Nail in the Coffin

Dial 33 for your local cleaner

Elie Ayache[1] trades in his philosopher's stone to take on local volatility yet again . . .

Do you remember the character Victor the Cleaner from the French movie *La Femme Nikita* (1989), by director Luc Besson? The Cleaner is the professional hit man in charge of "cleaning up after" other hit men, when they mess up their mission. Often, cleaning up the mess means not only finishing off the target that the first hit men may have failed to terminate properly, but retiring the hit men themselves, as they can no longer be left on the scene or made to leave the scene. The Cleaner just "takes care" of everybody then literally cleans up the place, with the help of the appropriate solvents and detergents. (I will spare you the gory details.)

The Cleaner has gained much popularity since his first incarnation in the shape of French actor Jean Reno in ... *Nikita*. In the Hollywood remake, *The Assassin* (1993), Harvey Keitel takes over the cleaning business. Probably the reason why Quentin Tarantino later hires him, in *Pulp Fiction* (1994), to play Winston 'The Wolf' Wolfe: a master cleaner with a touch of class who does wonders after the two sympathetic leading (hit) men, Vincent and Jules, accidentally spray blood and brain inside their escape car. For all this, the character of the Cleaner seems to have stuck to Jean Reno, as Luc Besson soon casts him as the leading role in the movie *Leon* (1994), also known as *The Cleaner*.

In the later *Godzilla* movie (1998), Jean Reno returns as yet another cleaner, only this time with a larger-than-life assignment. He is undercover agent Philippe Roche, of the French Secret Service, and his job is to gun down a giant lizard that is wreaking havoc in Manhattan. The mutant monster was created by radiation following the French atomic bomb tests in the South Pacific, subsequently swimming its way to the cozier atmosphere of the Big Apple to breed. When asked why he was assigned this mission, Roche famously replies: "I am a patriot. I love my country. These tests, done by my country, left a terrible mess. I am here to clean it up."

Sir! Step away!

The lizard of Wall Street

Godzilla is not the only mutant creature that the French have produced after messing around with the laws of nature. Local volatility is another and it has also found its way to Wall Street. It was crystallized by a formula derived by French derivative pricing expert Bruno Dupire in 1994, and subsequently spawned across all volatility arbitrage desks. Being a French patriot myself, I feel it is my duty today to come and clean up the mess that the local volatility model is leaving behind.

Not even *Wilmott* magazine, or the forum reserved to its columnists, are safe from the mess. Indeed, no sooner did I identifiy the variance swap as the new volatility instrument which would definitely reject the local volatility model than a column, appearing next to mine[2], recalled the local volatility model to our attention and reinvested it with the power to impede and distract us again.

As if the problem of the smile dynamics, which is bound up with the problem of exotic option pricing and the general problem of dynamic hedging, did not make it clear enough that the local volatility model is a dead end! As if the articles that Philippe and I have been publishing in this magazine did not expose the fallacy of the inhomogeneous models loudly enough or demonstrate the superiority of the homogenous models clearly enough![3] Not mentioning that inferring the local volatility surface from a scatter of vanilla options quotes is an ill-posed inverse problem, and that spline interpolation/extrapolation of the implied volatility surface, no matter how cubic or how smoothed or how clever, is almost guaranteed to generate arbitrageable option prices!

Brecher recognizes "this is a real problem" and that "the results depend very sensitively on the interpolation method." But he argues that "such effects are localized and do not lead to large differences in option prices."

Prices of vanilla options, that is. For, surely enough, instability in the volatility surface, even if it is localized, *can* lead to noticeable pricing differences for options with very rapidly changing gammas, like the barriers. This, Brecher recognizes too when he writes: "Although not necessarily true for more exotic structures, the precise form of the local volatility function in many cases seems to have little impact on plain vanilla European and American option prices." He therefore condemns his own technique to be no more than a fancy exercise in cubic spline interpolation and surface smoothing, with the pricing of vanillas as its only purpose. What good is a smile model if it is unable to price the exotics or *hedge* the vanillas?

Reprise

To repeat, the observed prices of vanilla options do not contain *any* information about the smile dynamics as they are just the snapshot of the present smile. In other words, from the prices of vanilla options (even a continuum thereof, in strike and maturity) we can only infer the probability distribution of the underlying price at the maturity dates of the options, as seen from today and from the spot price. *Only in a non-parametric local volatility model* does this impose the *conditional*, or forward, probability distributions, i.e. the probability of ending up at a certain price at a certain future date, conditionally on starting at a certain different price at a certain *future* – although nearer – date. Outside the local volatility model, for instance in non-parametric jump-diffusion and stochastic volatility models, the forward probability distributions are underdetermined and the prices of options that are sensitive to those forward probabilities, typically path dependent structures such as the barrier option or the cliquet option or the variance swap, will be out of control even when the model is calibrated to the continuum of vanillas. Likewise, the dynamic hedging strategies, even of the vanillas, are undetermined.[4]

So when Brecher writes in his introduction (just before reviewing the variety of smile models that have been proposed in the literature) that "the real motivation for considering the effects of the volatility smile in option pricing is precisely this: to calibrate one's pricing of exotic, or even just American, options to the market-observed prices of European plain vanillas," he is in fact *only* thinking of the local volatility model. Yes, the real motivation of a model for the volatility smile is to price the exotics (and work out reliable hedges for all derivative instruments, including the vanillas). But no, it certainly is not to calibrate the pricing of exotics to the prices of the vanillas. There is nothing in common between a vanilla option, a barrier option, and a variance swap. The price of each one of these structures contains probability information that is irreducible to the other. And you have no choice but to include the exotics in the calibration of your model. Because the local volatility model doesn't give you the choice "to have no choice but to include the exotics in the calibration," it locks itself up in an insane logic where the vanillas are the beginning and the end of everything. (They are the unique reference against which you have to calibrate everything.)

No matter how many nails I have driven into the coffin of the local volatility model in previous articles, the beast just wouldn't go away. Perhaps my arguments were too subtle and too oblique to finish it off. I used to be called the Philosopher, on this forum. From now on, you can call me the Cleaner.

Figure 1: Godzilla too hard to swallow? How about with a little pre-smoothing?

