Can anyone solve the smile problem?

Elie Ayache
ITO33 SA, 39 rue Lhomond, 75005 Paris, France,
eMail: NumberSix@ito33.com
Philippe Henrotte
ITO33 SA, 39 rue Lhomond, 75005 Paris, France,
eMail: NewNumberTwo@ito33.com
Sonia Nassar
ITO33 SA, 39 rue Lhomond, 75005 Paris, France, eMail: Sonia@ito33.com
Xuewen Wang
ITO33 SA, 39 rue Lhomond, 75005 Paris, France, eMail: Wang@ito33.com

Abstract
One of the most debated problems in the option smile literature today is the so-called "smile dynamics." It is the key both to the consistent pricing of exotic options and to the consistent hedging of all options, including the vanillas. Smiles models(e.g. local volatility, jump-diffusion, stochastic volatility, etc.) may agree on the vanilla prices and totally disagree on the exotic prices and the hedging strategies. Smile dynamics are heuristically classified as "sticky-delta" at one extreme, and "sticky-strike" at the other, and the classification of models follows accordingly. The real question this distinction is hinging upon, however, is space homogeneity vs. inhomogeneity. Local volatility models are inhomogeneous. The simplest stochastic volatility models are homogeneous. To be able to control the smile dynamics in stochastic volatility models, some authors have reintroduced some degree of inhomogeneity, or even worse, have proposed "mixtures" of models. We show that this is not indispensable and that spot homogeneous models can reproduce any given smile dynamics, provided a step is taken into incomplete markets and the true variable ruling smile dynamics is recognized. We conclude with a general reflection on the smile problem and whether it can be solved.

1 Introduction

The smile problem has raised immense interest among practitioners and academics. Since the market crash in October 1987, the volatilities implied by the market prices of traded vanillas have been varying with strike and maturity, revealing inconsistency with the Black-Scholes (1973) model which assumes a constant volatility. Ever since, a multitude of volatility smile models have been developed. The earliest of the volatility models were the local volatility models[1]. They inferred a volatility dependent on the stock price level and time that accommodates the market price of vanillas within the Black-Scholes framework (Dupire (1994), Derman & Kani (1994), Rubinstein (1994)). Indeed, local volatility models postulate that the underlying follows a lognormal diffusion process equation

$$\frac{dS}{S} = \pi(t)dt + \sigma(S, t)\,dW$$

yielding the following partial differential equation (PDE) for derivative instruments:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2(S,t)S^2\frac{\partial^2 V}{\partial S^2} + r(t)S\frac{\partial V}{\partial S} = r(t)V$$

They are so to speak an extension of the Black-Scholes lognormal diffusion process with constant volatility to a process where the volatility is dependent on both the share price level and time. Under these assumptions, the unique local volatility surface is backed out through forward induction from the smile of vanilla option prices. Once the local volatility surface is known, it is used to value and hedge any type of option on the same underlying. The implied volatility of an option with a given strike and a given maturity can be seen as an average over all local volatilities that the underlying may have as time evolves until the maturity date. Local volatility models accommodate the smile and are theoretically self-consistent as it is possible to hedge, and as a matter of fact perfectly replicate options in order to price them, as done in the Black-Scholes framework. In other words, they retain the market completeness.

Unfortunately, as shown in Figure 2, the shape of the local volatility surface, inferred from the market vanilla smile represented in Figure 1 may sometimes look very surprising and unintuitive, with no easily explainable trend either along the underlying share price direction or in the time direction. For instance, far in the future, local volatilities are roughly constant, i.e. the model predicts a flattening of the smile, which seems inconsistent with the omnipresence of the skew or smile observed for the last 15 years. Not mentioning the numerical efforts in order to interpolate and extrapolate the sparse empirical smile data, then to smooth the surfaces of interest. This is computationally known as an "ill-posed inverse problem."

