

Lecture 9: Patterns of Volatility Change

- Regimes of volatility: how implied index volatilities move
- The best average hedge?


Copyright Emanuel Derman 2008

9.1 Patterns of Volatility Change

How do implied volatilities of options with definite strikes actually behave as stock price changes and time passes?

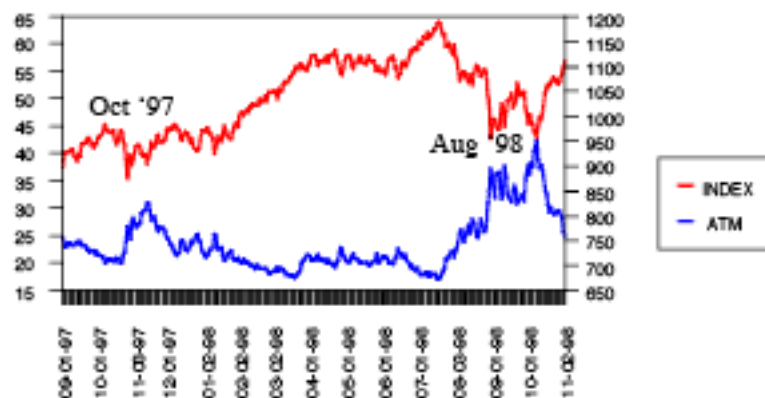
The implied volatility surface of index options - a negative skew

FIGURE 1. The mid-market S&P 500 implied volatility surface at the market close on (a) September 27, 1995 and (b) December 3, 1998.



A negative correlation between *at-the-money* volatility and index level

Three-Month Implied Volatilities of SPX Options



But you cannot buy at-the-money implied volatility, so let's look at the volatility of definite strikes, which is much more complex. What is going on? Here's some old data during periods of crisis.

Copyright Emanuel Derman 2008

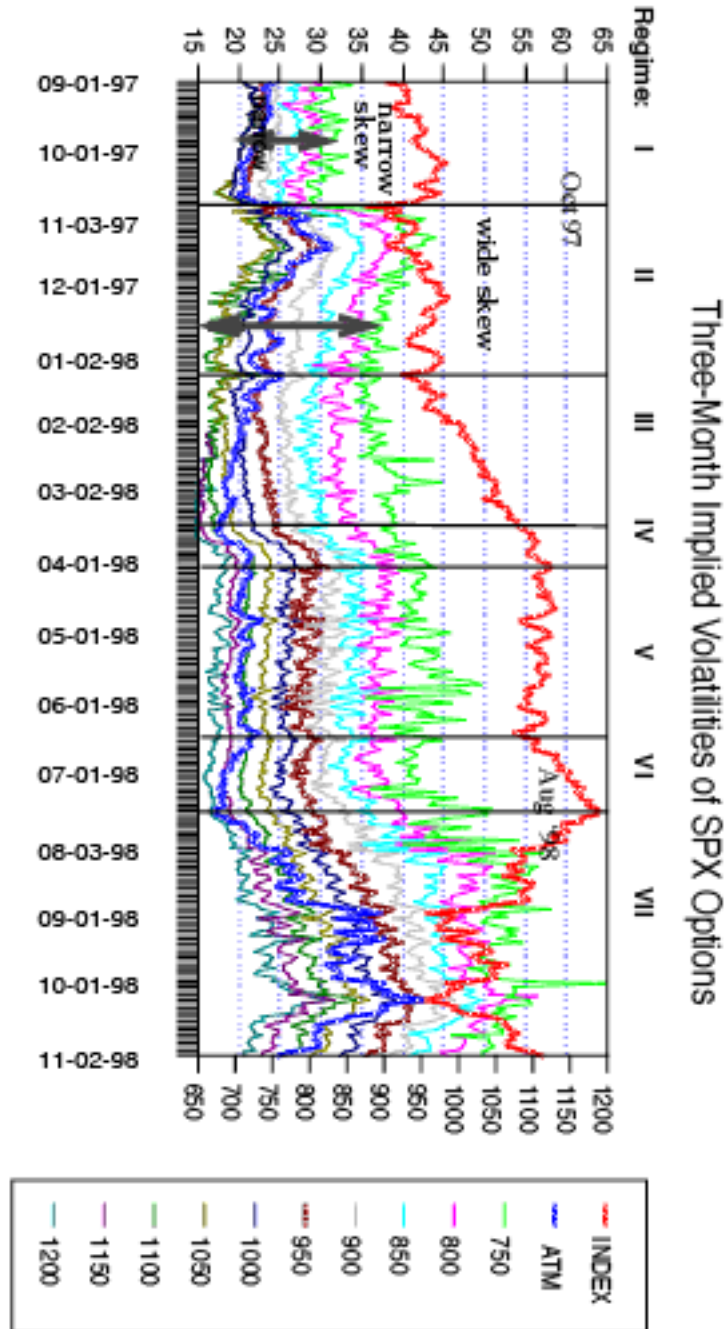


FIGURE 3. The implied volatilities of S&P 500 options from September 1997 through October 1998.