Local volatility has no meaning

What other arguments might you have left at this point to further resist the obligation to include the exotics in the calibration of the smile model? Maybe you prefer to use a robust and parsimonious model and to personally take a view? Maybe you are happy with the model-dependent relation that your parsimonious model imposes between the vanillas and the exotics – so long as you know the problem and the ins and outs of your model, you will know your way around the problem – and you are more than happy to calibrate your model to a bunch of vanillas, just in order to get it a little closer to the market? Maybe so, and I am certainly no advocate of calibration bulimia. But then go ahead and take a look at the shape of the local volatility surfaces produced by non-parametric calibration (see Figure 1). Do you call them parsimonious and natural?

In the "Nail in the Coffin" column that followed mine[5], Philippe Henrotte delivered the *coup de grâce* to the local volatility model, in terms of a very simple question. He asked: "What if vanilla options were just an accident of history, and volatility trading had started instead with non linear payoffs, such as the log contract or the power contract or the variance swap, which differed from the vanillas in that they exhibited *no singularity* in space at the strike price?" (This "What if?" is typically a philosophical question whose only purpose is to make us think *today*. As a matter of fact, history couldn't have unfolded differently because at the time when Black and Scholes priced the vanilla options, nobody had any idea that volatility trading would ensue.) It isn't surprising, Philippe argued, that the vanilla option, which is basically an inhomogeneous instrument with a localized strike price, should suggest the idea of a local volatility model. Now imagine that all we had instead, as "implied volatility" quotes, were surfaces of volatility strikes of forward starting variance swaps and options on variance swaps, with different starting dates and maturity dates. Who would have ever thought of the local volatility model as the way to capture the pricing of these perfectly homogeneous derivative instruments?

Think of it this way: volatility is a statistical measure pertaining to the underlying share price, not to the option. It is a *single* number whose past realization you can infer from the time series of underlying prices, with controllable error bounds that depend crucially on the data generating process you assume. In the simulation of Brownian motion, there isn't just one, but *many* underlying paths that correspond to that volatility number. The great achievement of the Black-Scholes option pricing theory is to have established a definite, one-to-one relation between the *single* volatility number of the underlying share price and the price of a *single* option. Whereas the relation between the share price and its volatility would require a whole path to become manifest (and for that matter, only statistically so), we are now able, thanks to Black-Scholes, to accurately and definitely link volatility to an unquestionable option price. This is the magic of Itô's lemma and of convexity. However, this should be no reason to turn the determination the other way round and pervert the logic to the fantastic limit where we think we can deterministically infer, from a *surface* of options prices, the one-to-one *mapping* of something called the "local volatility surface"! How would we infer such a local volatility surface if the only observables we had were the paths of the underlying? What I am saying is that volatility loses its meaning of a statistical variable in the local volatility model, that is to say, it loses *all meaning*. To put it differently, we should only be allowed to infer a *single* volatility number from option prices, in the Black-Scholes setting. And if we are looking for a better fit of those prices, we must add *other* factors to our model, such as jumps, or stochastic volatility.

Local volatility shouldn't exist

Ever since Philippe published his article, thus deconstructing the vanilla options, I couldn't help thinking that the local volatility model never actually existed or was never supposed to exist. Philippe described the variance swap as a creature whose birth and subsequent spawning was mainly due to the local volatility model, as it critically depended on the diffusion assumption. For a large class of diffusion models, theory shows indeed that the variance swap is statically replicable by the vanilla options and – what's even more remarkable – that the replication strategy does not depend on the particular model falling in that class. Only jumps could invalidate the replication result so the local volatility fanatic was now happy, not only to be able to price a new instrument in his framework – as a matter of fact, he needs no model to price the variance swap but only the availability of the vanilla option prices – but to be able, *precisely*, to reproduce those available vanilla option prices *without* invalidating the replication result. It is as if the pricing of the variance swap had re-established the reality of the vanilla smile, and as if this "new" reality was called to mind by the local volatility fanatic as a "further" vindication of his favorite model.

Comes the day when the variance swap, now alive and kicking in a very tight and very liquid market of its own, no longer agrees with the value of the static replication strategy. No matter how well, or how completely, the diffusion models reproduce the vanilla smile surface, the unfailing discrepancy in the price of the variance swap can only mean that they are missing something else. And what they are missing are the jumps. This sentences the end of all pure diffusion models. But while a homogenous stochastic volatility model, like the Heston model, is rejected because it fails to explain the price of the variance swap and the vanilla options simultaneously, the local volatility model, although initially rejected for the same reason, cannot even remain in the garbage can, as we now wonder how the thought ever came to us to conceive of such monstrosity in the first place.

Sci-fi horror

Philippe Henrotte was not aware of Brecher's article at the time he wrote his. As final as the blow he delivered may have been, its timing could not stop the lizard from making just one last scary appearance, as is customary in the horror movies that we enjoy most. Even I, the Cleaner, have become accustomed to a world clean of local volatility, and find it hard to return to a place where spline extrapolation of implied volatility can result in negative numbers at the boundary of the grid or a place where local variance itself can come out negative from Dupire's formula.

"The local volatility model helped give birth to a liquid equity variance swap market," Philippe writes. "But as its market matured, it freed itself from the matrix of the theory ... A frontal collision is looming between the variance swap and the local volatility model ... As in a bad horror movie, the beast is about to devour its creator." Philippe calls this the "Frankenstein syndrome."

In my scenario, the roles are reversed. Local volatility is the monster beast (not the variance swap) and as Philippe and I are set on our voyage to the new world of variance swaps and homogenous jump-diffusion/stochastic volatility models, the sudden appearance of Brecher's article strikes us like the *Alien* that still managed to creep into the ship. Or maybe it is a giant lizard, as the movie *Godzilla* suggests, and it will destroy the whole city. The *Alien* and *Godzilla* have the eggs in common, and the scene of the hatching is a delectable supplement. The *Predator* is a sort of solitary Cleaner, and in a recent Hollywood sequel he confronts the *Alien*. Yet, the *Predator* has also something reptilian about him (way he climbs the trees, the scales). I have seen no eggs, though. Dinosaurs lay eggs, and at the end of Philippe's piece, the dominance of the inhomogeneous species, which includes the vanilla options and the local volatility model they have inspired, is compared to the reign of the dinosaurs. Or maybe the species will not go extinct and a genetically reengineered local volatility model will be able to swallow the variance swap. A new creature will then emerge from the volatility swamp. Much darker and scarier, no doubt. And then a special cleaning squad will...

