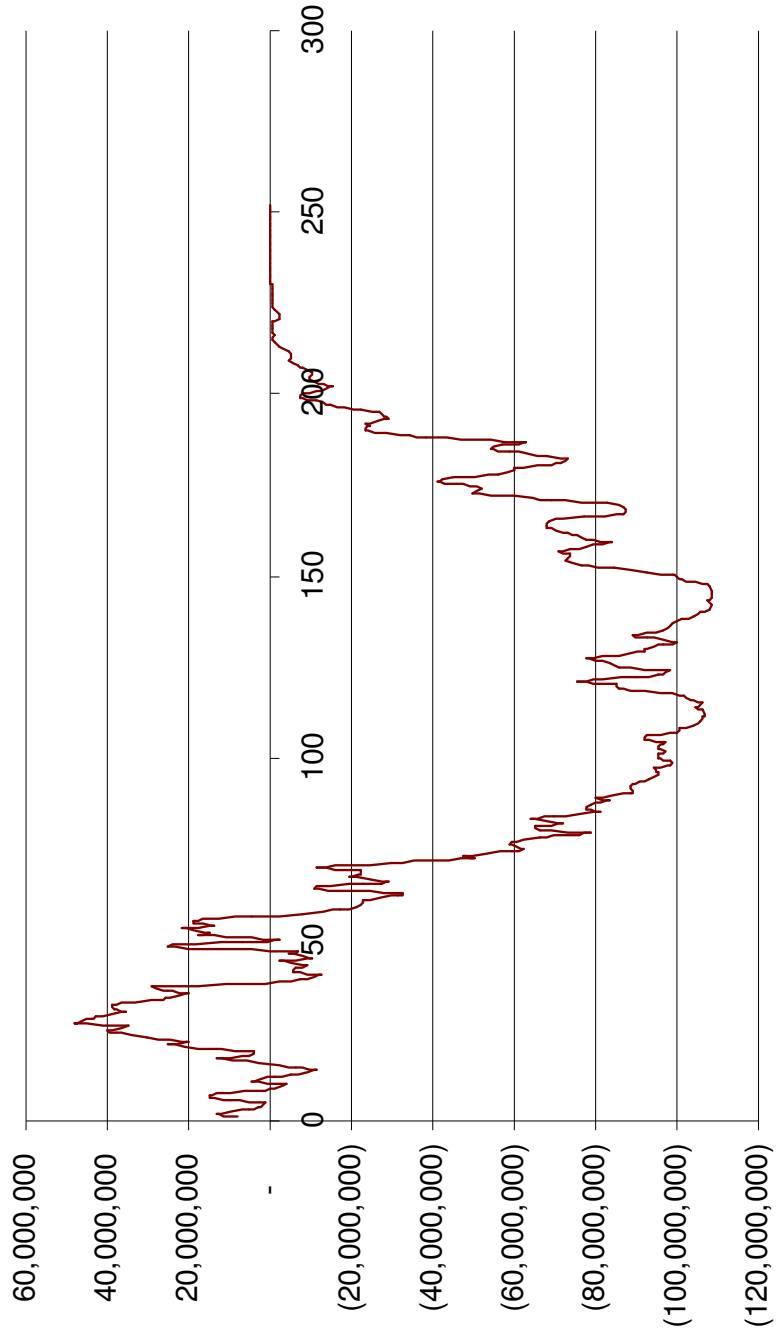
- Let's now see what effect hedging at the wrong volatility has on the delta.
 - We look at the difference between Δ -delta computed at 20% volatility and Δ -delta computed at 12% volatility as a function of time.
- In the range scenario, the difference between the deltas persists throughout the hedging period because both gamma and vega remain significant throughout.
- On the other hand, in the trend scenario, as gamma and vega decrease, the difference between the deltas also decreases.

Range Scenario



Trend Scenario

\$ Delta Difference (20% vol. - 12% vol.)




Fixed and Swimming Delta

- Fixed (sticky strike) delta assumes that the Black-Scholes implied volatility for a particular strike and expiration is constant. Then
- Swimming (or floating) delta assumes that the at-the-money Black-Scholes implied volatility is constant. More precisely, we assume that implied volatility is a function of relative strike K/S only. Then


$$\delta_{FIXED} = \frac{\partial C_{BS}}{\partial S}$$

$$\delta_{SWIM} = \frac{\partial C_{BS}}{\partial S} + \frac{\partial C_{BS}}{\partial \sigma_{BS}} \frac{\partial \sigma_{BS}}{\partial S} = \frac{\partial C_{BS}}{\partial S} - \frac{K}{S} \frac{\partial C_{BS}}{\partial \sigma_{BS}} \frac{\partial \sigma_{BS}}{\partial K}$$

An Aside: The Volatility Skew



Volatility vs x




Observations on the Volatility Skew

- Note how beautiful the raw data looks; there is a very well-defined pattern of implied volatilities.
- When implied volatility is plotted against $\frac{\ln(K/F)}{\sqrt{T}}$, all of the skew curves have roughly the same shape.