The S&P 500 implied volatility skew at any instant when the index is at level S_0 is approximately described by the linear approximation

$\Sigma(K, t_0) = \Sigma_{atm}(S_0, t_0) - b(t_0)(K - S_0)$, which is not too bad an approximation when the strike is near the money.

In what follows, we will write $\Sigma_{atm}(S_0, t_0) \equiv \Sigma_0$ and $b(t_0) \equiv b$, so that

$$\Sigma(K, t_0) = \Sigma_0 - b(K - S_0) \quad \text{Eq.9.1}$$

This simple linear dependence on strike K tells us nothing about how implied volatilities will vary when the index level moves away from S_0 ; the $(K - S_0)$ term is intended to describe only the variation in K .

Nevertheless, anyone interested in option value wants to know what will happen to the option's implied volatility when the index moves to a level S . What is the S -dependence of $\Sigma(S, t, K, T)$ given its observed K -dependence? This question is important for obtaining appropriate hedge ratios as well as options values. You can take a look at the paper *Regimes of Volatility, Risk 1999*, or on my website, to see a study of how implied volatilities actually moved in the late 1990s during periods of market crisis.

9.1.1 Rules, Models, Theories for the Variation of Implied Volatility

Given the current index implied volatility skew, one wants to know how it will vary with index level in the future. It's often easier to formulate models of change by stating what doesn't change. In physics or mathematics, such quantities are called invariants. In the world of options trading, it's become customary for traders to refer to what doesn't change as "sticky." There are at least three different views on which aspects of the current skew are sticky as the index moves over short times. These are models or rather *descriptions* of the Black-Scholes quoting parameter called implied volatility Σ rather than the underlying S and its stochastic evolution. Obviously, being descriptions, none of them should be regarded as absolutely correct, especially over a long period. Our aim will be to see to what extent these rules hold sway over the options market at different times.

The three rules are:

1. The Sticky Strike Rule.
2. The Sticky Delta Rule.
3. The Sticky Implied Tree Rule.

9.1.2 The Sticky Strike Rule

Given the current skew, one possibility is that as the index moves, each option of a definite strike maintains its initial implied volatility – hence the “sticky strike” appellation. Obviously, this cannot hold indefinitely. This is the simplest “model” of implied volatility.

Mathematically, the sticky strike rule is

$$\Sigma(S, K, t) = \Sigma_0 - b(K - S_0) \quad \text{Sticky Strike Rule} \quad \text{Eq.9.2}$$

This is equivalent to assuming that Equation 9.2 holds true for any index level S , that is, that implied volatility has no dependence on S . The value S_0 is present simply to indicate the current skew’s dependence on moneyness. Because of our assumption of stickiness, we have assumed that $b(t) = b$, independent of t . Of course, b can change, even dramatically during crisis periods.

Intuitively, “sticky strike” is a poor man’s attempt to preserve the Black-Scholes model. It allows each option an independent existence, and doesn’t worry about whether the collective options’ market view of the index is consistent. It models the current skew by attributing to each option of a definite strike its own future Black-Scholes-style tree of constant instantaneous volatility. Then, as the index moves, each option keeps the exactly the same constant future instantaneous volatility in its future Black-Scholes valuation tree by moving the previously current tree so that its root now sits at the current index level.

Equation 9.2 shows that the implied volatility for an option of strike K is independent of index level S , and therefore the delta of the option is the same as the Black-Scholes delta.

Table 1 summarizes the behavior of volatilities under the sticky-strike rule.

TABLE 1. Volatility behavior using the sticky-strike rule.

Quantity	Behavior
Fixed-strike volatility:	is independent of index level
At-the-money volatility $\Sigma_{atm}(S)$:	$\Sigma_{atm}(S, t) = \Sigma_0 - b(S - S_0)$ which decreases as index level increases
Exposure Δ :	$= \Delta_{BS}$

You can think of this model as representing Irrational Exuberance. When the index market rises implied volatility falls, so that you trade at-the-money options at progressively lower volatilities as the index rises. If you are a market maker, and you lower the implied volatility as the market rises, this implies that you think that the higher the market gets, the less likely a future catastrophe will occur. That's irrational and it can't go on for too long or you will end up at zero volatility.