Enough. *Arrêtez le massacre!*

"A lizard-free tomorrow"[6]

Rather than trying to figure the bloodiest scenario and the ugliest beast, better to clear our minds and engage in positive thinking. The best defense against local volatility is, like I said, to really believe it never existed. (Or will that be "the greatest trick the devil ever pulled" so the beast will always haunt us?) This is why, for the remainder of the article, I will guide you through a world free of local volatility, where calibration to all kinds of payoff structures (and not just the vanillas) is possible and is meaningful. That is to say, it is truly *informative*; every new structure I will calibrate against will help define and determine the model in ways that are not possible without it. I will no longer mention the local volatility model yet I ask you to recall it to mind (yes, like the devil's temptation!) every time I perform a new calibration, and to see for yourself how absolutely unthinkable that it should reach to this level.

Equity-to-credit

First, consider the equity volatility smile shown in Table 1. The day is February 3rd 2003. The options expire on January 16th 2004. The equity is Tyco and the underlying stock price is $16. We assume no dividends and the interest rate is 2.75 per cent. The data is scarce; however the massive volatility skew on the 7.5 put suggests default risk. We consequently calibrate a very simple equity-to-credit model to the data. This is a jump-diffusion model where the underlying diffuses with constant Brownian volatility and can go to default, i.e. jump to zero, with constant Poisson intensity (a. k. a. hazard rate). Calibration is achieved by least squares minimization of the difference between the market option prices and the model prices. (It is understood that the Poisson intensities of all the jump processes that we will henceforth imply from the market are expressed in the risk-neutral measure.)

The results are shown in Table 2.

Now let us value the rest of the options of same maturity under this jump-diffusion model and compute the equivalent Black-Scholes implied volatility. The resulting theoretical volatility smile is exhibited in Table 3.

Notice the implied volatility on the 5 put. I doubt any spline extrapolation could have predicted this value. The suggestion to extrapolate to constant implied volatility (in this case) outside the grid of market data is, of course, even worse!

Table 1: Implied volatility smile of Tyco. The options maturity is 16 Jan 2004. The day is 03 Feb 2003. The underlying spot price is $16. The options are assumed to be European

	Strike											
	5.00	7.50	10.00	12.50	15.00	17.50	20.00	22.50	25.00	30.00	35.00	45.00
Implied volatility		75.40%			51.40%	47.10%	45.00%					

Actually, more market data *was* available, and if we now compare the model prices with the larger set of option market prices we get the picture in Table 4 and Figure 2.

Impressive, don't you think? As far as the out-of-the-money puts are concerned, the market implied volatility skew seems to be perfectly explainable by default risk. (Go ahead and try the local volatility model, here.)

Introducing the CDS

Not only is default risk the most probable explanation, but it has become a manifest object nowadays, thanks to the presence of credit default swap (CDS) quotes. So we may actually complete our data set with the five-year CDS spread that was quoted that day on Tyco: 435 bps.

And the next thing that pops to mind is the desire to calibrate the model both to the option volatility smile *and* the CDS spread. The hazard rate we found previously would imply a 620 bps CDS spread (assuming 30 per cent recovery). So we should expect the new parameters to be different. Indeed, they are (see Table 5.)

Table 2: Calibrated parameters of the simple equity-to-credit process

	Brownian volatility	Hazard rate
E2C process	38.57%	8.74%

And if we re-compute the option model prices and CDS spread and compare to the market prices we get the new picture in Tables 6 and 7 and Figure 3.

The model is still doing fine for the at-the-money options and the far-out-of-the-money put. Indeed, the first are most sensitive to the Brownian volatility parameter and the second to the hazard rate parameter. Therefore, it is not surprising that they weigh most in the calibration procedure, relatively to the conflicting CDS quote. However, the nearer out-of-the-money puts and the far-out-of-the-money call are undervalued by the model, and the nearer out-of-the-money calls are overvalued. (Of course, we could have weighted the options and the CDS differently, relatively to each other, and the calibration procedure would have produced slightly different results.)

Figure 2: Comparison between the market volatility smile and the theoretical smile of options maturing on 16 Jan 2004. Today is 03 Feb 2003

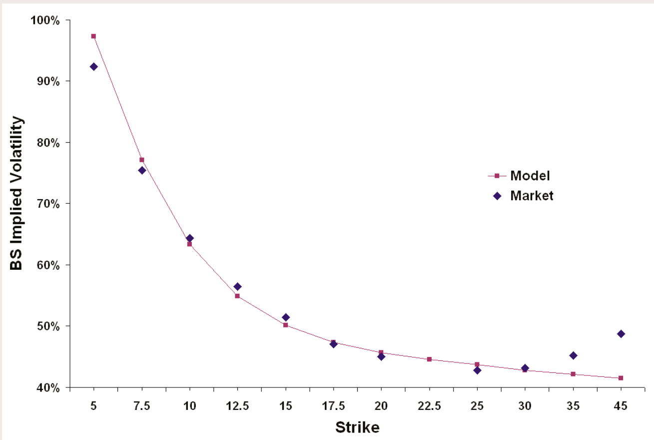

Table 3: Theoretical volatility smile of options maturing on 16 Jan 2004. Today is 03 Feb 2003

	Strike											
	5.00	7.50	10.00	12.50	15.00	17.50	20.00	22.50	25.00	30.00	35.00	45.00
Model	97.31%	77.10%	63.33%	54.91%	50.13%	47.37%	45.67%	44.55%	43.76%	42.77%	42.15%	41.48%

Introducing stochastic volatility

Fat tails on the nearer out-of-the-money puts suggest that default may not be the only

Table 4: Comparison between the market volatility smile and the theoretical smile of options maturing on 16 Jan 2004. Today is 03 Feb 2003

	Strike											
	5.00	7.50	10.00	12.50	15.00	17.50	20.00	22.50	25.00	30.00	35.00	45.00
Market	92.40%	75.40%	64.40%	56.50%	51.40%	47.10%	45.00%		42.80%	43.20%	45.20%	48.70%
Model	97.31%	77.10%	63.33%	54.91%	50.13%	47.37%	45.67%	44.55%	43.76%	42.77%	42.15%	41.48%

Table 5: Calibrated parameters of the equity-to-credit process using options and CDS

	Brownian volatility	Hazard rate
E2C process	41.07%	6.74%

source of volatility smile. Stochastic volatility may be another factor, especially when it is correlated with the underlying.