How Big are the Delta Differences?

- We assume a skew of the form
$$\frac{\partial \sigma_{BS}}{\partial x} = -\frac{0.3}{\sqrt{T}}$$
- From the following two graphs, we see that the typical difference in delta between fixed and swimming assumptions is around \$100mm. The error in hedge volatility would need to be around 8 points to give rise to a similar difference.
- In the range scenario, the difference between the deltas persists throughout the hedging period because both gamma and vega remain significant throughout.
- On the other hand, in the trend scenario, as gamma and vega decrease, the difference between the deltas also decreases.

Range Scenario



Trend Scenario



Summary of Empirical Results

- Delta hedging always gives rise to hedging errors because we cannot predict realised volatility.
- The result of hedging at too high or too low a volatility depends on the precise path followed by the underlying price.
- The effect of hedging at the wrong volatility is of the same order of magnitude as the effect of hedging using swimming rather than fixed delta.
- Figuring out which delta to use at least as important than guessing future volatility correctly and probably more important!

Some Theory

- Consider a European call option struck at K expiring at time T and denote the value of this option at time t according to the Black-Scholes formula by $C_{BS}(t)$. In particular, $C_{BS}(T) = (S_T - K)^+$.
- We assume that the stock price S satisfies a SDE of the form
$$\frac{dS_t}{S_t} = \mu_t dt + \sigma_{t,S_t} dZ$$
 where σ_{t,S_t} may itself be stochastic.
- Path-by-path, we have

$$\begin{aligned}
C_{BS}(T) - C_{BS}(0) &= \int_0^T dC_{BS} \\
&= \int_0^T \left\{ \frac{\partial C_{BS}}{\partial S_t} dS_t + \frac{\partial C_{BS}}{\partial t} dt + \frac{\sigma_{S_t}^2 S_t^2}{2} \frac{\partial^2 C_{BS}}{\partial S_t^2} dt \right\} \\
&= \int_0^T \left\{ \frac{\partial C_{BS}}{\partial S_t} dS_t + \frac{1}{2} (\sigma_{S_t}^2 - \hat{\sigma}_t^2) S_t^2 \frac{\partial^2 C_{BS}}{\partial S_t^2} dt \right\}
\end{aligned}$$


where the forward variance $\hat{\sigma}_t^2 = \frac{\partial}{\partial t} \{ \sigma_{BS}^2(t) \cdot t \}$.

So, if we delta hedge using the Black-Scholes (fixed) delta, the outcome of the hedging process is:


$$C_{BS}(T) - C_{BS}(0) = \int_0^T \frac{1}{2} (\sigma_{S_t}^2 - \hat{\sigma}_t^2) S_t^2 \frac{\partial^2 C_{BS}}{\partial S_t^2} dt$$

- In the Black-Scholes limit, with deterministic volatility, delta-hedging works path-by-path because $\sigma_{S_t}^2 = \hat{\sigma}_t^2 \forall S_t$
- In reality, we see that the outcome depends both on gamma and the difference between realised and hedge volatilities.
 - If gamma is high when volatility is low and/or gamma is low when volatility is high you will make money and *vice versa*.
- Now, we are in a position to provide a counterexample to trader intuition:
 - Consider the particular path shown in the following slide:
 - The realised volatility is 12.45% but volatility is close to zero when gamma is low and high when gamma is high.
 - The higher the hedge volatility, the lower the hedge P&L.
 - In this case, if you price and hedge a short option position at a volatility lower than 18%, you lose money.

A Cooked Scenario



Cooked Scenario: P&L from Selling and Hedging at the Same Volatility



Conclusions

- Delta-hedging is so uncertain that we must delta-hedge as little as possible and what delta-hedging we do must be optimised.
- To minimise the need to delta-hedge, we must find a static hedge that minimises gamma path-by-path. For example, Avellaneda *et al.* have derived such static hedges by penalising gamma path-by-path.
- The question of what delta is optimal to use is still open. Traders like fixed and swimming delta. Quants prefer market-implied delta - the delta obtained by assuming that the local volatility surface is fixed.