9.1.3 The Sticky Delta/Moneyness Rule

The sticky-delta rule is a more subtle view of what quantity remains invariant as the index moves. It's easier to start by explaining the related concept of sticky moneyness.

Sticky moneyness means that an option's volatility depends only on its moneyness K/S ; the option's dependence on index level and strike derives entirely from its dependence on the moneyness. In mathematical terms, this means we generalize Equation 9.1 to

$$\Sigma(S, K, t) = \Sigma_0 - b(K/S - 1)S_0 \quad \text{Sticky Moneyness Rule} \quad \text{Eq.9.3}$$

where S_0 is the initial index level at which the skew is first observed. Expanding this for S close to S_0 close to K , we obtain to leading order in $(K - S)$ for S and K both close to S_0 , the approximation

$$\Sigma(S, K, t) = \Sigma_0 - b(K - S) \quad \text{Eq.9.4}$$


The sticky moneyness rule is an attempt to shift the skew as the stock price moves by adjusting for moneyness. It quantifies the intuition that the current level of at-the-money volatility, the volatility of the most liquid option, should stay constant as the index moves. Similarly, for example, the option that is 10% out of the money after the index move should have the same implied volatility as the correspondingly out-of-the-money option before the index move.

This model assumes that the market mean reverts to a definite at-the-money volatility Σ_0 independent of market level. It's a model of common sense and moderation.

For a roughly linear skew, Equation 9.4 indicates that

$$\Sigma \approx \Sigma(S - K)$$

so that implied volatility is the same for all values of $S - K$, and in particular of $S - K = 0$, i.e. at the money. You can see from this, in the graph below, that implied volatility must rise when S rises.



In the Black-Scholes model, the exposure Δ_{BS} depends on K and S through the moneyness K/S , so that “sticky moneyness” is equivalent to “sticky delta,” with an at-the-money option corresponding approximately to $|\Delta_{BS}| = 0.5$. Options market participants often find it very convenient to think of the value of the $(|\Delta_{BS}| - 0.5)$ as a measure of an option’s out-of-the-moneyness expressed in units of its volatility until expiration. But, strictly speaking, sticky delta means that the implied volatility must be purely a function of Δ_{BS} , not just

K/S . Δ_{BS} depends on $\frac{\ln S/K}{\Sigma \sqrt{\tau}}$.

Equation 8.12 as well as the illustration above shows that in a negatively skewed market the implied volatility for an option of strike K increases with index level S , and therefore the delta of the option is greater than the Black-Scholes delta for an option with the same Black-Scholes volatility. Table 2 summarizes the behavior of volatilities under the sticky-delta rule.

TABLE 2. Index Volatility behavior using the sticky-delta/moneyness rule.


Quantity	Behavior
Fixed-strike volatility:	increases as index level increases
At-the-money volatility:	is independent of index level
Exposure Δ :	$> \Delta_{BS}$

9.1.4 The Sticky Implied Tree Model: One Index, One Local Volatility Function, One Tree.

You can interpret all current index options prices as determining a single consistent unique tree – the implied tree – of future instantaneous index volatilities consistent with the current market and its expectations of future volatilities. This consistency contrasts with the two previous stickiness rules, where each option is described by a different Black-Scholes tree for the index, because the B-S volatility changes with strike and market level, even though all options have the same index underlying.

Figure 9.1 shows a schematic view of the implied tree and its local volatilities consistent with a particular implied volatility surface. These *local volatilities*, which vary with future index level and future time, bear the same relation to current implied volatilities as forward rates bear to current bond yields.

FIGURE 9.1. The implied tree corresponding to a given implied volatility surface.



In the tree below, local volatility increases (twice as fast as the skew, we've seen) as the index level decreases, because the implied tree model attributes the implied volatility skew to the market's expectation of higher realized (local) volatilities and higher implied volatilities in the event that the index moves down. You can also approximately think of this aversion to increased volatilities on a downward index move as an aversion to downward index jumps.

Once you have determined the future index tree implied by the current skew and current index level, you can isolate future subtrees within it at different future index levels to compute the options market's expectation of future skews. This is similar to rolling along the curve of forward rates to compute the bond market's expectation of future yields.

Extracting Local Volatilities From Implied Volatilities.