You know that I know that you know what a local volatility fanatic would want to do at this juncture: Make the Brownian volatility a deterministic function of stock price, and match the remaining options prices.

Figure 3: Comparison between the market volatility smile and the theoretical smile of options maturing on 16 Jan 2004. Today is 03 Feb 2003

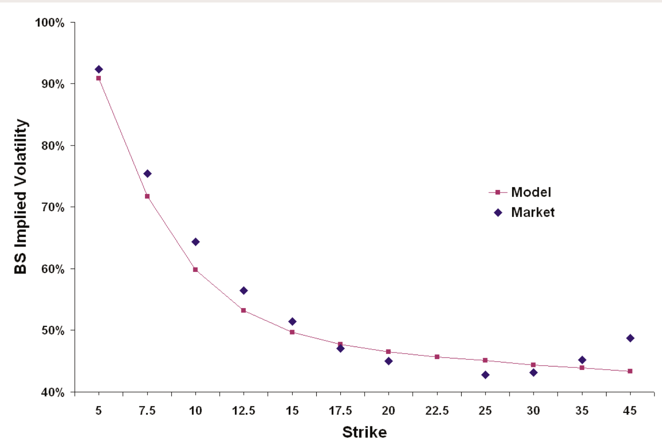

Table 6: Comparison between the market volatility smile and the theoretical smile of options maturing on 16 Jan 2004. Today is 03 Feb 2003

	Strike											
	5.00	7.50	10.00	12.50	15.00	17.50	20.00	22.50	25.00	30.00	35.00	45.00
Market	92.40%	75.40%	64.40%	56.50%	51.40%	47.10%	45.00%		42.80%	43.20%	45.20%	48.70%
Model	90.86%	71.74%	59.83%	53.23%	49.71%	47.71%	46.49%	45.67%	45.09%	44.36%	43.89%	43.37%

Table 7: Comparison between the market five-year CDS spread and the theoretical spread (assuming 30 per cent recovery)

	Market	Model
CDS spread	435 bps	475 bps

But you know that I know that we both know better than this (not mentioning that we both believe that local volatility doesn't exist). So here is my suggestion: Consider a very simple two-regime switching model instead, where both Brownian volatility and the hazard rate admit of two discrete states. This is the simplest, most parsimonious, most robust way of making the previous equity-to-credit model stochastic, while retaining space and time homogeneity.

The Brownian volatility and the hazard rate have values σ_1 and λ_1 in regime 1, σ_2 and λ_2 in regime 2. Migration from regime 1 to regime 2 is triggered by a Poisson process of intensity $\lambda_{1,2}$. When this occurs, the underlying stock price undergoes a proportional jump of size $y_{1,2}$. The intensity of migrating from regime 2 back to regime 1 is $\lambda_{2,1}$, and the corresponding return jump is $y_{2,1}$. By convention, regime 1 is the present regime. (Like I said, we wish to keep it simple. In a more general regime-switching model, we can assume any number of regimes, and any number of jump processes, of fixed size and fixed intensity, triggering the transition between a given pair of regimes.)

Calibration of this homogeneous eight-parameter, stochastic volatility/stochastic hazard model to the one-year option volatility smile and to the five-year CDS spread yields the result in Table 8, which I will call solution #1. From Tables 9 and 10 and Figure 4, you can see that both the volatility smile and the CDS spread are perfectly matched.

Stochastic volatility and its correlation with the return jumps seem to explain the fat tails. Brownian volatility is noticeably higher in regime 2, yet strangely, credit risk seems to vanish in that regime. However, the Poisson intensity of reverting back to regime 1 is twice as high as the intensity of switching to regime 2. Also notice that the underlying can only jump downwards as we switch between regimes, not mentioning the jump to default. All in all, we have a combination of effects that can separately and jointly explain the volatility smile (the negative jumps, stochastic volatility negatively correlated with the underlying), and it is hard to tell whether the solution we have found is realistic or is even unique. Are we really certain volatility is higher in regime 2?

Only an option of longer maturity can give us information about the long-term behavior of volatility. Using the last parameters we found, we can compute the theoretical value of such an option, say the 15 put expiring the following year, on January 21st 2005. For that put, the model predicts a 56.33 per cent equivalent Black-Scholes implied volatility. And now we can lay another card on the table and reveal the real market implied volatility of that option: 49.40

Table 8: Solution #1. Calibrated parameters of the two-regime switching model

	Brownian volatility	Hazard rate
Regime 1	25.84%	8.36%
Regime 2	73.66%	0.00%
Switching	**Intensity**	**Jump size**
Regime 1 ⟶ Regime 2	0.72	-16.39%
Regime 2 ⟶ Regime 1	1.78	-20.99%

Figure 4: Solution #1. Comparison between the market volatility smile and the theoretical smile of options maturing on 16 Jan 2004. Today is 03 Feb 2003

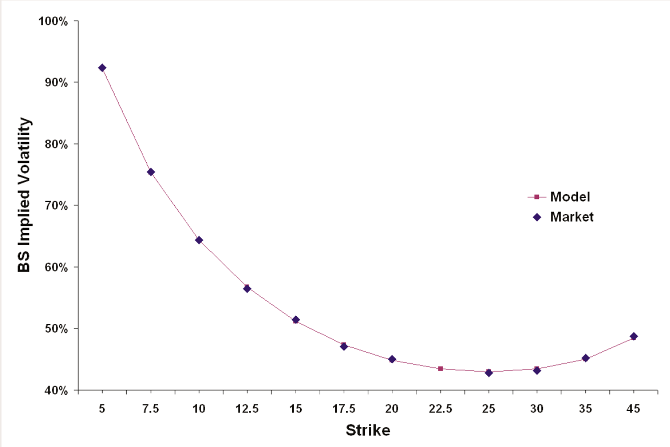

per cent. Our last calibration is obviously found wanting.