Another Digression: Local Volatility

- We assume a process of the form:

$$\frac{dS_t}{S_t} = \mu_t dt + \hat{\sigma}_{t,S_t} dZ$$

- with $\hat{\sigma}_{t,S_t}$ a deterministic function of stock price and time.
- Local volatilities $\hat{\sigma}_{T,K}$ can be computed from market prices of options C using

$$\frac{\hat{\sigma}_{T,K}^2}{2} = \frac{\frac{\partial C}{\partial T}}{\frac{\partial^2 C}{\partial K^2}}$$

- Market-implied delta assumes that the local volatility surface stays fixed through time.

We can extend the previous analysis to local volatility.

$$\begin{aligned}
C_L(T) - C_L(0) &= \int_0^T dC_L \\
&= \int_0^T \left\{ \frac{\partial C_L}{\partial S_t} dS_t + \frac{\partial C_L}{\partial t} dt + \frac{\sigma_{t,S_t}^2 S_t^2}{2} \frac{\partial^2 C_L}{\partial S_t^2} dt \right\} \\
&= \int_0^T \left\{ \frac{\partial C_L}{\partial S_t} dS_t + \frac{1}{2} (\sigma_{t,S_t}^2 - \hat{\sigma}_{t,S_t}^2) S_t^2 \frac{\partial^2 C_L}{\partial S_t^2} dt \right\}
\end{aligned}$$

So, if we delta hedge using the market-implied delta $\frac{\partial C_L}{\partial S_t}$, the outcome of the hedging process is:

$$C_L(T) - C_L(0) = \int_0^T \frac{1}{2} (\sigma_{t,S_t}^2 - \hat{\sigma}_{t,S_t}^2) S_t^2 \frac{\partial^2 C_L}{\partial S_t^2} dt$$

Define $\Gamma_L(S_t) \equiv S_t^2 \frac{\partial^2 C_L}{\partial S_t^2}$

$$\text{Then } C_L(T) - C_L(0) = \int_0^T \frac{1}{2} (\sigma_{t,S_t}^2 - \hat{\sigma}_{t,S_t}^2) \Gamma_L(S_t) dt$$

If the claim being hedged is path-dependent then $\Gamma_L(S_t)$ is also path-dependent. Otherwise all the $\Gamma_L(S_t)$ can be determined at inception.

Writing the last equation out in full, for two local volatility surfaces $\tilde{\Sigma}$ and $\hat{\Sigma}$ we get

$$C_L(\tilde{\Sigma}) - C_L(\hat{\Sigma}) = \frac{1}{2} \int_0^T dt \int_0^\infty dS_t (\tilde{\sigma}_{t,S_t}^2 - \hat{\sigma}_{t,S_t}^2) E[\Gamma_L(S_t)|S_t] \varphi(S_t|S_0) dt$$

Then, the functional derivative $\frac{\delta C_L}{\delta \hat{\sigma}_{t,S_t}^2} = \frac{1}{2} E[\Gamma_L(S_t)|S_t] \varphi(S_t|S_0)$

In practice, we can set $\frac{\delta \mathcal{C}_L}{\delta \sigma_{t,S_t}} = 0$ by bucket hedging.

Note in particular that European options have all their sensitivity to local volatility in one bucket - at strike and expiration. Then by buying and selling European options, we can cancel the risk-neutral expectation of gamma over the life of the option being hedged - a static hedge. This is *not* the same as cancelling gamma path-by-path.

If you do this, you still need to choose a delta to hedge the remaining risk. In practice, whether fixed, swimming or market-implied delta is chosen, the parameters used to compute these are re-estimated daily from a new implied volatility surface.

Dumas, Fleming and Whaley point out that the local volatility surface is very unstable over time so again, it's not obvious which delta is optimal.

Outstanding Research Questions

- Is there an optimal choice of delta which depends only on observable asset prices?
- How should we price
 - Path-dependent options?
 - Forward starting options?
 - Compound options?
 - Volatility swaps?

Some References

- Avellaneda, M., and A. Parás. Managing the volatility risk of portfolios of derivative securities: the Lagrangian Uncertain Volatility Model. *Applied Mathematical Finance*, 3, 21-52 (1996)
- Blacher, G. A new approach for understanding the impact of volatility on option prices. *RISK 98 Conference Handout*.
- Derman, E. Regimes of volatility. *RISK April*, 55-59 (1999)
- Dumas, B., J. Fleming, and R.E. Whaley. Implied volatility functions: empirical tests. *The Journal of Finance Vol. LIII, No. 6*, December 1998.
- Gupta, A. On neutral ground. *RISK July*, 37-41 (1997)