The implied tree model allows the detailed numerical extraction of future local and implied volatilities from current implied volatilities. However, just as the yield to maturity of a bond is the linear average of forward rates until the bond's maturity, so you can think of an option's Black-Scholes implied volatility as the approximately linear average of local volatilities between the current index level and the option's strike. Table 3 displays a simple example in which we use this approximation to extract future local and implied volatilities from the current skew. For strikes not far from the money, this method works surprisingly well when compared with more exact numerical methods.

Table 3: Extracting local volatilities from a sample of hypothetical implied volatilities

Index = 100			
Strike	BS Vol. (%)	Index	Local Vol. (%)
100	20%	100	20%
99	21%	99	22%
98	22%	98	24%
97	23%	97	26%

The first two columns in Table 3 show the current implied volatility skew when the index is at a level of 100. The skew is taken to be linear and negative, increasing at one volatility point per strike point. When the index is at 100, the 100-strike at-the-money volatility in column 2 is 20% per year. This is also the value of the local volatility at index level of 100 in column 4, because local

volatility at some index level *is* effectively the short-term at-the-money implied volatility at that index level. The 99-strike volatility in column 2 is 21%. Therefore, the expected at-the-money (local) volatility at an index level of 99 must be 22%; it is the average of a local volatility of 20% at an index level of 100, and 22% at an index level of 99, that averages to 21% for the 99-strike option when the index is at 100.

Column 4 shows the local volatilities corresponding to this skew, computed using this averaging procedure. The averaging necessitates that local volatilities must increase twice as fast with index level as the implied volatilities increased with strike level. In the bond world, the analogous statement is that forward rates increase twice as fast with future time as bond yields increase with maturity.

Given the local volatilities in column 4, we can use them to reconstruct the implied volatilities at a different index level, say 99. The implied volatility is 22% at a strike level of 99, equal to the local volatility in column 4 at an index level of 99. The implied volatility is 23% for a strike of 98, the average of the local volatility of 22% at 99 and 24% at 98 in column 4.

Implied Volatility In The Sticky Implied Tree Model.

Table 3 illustrates the use of the one consistent implied tree. As the index level within the tree rises, you can see that the local volatilities decline, monotonically and (roughly) linearly, in order to match the linear strike dependence of the negative skew. In the sticky implied tree model, the implied volatility is an average of the local volatilities between S and K , so that, in the linear approximation to the skew,

$$\Sigma(S, K, t) = \Sigma_0 - b(K + S - 2S_0) \quad \text{Sticky Implied Tree Model}$$


At-the-money volatility is given by

$$\Sigma(S, S, t) = \Sigma_0 - 2b(S - S_0)$$

This equation shows that implied volatilities decrease as K or S increases. At-the-money implied volatility, for which K equals S , decreases twice as fast with index level as the skew slope with respect to strike. Because volatility decreases as you move to higher index levels in the tree, an option's exposure delta in the model is smaller than the Black-Scholes delta of an option with the same volatility.

In the linear approximation for local volatility models you can write $\Sigma \approx f(K + S)$ with Σ a function of $(K + S)$, which tells you how to relate the skew at different strikes and spot levels to each other via an invariance principle.

Let's draw this for a roughly linear negative skew in the following figure.



In order to satisfy $\Sigma(100, 80) = \Sigma(80, 100)$ with a negative skew, you can see that implied volatility must rise as the stock price decreases.

Table 4 summarizes the variation of implied volatility in the sticky implied tree model. Volatilities of options are anti-correlated with the index, rising as the index falls and falling as it rises. At-the-money volatilities respond to index moves at twice that rate.

TABLE 4. Equity index volatility behavior in the sticky implied tree model.

Quantity	Behavior
Fixed-strike volatility:	decreases as index level increases
At-the-money volatility:	decreases twice as rapidly as index level increases
Exposure Δ :	$< \Delta_{BS}$

In this regime the options market experiences fear. The sticky implied tree model assumes the skew arises from a fear of higher market volatility in the event of a fall, or perhaps jumps, and assumes that after the fall, market volatility will rise twice as fast.

9.1.5 Summary of the Rules

Assume the current stock price is S_0 , and that the current skew at this instant of time is linear, of the form $\Sigma(S_0, K) = \Sigma_0 - b(K - S_0)$. Here is a summary of the different “stickiness” and the models which correspond to them.