Introducing the longer-term option

We repeat the calibration of the two-regime switching model by adding the longer maturity option to our data set. We now find the parameters in Table 11, which I will call solution #2.

Again the market data is perfectly matched (see Table 12), and the results seem more realistic this time with a perfectly plausible explanation in terms of corporate events. Today, Tyco is in a regime of high equity price volatility (94.66 per cent) and high hazard rate (14.16 per cent), however it will, with very high probability (9.24 Poisson intensity), switch to a regime of much lower volatility (13.92 per cent) and lower hazard rate (5.44 per cent). This, we may take as the very definition of a *distressed debt*. A state where the company is probabilistically split between two spectacularly opposed scenarios: very high volatility and default risk on the one hand, imminent recovery and restructuring on the other.

Local minima

How could two solutions, as structurally different from each other as solutions #1 and #2, explain the same one year volatility smile and the five-year CDS spread? The answer is that calibration of the two-regime switching model, using options of a single maturity date and a single CDS is an ill-posed problem. Solutions #1 and #2 are two local minima of that tentative calibration, and it is only by chance that solution #1 was found first.

Only when the two-year option was added to the calibration set was solution #1 rejected and solution #2 selected as the answer to the extended calibration problem.

You start getting a sense of the real meaning of calibration, and of how it should be used to add *significant* information to the problem at hand. A few key market observables are added each time to the calibration set (the CDS spread, the two-year option price) and they make a big difference.

Realism and intuitive grasp of the parameters are not the only difference between solutions #1 and #2. Of course, the two solutions massively disagree on the pricing of the two-year option and this should help discriminate between them. But this poses at once a philosophical problem. Why should you be interested in the two-year option at all? Maybe your interest stops at the one year horizon and anything falling beyond it is irrelevant to you.

In thinking that way, you remain the hostage of a state of mind that you have inherited from the local volatility methodology and its forward induction procedure. It is the local volatility model, not you, who can make no use of the two-year option quote, *because it is unable to extrapolate the full two-year volatility smile from a single data point*.

By definition, space inhomogeneous models require information from the outermost and innermost recesses of space in order to complete their information set. When no such information is available, the suggestion is either to stop short of it or to create it artificially. Thus, inhomogeneous models lead us to the rather awkward situation where we have to invent what we want to explain, and then explain it.

By contrast, homogenous models are holistic and any additional information can alter their infinitesimal generator and affect the whole dynamics, no matter which part of space or which horizon it comes from. So chances are a difference will show, between solutions #1 and #2, *even in the short term*. As a matter of fact, the hedging will be different!

The hedging problem and the smile dynamics

In the regime-switching model, we are confronted with jumps and stochastic volatility, and dynamic hedging with the underlying can only make sense as optimal hedging. We compute optimal hedges in the sense that the self-financing dynamic hedging strategy guarantees that the portfolio com-

Table 9: Solution #1. Comparison between the market volatility smile and the theoretical smile of options maturing on 16 Jan 2004. Today is 03 Feb 2003

Strike	5.00	7.50	10.00	12.50	15.00	17.50	20.00	22.50	25.00	30.00	35.00	45.00
Market	92.40%	75.40%	64.40%	56.50%	51.40%	47.10%	45.00%		42.80%	43.20%	45.20%	48.70%
Model	92.30%	75.32%	64.34%	56.70%	51.15%	47.31%	44.85%	43.48%	42.93%	43.43%	45.03%	48.50%

Table 10: Solution #1. Comparison between the market five-year CDS spread and the theoretical spread (assuming 30% recovery)

	Market	**Model**
CDS spread	435 bps	435 bps

Table 11: Solution #2. Calibrated parameters of the two-regime switching model

	Brownian volatility	**Hazard rate**
Regime 1	94.66%	14.16%
Regime 2	13.92%	5.44%
Switching	**Intensity**	**Jump size**
Regime 1 ⟶ Regime 2	9.24	1.71%
Regime 2 ⟶ Regime 1	0.54	-26.31%

Table 12: Solution #2. Comparison between the market volatility smile and the theoretical smile. Today is 03 Feb 2003

	Strike											
	5.00	7.50	10.00	12.50	15.00	17.50	20.00	22.50	25.00	30.00	35.00	45.00
16 Jan 2004												
Market	92.40%	75.40%	64.40%	56.50%	51.40%	47.10%	45.00%		42.80%	43.20%	45.20%	48.70%
Model	92.40%	75.40%	64.33%	56.67%	51.24%	47.27%	44.73%	43.45%	43.06%	43.63%	44.69%	47.05%
21 Jan 2005												
Market					49.40%							
Model					49.40%							

Table 13: Comparison between the optimal hedging ratios of solutions #1 and #2 for options maturing on 16 Jan 2004. Today is 03 Feb 2003. The two solutions agree on the options prices because they fit the same smile. This, by Euler's theorem, entails they also agree on the options deltas. HERO is the minimized standard deviation of P&L of the hedged option

16 Jan 2004	**Call 16**				**Put 12.5**			
	Price	**Delta**	**Hedge**	**HERO**	**Price**	**Delta**	**Hedge**	**HERO**
Solution #1	$3.21	71.06%	49.84%	$1.43	$1.51	-12.75%	-32.34%	$1.35
Solution #2	$3.22	71.25%	65.70%	$1.33	$1.51	-12.89%	-18.36%	$1.28

posed of option and hedge breaks even on average and that the standard deviation of its P&L is minimized. Hedging has to be dealt with in the real (historical) probability measure, therefore we require the independent input of the Sharpe ratio of the underlying in order to revert from the risk-neutral probability to the real probability. Assuming a Sharpe ratio of 0.5, we compute the hedging ratios in solutions #1 and #2. The results are summarized in Table 13.

That the hedges turn out different makes perfect sense. The one-year volatility smile is explained by default risk, jumps and stochastic volatility. Although the two solutions agree on the CDS and the corresponding default component, they massively disagree on the dynamics of stochastic volatility, therefore on the smile dynamics, therefore on the hedging strategies of all vanilla options.