Sticky	General functional form for future implied volatility	Linear approximation: Future skew when stock price is S	Model with this property
Strike	$\Sigma(S, K) = f(K)$	$\Sigma(S, K) = \Sigma_0 - b(K - S_0)$	Black-Scholes ^a
Moneyness	$\Sigma(S, K) = f(K/S)$	$\Sigma(S, K) \approx \Sigma_0 - b(K - S)$	Stochastic volatility ^b , jump diffusion
Implied tree/ local volatility	$\Sigma(S, K) = f(K, S)$	$\Sigma(S, K) \approx \Sigma_0 - b(K + S - 2S_0)$ because Σ is approximately the average of the local volatilities between spot and strike.	Local volatility ^c
Delta	$\Sigma(S, K) = f(\Delta)$	$\Sigma(S, K) \approx \Sigma_0 - b[0.5 - \Delta_{\text{call}}(S, K, t, T)]$ or $\Sigma(S, K) \approx \Sigma_0 - b' \left[\frac{\ln K/S}{\Sigma \sqrt{\tau}} \right]$ Note that Δ is itself a function of Σ !	?

- The Black-Scholes model corresponds roughly to the sticky strike rule of thumb, but it cannot honestly accommodate a skew, because all implied volatilities are the same irrespective of strike in the Black-Scholes model. So, although people use it, it's not really consistent from a theoretical point of view.
- In stochastic volatility models, there is another stochastic variable, the volatility itself, and so $\Sigma(S, K) = f(K/S)$ only if the other stochastic variable doesn't change.
- Crepey, *Quantitative Finance 4* (Oct. 2004) 559-579, argues that the local volatility hedging is the best for equities markets, in that it gets things right when the market moves a lot and isn't very wrong otherwise. See next page.

9.1.6 Regimes of Volatility

None of these rules describe the behavior of implied volatility for long periods. They may however be useful in describing the behavior of implied volatilities over shorter periods. You can see more about these rules and the extent to which the options market empirically satisfied different rules at different times in the paper *Regimes of Volatility* in Risk, 1999.

What seemed to happen is that during calm upward-trending periods, the market satisfied the sticky strike rule, and during fearful periods it comes closer to satisfying the sticky implied tree rule. Over time, the combination of these behaviors leads to an at-the-money volatility that seems to revert to the same level. You can notice this in the earlier graph of S&P implied volatilities for

individual strikes at the start of the section “Patterns of Volatility Change” on page 2.

Crepey¹ has carried out a similar analysis of the FTSE 100. Note that though at-the-money volatility is strongly negatively correlated with the FTSE, fixed-strike volatility remained more or less constant as the FTSE rose, and then increased sharply as the FTSE dropped. It is difficult to generalize since volatility is fundamentally stochastic.




Figure 1. Inverted mirror picture on the FTSE 100, October 1, 1999 to March 1, 2000.

1. S. Crepey, Quantitative Finance 4 (Oct. 2004) 559-579:

9.2 Problems and Benefits of Local Volatility Models

9.2.1 Inadequacy of the Short-Term Skew

Setting aside questions of implementation, the main problem with local volatility models (for index options) is that, once calibrated to the current smile, (which flattens in moneyness at large expirations), future short-term local volatilities have less skew than current short-term implied volatilities. Therefore the short-term future skew in a local volatility model is too flat when compared with the index's perennially steep short-term skew. This suggests that local volatility models may not be the right model for the short-term index skew.

On the other hand, all financial models need recalibration; not one of them can model the future accurately given an initial state of the world; that's why, even in Black-Scholes, the implied volatility changes from day to day. Local volatility models are still popular when recalibrated regularly; they allow the valuation of exotic options consistent with the volatility surface for vanilla options, and are widely used as a means of valuing exotics.

9.2.2 Better Hedge Ratios During Volatility Regimes

Local volatility models may well produce better hedge ratios than Black Scholes. An argument of Crepey's, summarized below, suggests that when the models are regularly recalibrated, the hedge ratios produced by local volatility models may be better than those of Black-Scholes.

First, given that no model is perfect, the best hedge is the one that minimizes the variance of the P&L of the hedged portfolio. If the replication were exact, the variance would vanish.