Arbitrage-free extrapolation

The next step is to produce the full two-year volatility smile and the term-structure of CDS spreads. Like I said, the homogeneous model is not confined to any portion of space or any time interval. In that respect too, it is the total opposite of the local volatility model. We can extrapolate option prices to any strike we want and any maturity we want; we can compute CDS curves of any length we want; and be guaranteed to generate prices which are arbitrage-free. The parameters of solution #2 predict the two-year volatility smile shown in Table 14 and Figure 5 and the CDS spread curve in Table 15.

In the market, the one-year CDS spread is quoted 500 bps. As for the two-year 7.5 put, it is not very liquid, and its bid-and-ask quotation is ($1.00-$1.12) which corresponds to (67 per cent-70 per cent) Black-Scholes implied volatility. Our model's prediction falls within the market range for the put; however, we can add the one-year CDS to our calibration set in order to pull our theoretical CDS curve towards the market. The new parameters are exhibited in Table 16.

Let us call this, solution #3. Interestingly, regime 2 is almost the same as before, however both the default intensity and the regime transition intensity have changed, in regime 1, in order to match the short term CDS spread. Indeed, the CDS curve is now almost perfectly matched and the volatility smile is almost unchanged (see Tables 17 and 18 and Figure 6).

Introducing the term-structure of volatility skew

Just to set our minds to rest, we perform one last calibration where we further extend our calibration data input *to include the market quote of the two-year 7.5 put*. So far, we have read the volatility term-structure and the CDS spread term-structure from the market, so why not also read the *term-structure of volatility skew* and calibrate against it? To really test the stability of our result, we elect to calibrate against the other extremity of the price range of the two-year 7.5 put, namely the 70 per cent implied volatility. The results are quite surprising. (See Tables 19, 20 and 21.)

Calling solution #4 this last set of parameters, we can see that it is not structurally different from solution #3. Regime 1 is still the present risky regime, and regime 2 is still the regime of expected restructuring and safety. After restructuring and transition to regime 2, we are still likely to fall back, with low probability, into the troubled regime 1, and this is also accompanied by a negative jump of the underlying stock price. However, the size of this negative jump is much larger in solution #4 (-86 per cent) and its intensity is much lower (4.02 per cent). It is akin to a jump to default (only it doesn't trigger the CDS payment), and this is what the model is actually using to explain the relatively higher volatility skew (70 per cent) of the longer-term out-of-the-

Table 14: Two-year theoretical volatility smile (options maturing on 21 Jan 2005). Today is 03 Feb 2003

	Strike											
	5.00	7.50	10.00	12.50	15.00	17.50	20.00	22.50	25.00	30.00	35.00	45.00
Model	79.20%	67.62%	59.07%	53.54%	49.40%	46.29%	43.82%	41.93%	40.52%	38.84%	38.17%	38.17%

Table 15: Theoretical CDS spread curve (assuming 30% recovery)

Maturity	**Spread (bps)**
1 Year	482
2 Year	453
3 Year	443
4 Year	438
5 Year	435

Figure 5: Two-year theoretical volatility smile (options maturing on 21 Jan 2005). Today is 03 Feb 2003

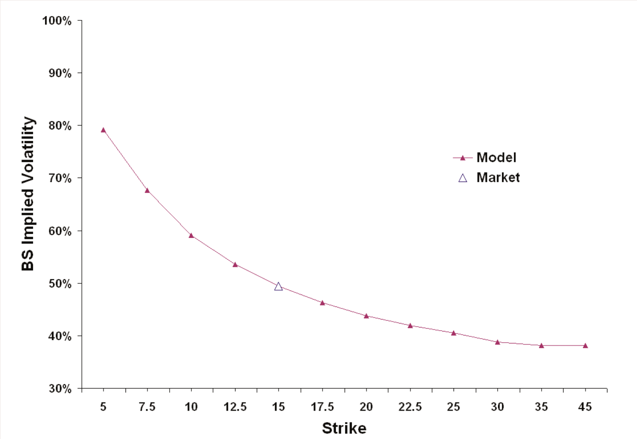

Figure 6: Solution #3. Comparison between the market volatility smile and the theoretical smile. Today is 03 Feb 2003

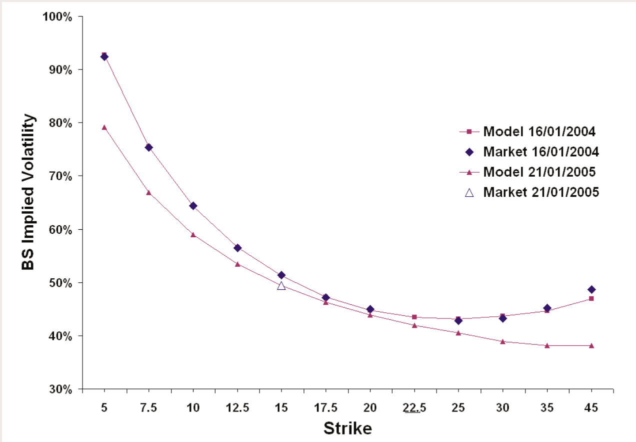

money put Of course, out-of-the money puts of the front maturity are also affected by this "equity default." As a result, Brownian volatility is reduced in regime 1; but most importantly, the hazard rate, or the probability of "true default," is reduced too, and the one-year CDS is consequently undervalued by solution #4.

You might argue that the phenomenon that we are uncovering here, the effect of the term-structure of volatility skew, is a second order adjustment and that solution #3 and #4 are not that different after all. True, the size of the jump out of regime 2 is different, but so is its probability. You might think they compensate each other on average. Don't the two solutions match the one-year volatility smile, the volatility term-structure and the five-year CDS perfectly? As for matching the rest, it must all depend on the relative weighting of the short term CDS and the long-term out-of-the-money put in the calibration procedure. "Calibration is an art anyway, not a science," you are tempted to think. Therefore, you feel it is up to you to decide which of these two instruments is more liquid, therefore more informative, therefore deserves more weight. Not mentioning that your horizon may, as before, not extend beyond the first year, so what do you care about matching exactly the longer-term volatility skew?