Now consider a call on the index S , with value $C(S, t, K, T, \Sigma(S, t, K, T))$, and consider the hedged portfolio $\Pi = C - \Delta S$, long the call and short Δ shares of the index. We can now compare the effect of using local volatility hedge ratios vs. Black-Scholes implied volatility hedge ratios. The relevant hedged portfolios in the two cases are

$$\begin{aligned}\Pi_{BS} &= C - \Delta_{BS}S \\ \Pi_{loc} &= C - \Delta_{loc}S\end{aligned}\tag{Eq.9.5}$$

Therefore, the realized difference between the Black-Scholes-hedged P&L and the local-volatility-hedged P&L is entirely due to the difference in the hedge ratios used, so that for a stock move δS ,

$$\delta\Pi_{loc} - \delta\Pi_{BS} = (-\Delta_{loc} + \Delta_{BS})\delta S \equiv \varepsilon \times \delta S \quad \text{Eq.9.6}$$

since the change in the market value of the option is the same in either case.

In a local volatility model, in a negatively skewed market,

$$\Delta_{loc} = \Delta_{BS} + \frac{\partial C \partial \Sigma}{\partial \Sigma \partial S} \approx \Delta_{BS} + \frac{\partial C \partial \Sigma}{\partial \Sigma \partial K} \leq \Delta_{BS}$$

and so $\varepsilon = (\Delta_{BS} - \Delta_{loc}) \geq 0$.

From the theory of hedging, the formulae for the respective changes in the P&Ls during time δt in each model, as we worked out in Lecture 2 using a Taylor expansion, is

$$\begin{aligned} \delta\Pi_{BS} &= \frac{1}{2}\Gamma_{BS}S^2[\sigma_R^2 - \Sigma^2]\delta t \\ \delta\Pi_{loc} &= \frac{1}{2}\Gamma_{loc}S^2[\sigma_R^2 - \sigma(S, \delta t)^2]\delta t \end{aligned} \quad \text{Eq.9.7}$$

since the implied volatility of the index at level S over the next instant is exactly the local volatility $\sigma(S, \delta t)$ in the tree at index level S for very short expirations. (The p.d.e. in local volatility models is the Black-Scholes equation with Σ in the equation replaced by $\sigma(S, \delta t)$.)

The B-S P&L is positive or negative depending on whether realized volatility is greater or less than implied volatility. The local volatility P&L is positive or negative depending on whether realized volatility is greater or less than short-term local volatility.

Which of these hedged P&Ls is closer to zero on average, i.e. which one produces a better hedging strategy?

Combining Equation 9.6 and Equation 9.7, we see that in a local volatility model,

$$\delta\Pi_{BS} = \delta\Pi_{loc} - \varepsilon \times \delta S = \frac{1}{2}\Gamma_{loc}S^2[\sigma_R^2 - \sigma(S, \delta t)^2]\delta t - \varepsilon \times \delta S \quad \text{Eq.9.8}$$

This is our formula for the P&L of a Black-Scholes hedging strategy in a world where the skew is negative.

Let's analyze this as Crepey does. There are two sources of Black-Scholes contributions to P&L in Equation 9.8, one from the gamma term which is quadratic and non-directional, and depends on the volatility mismatch, and one

from the hedging mismatch which is directional and depends on the sign of δS .

Crepey discusses four different market regimes, grouped along two axes, volatility and direction, as displayed in Table 5. Indexes can move up or down, with high or low realized volatility compared to instantaneous local volatility.

For volatile down markets (a fast sell-off, as Crepey calls it), both terms in Equation 9.8 are positive, and the errors to $\delta \Pi_{BS}$ are additive. The Black-Scholes P&L differs from zero (the perfect hedge value) due to two additive contributions.

For non-volatile up markets (slow rise), both terms are negative and the same is true.

For slow sell-offs or fast rises, the two terms in Equation 9.8 have opposite signs, and the hedging errors tend to cancel.

TABLE 5. Types of Markets: Equity index markets have the characteristics of the yellow cells.

Volatility Direction	volatile	non-volatile
up	$\sigma_R > \sigma(S, \delta t), \delta S > 0$	$\sigma_R < \sigma(S, \delta t), \delta S > 0$
down	$\sigma_R > \sigma(S, \delta t), \delta S < 0$	$\sigma_R < \sigma(S, \delta t), \delta S < 0$

Therefore, the Black-Scholes hedging strategy will perform worst in fast sell-offs or slow rises. The local-volatility hedging strategy performs worst in slow sell-offs and fast rises.

Negatively skewed equity index markets are precisely characterized by slow rises and fast sell-offs. Therefore the Black-Scholes hedging strategy is worse than the local volatility strategy in these characteristic regimes. Crepey has also backtested the hedging of actual options to show that the P&L of a hedged portfolio has less variance under the local volatility hedging strategy.

Therefore, Crepey argues, for negatively skewed index markets, local volatility models have hedge ratios that, on average, tend to work better than pure Black-Scholes hedge ratios.