Enter the variance swap

One instrument whose value decisively depends on the size of the jumps is the variance swap. And this is why, in this last act of the kill, the variance swap has to make a sharp and clean entrance.

Variance swaps may not be as liquidly traded on equity single names as they are on equity indices. However, I contend that their space homogeneous nature makes them the ideal candidate for volatility betting and trading, and even more so in the single name situation than in the index.

Once again, think of the *distressed debt* situation. Equity-to-credit players specialize in the so-called "busted bonds" because of their sensitivity to both equity and credit. Equity-to-credit arbitrage, which involves trading the single name credit derivative (bond, CDS) against the out-of-

Table 16: Solution #3. Calibrated parameters of the two-regime switching model

	Brownian volatility	Hazard rate
Regime 1	78.29%	13.03%
Regime 2	13.82%	5.24%
Switching	**Intensity**	**Jump size**
Regime 1 ⟶ Regime 2	6.09	1.66%
Regime 2 ⟶ Regime 1	0.54	-27.12%

Table 18: Solution #3. Theoretical CDS spread curve (assuming 30% recovery)

Maturity	Spread (bps)
1 Year	495
2 Year	458
3 Year	445
4 Year	439
5 Year	435

Table 19: Solution #4. Calibrated parameters of the two-regime switching model

	Brownian volatility	Hazard rate
Regime 1	50.47%	7.07%
Regime 2	10.43%	5.87%
Switching	**Intensity**	**Jump size**
Regime 1 ⟶ Regime 2	1.07	10.44%
Regime 2 ⟶ Regime 1	0.04	-85.80%

Table 17: Solution #3. Comparison between the market volatility smile and the theoretical smile. Today is 03 Feb 2003

	Strike											
	5.00	7.50	10.00	12.50	15.00	17.50	20.00	22.50	25.00	30.00	35.00	45.00
16 Jan 2004												
Market	2.40%	75.40%	64.40%	56.50%	51.40%	47.10%	45.00%		42.80%	43.20%	45.20%	48.70%
Model	92.76%	75.50%	64.29%	56.64%	51.22%	47.28%	44.76%	43.48%	43.08%	43.62%	44.66%	46.97%
21 Jan 2005												
Market					49.40%							
Model	79.19%	66.93%	59.03%	53.49%	49.42%	46.30%	43.85%	41.97%	40.55%	38.85%	38.16%	38.13%

>

Table 20: Solution #4. Comparison between the market volatility smile and the theoretical smile. Today is 03 Feb 2003

	Strike											
	5.00	7.50	10.00	12.50	15.00	17.50	20.00	22.50	25.00	30.00	35.00	45.00
16 Jan 2004												
Market	92.40%	75.40%	64.40%	56.50%	51.40%	47.10%	45.00%		42.80%	43.20%	45.20%	48.70 %
Model	93.10%	75.73%	64.39%	56.71%	51.21%	47.24%	44.76%	43.60%	43.25%	43.50%	43.99%	45.16%
21 Jan 2005												
Market		70.00%			49.40%							
Model	81.64%	69.83%	61.24%	54.61%	49.40%	45.33%	42.30%	40.23%	38.97%	38.01%	38.00%	38.70%

Table 21: Solution #4. Theoretical CDS spread curve (assuming 30% recovery)

Maturity	Spread (bps)
1 Year	469
2 Year	454
3 Year	445
4 Year	440
5 Year	436

Table 22: Comparison of the values, given by solutions #3 and #4, of the one-year variance swap, starting on 03 Feb 2003 and ending on 16 Jan 2004

	Solution #3	Solution #4
Volatility strike	47.58%	50.20%

the-money puts, is very popular among hedge funds nowadays. Yet those players complain that the standard equity options become useless as volatility instruments in the penny stock situation, because their strikes, no matter whether of calls or puts, end up so much higher than the stock price! By contrast, the variance swap remains at-the-money, and has, in such situations, literally exploded following the downgrade, because of the exceptional volatility of the penny stock.

Potentially, all stocks are penny stocks (at least the ones that carry substantial CDS spreads and attract, for this reason, the equity-to-credit play). So the pricing and the trading of the single name variance swap shouldn't remain a luxury, but should become common practice in equity volatility arbitrage.

The reason it didn't is that the pricing methodology everyone is using – static replication of the variance swap with the vanillas – breaks down completely in case of jumps of the underlying, not to mention default risk which makes the formula literally diverge![7]

In our regime-switching model, jumps and default are omnipresent so we have no choice but to value the variance swap as the true path-dependent instrument that it is, that is to say, we value it directly, without proxy or static replication: as the average of daily square returns specified in its term sheet.

How we deal with the default event is by removing the -100 per cent jump from the observation sample of the variance swap (indeed, the term sheet provides that trading of the stock is suspended in such exceptional event), however, we assume the penny stock remains volatile (and very much so!) after the default event. Mathematically speaking, although the stock price reaches almost zero at the time of default and we are no longer able to sample paths of the underlying beyond that point, we simply patch up the variance of the penny stock with the variance that has been realized on the path prior to default. (As variance is homogeneous and independent of the stock price, it can remain strictly positive even in the limit of the stock price reaching zero.) In other words, the variance of the penny stock is a user-defined parameter that we independently require. It is constant, of course, as default is an absorbing regime.

Going back to solutions #3 and #4, and assuming a 100 per cent penny stock volatility, we value the variance swap whose averaging period starts today (February 3rd 2003) and ends at the maturity date of the one-year options (January 16th 2004). The volatility strikes we find are reported in Table 22.

Again, this demonstrates the superiority of the homogeneous model.

Although they seem to agree on everything (except the longer-term out-of-the-money put which everyone had thought was outside our scope), solutions #3 and #4 do not in fact agree on the exact dynamics, as is apparent from their disagreement over the jump out of regime 2. Because of time and space homogeneity, those "exact dynamics" are not confined to a specific region of space or a specific date. The difference they make can be felt as early and as urgently as we wish, provided we find the right instrument to reflect it.

The variance swap happens to be such an instrument. Quoting Philippe Henrotte: "Serious models should attempt to capture the stochastic nature of the jumps... Whereas short term options only depend on the short term jump predictions, the long maturity variance swap will integrate through time the stochastic behavior of the jumps." Remarkably, we did not even need a long maturity swap to discriminate between solutions #3 and #4 as far as the stochastic nature of the jumps was concerned. The one-year variance swap was sufficient. So stability was not achieved yet and solutions #3 and #4 were very different after all. They agreed on everything except the stochastic behavior of the jumps trough time. What looks like fine tuning from the point of view of the vanillas may still conceal a major cleft. (Think how unmindful of all this the local volatility approach is.) Note, however, that the difference is even more pronounced for the two-year variance swap. (See Table 23.)

A supreme conclusion

Our little experiment with the two-regime switching model (a prototype of a jump-diffusion/stochastic volatility/stochastic hazard rate model) has shown that two very different variance swap valuations can coexist with the same given vanilla smile. What this means is that the day when the market price of the variance swap starts deviating from the common view (the vanilla static replication), we will have to start looking for explanations in the spirit of the regime-switching model. This, by itself, sentences to death the static replication argument. As Philippe Henrotte had predicted, "the variance swap prices produced by the local volatility model are bound to collide with the market quotes which do integrate the possibility of jumps. Local volatility models [or, more generally, pure diffusion models] will not survive this collision."

Above all, I think what this "nail in the coffin" has achieved is open our mind to the way calibration should really be thought of, and open the road to actually achieving it. Derivative pricing models and their calibration routines are tools to probe the market and help us interact with it, that is to say, help us price other derivatives and hedge the derivatives whose price we are either reading from the market or generating ourselves. In this respect, every piece of information we can get from the market is valuable and can make a huge difference. Had the quote of the one-year-variance swap been available, it would have lifted the uncertainty between solution #3 and #4 and helped us price the rest of derivatives more accurately and reliably. (Yes, I am implying one must calibrate one's model against the variance swap. We do this fluently, don't you?)

Yet some people insist that the difference, the deviation, the quirk, that the market is kind enough to display for us today and which may hold the key to subsequent, contrasting developments, be trivialized and emptied of all informational content. *This is exactly what the local volatility model does*. Not only must it locally explain (i.e. not explain at all) the variety of option prices

Table 23: Comparison of the values, given by solutions #3 and #4, of the two-year variance swap, starting on 03 Feb 2003 and ending on 21 Jan 2005

	Solution #3	Solution #4
Volatility strike	46.06%	50.67%

that are available to it and must it make up the missing prices that it will subsequently spend even more efforts trying to (even less) explain, but the only outcome of this whole trouble is to entangle itself in the net of the local volatility function (see Figure 7) with no hope of ever being able to reach to the liberal price of a barrier option or of a variance swap. In a word, to calibrate a local volatility model is to calibrate against nothing, for nothing.

Thus, I hope I have shown you that meaningful and useful calibration should, in every way and in every step, proceed to the exact contrary of the local volatility model. It is a hermeneutical "negotiation" process, progressively deciphering the market like we did (credit risk skew, then fat tails, then term-structure of volatility, then term-structure of CDS spreads, and finally term-structure of volatility skew). Since the local volatility model is exactly nothing and doesn't exist, this means I have un-concealed the sway of *being* and of *existence*: philosophy's supreme prize. To put it in Heidegger's words, Being itself shines out through the "clearing" (or should I say, the "cleaning"?).

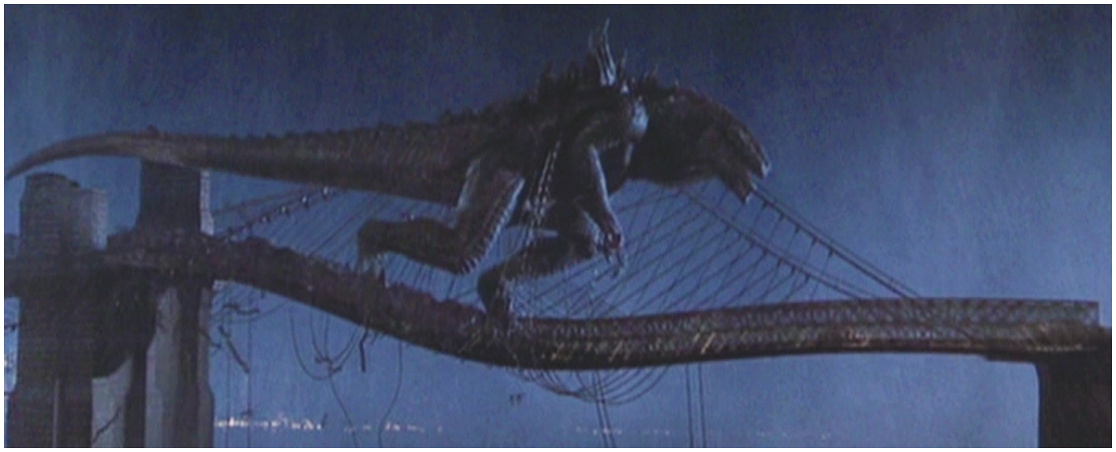

Figure 7: The end of local volatility

FOOTNOTES

1. ITO 33, 36 rue Lacépède, 75005 Paris, numbersix@ito33.com.
2. Elie Ayache, "The Irony in the Variance Swaps", *Wilmott* magazine, September 2006: 16-23. Dominic Brecher, "Pushing the Limits of Local Volatility in Option Pricing", *Wilmott* magazine, September 2006: 6-15.
3. Elie Ayache, Philippe Henrotte, Sonia Nassar and Xuewen Wang, "Can Anyone Solve the Smile Problem?'", *Wilmott* magazine, January 2004: 78-96. Elie Ayache, "What Is Implied by Implied Volatility?", *Wilmott* magazine, January 2006: 28-35. Philippe Henrotte, "The Case for Time Homogeneity", *Wilmott* magazine January 2004: 71-75.
4. Elie Ayache, "Can Anyone Solve the Smile Problem? A Post-Scriptum", *Wilmott* magazine, January 2005: 6-11.
5. Philippe Henrotte, "How Exotic Is the Variance Swap?", *Wilmott* magazine, November 2006: 24-26.
6. To put it in the words of Charles Caiman.
7. See Philippe Henrotte, "How Exotic Is the Variance Swap?", op. cit.