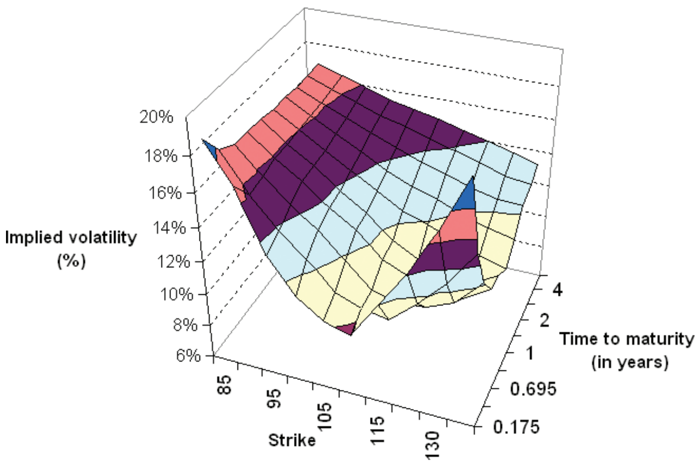

Figure 1: Implied volatility surface inferred from vanilla options market prices. Source: S&P 500 index on October 1995 [1]

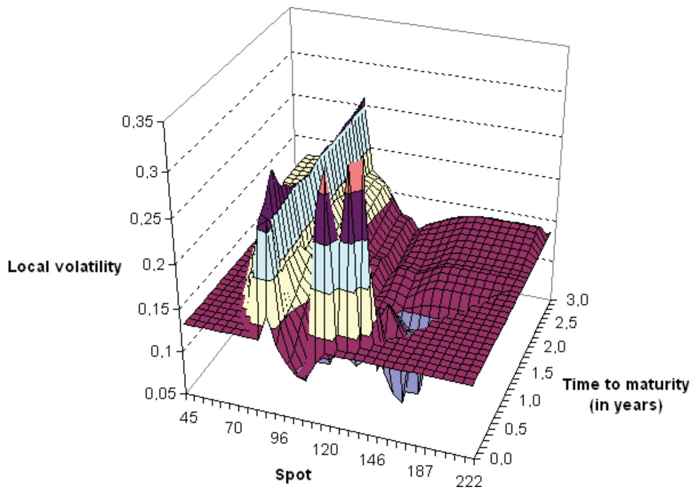

Figure 2: Local volatility surface inferred from vanilla options market prices

2 Is the local volatility model really a model?

2.1 The sirens of "tweaking"

When you think about it, the local volatility models just provide numerical methods for finding a volatility surface $\sigma(S,t)$ that fits the market data of the options, $C(K,T)$, by exploiting the mechanics of the pricing equations or the PDEs. To our mind, they do not really provide a (physical) explanation of the smile phenomenon. Dupire has not discovered a smile model. His great discovery was the forward PDE for pricing vanilla options of different strikes and different maturities in one solve. Tweaking the diffusion coefficient in the Black-Scholes PDE in order to match a given set of vanilla option prices is reminiscent of the method of "epicycles" which was the only way to account for the movement of celestial bodies when the real scientific explanation was lacking. (See Henrotte (2004) in the present issue of Wilmott Magazine for a defence of homogeneous models against the dangers of "tweaking" and Ayache (2001) for an early version of the argument). Local volatility models do not intend to explain the volatility smile problem by introducing new dynamics for the underlying stock. And by "new dynamics" we mean something original, like jumps or stochastic volatility or default. Suggesting that smiles are caused by jumps in the underlying or by stochastic volatility (or both) not only sounds realistic and informative, but may qualify as an explanation. Think how incredible it must sound, in comparison, that volatility should locally rise at a given point in time and space, then drop at some other point, for the sole purpose of matching today's option prices! It really sounds as if somebody was trying to force an interpretation in terms of local volatility on a phenomenon which has different and deeper origins. As a matter of fact, Jim Gatheral (2003) has provided what is to our mind the right interpretation of local volatility. He shows that local volatility is but the local expected variance of the underlying in general stochastic volatility models (that is to say, in "realistic" models).

2.2 The "natural" local volatility surface

Another reason why we should be suspicious of the local volatility model and why it falls in a class of its own (which may simply be the class of "not being a model") is that it is non parametric in essence or else arbitrarily parametric. Dupire's derivation essentially shows that any smile surface can be fitted by local volatility provided the model is non parametric, and it basically provides the non parametric formula. On the other hand, methods consisting in parameterizing the local volatility surface a priori (through spline functions or any other convenient representation), and in fitting the smile surface by minimization of a loss function (Coleman, Li, Verma (1999), Jackson, Sueli, Howison (1998)), suffer from the arbitrariness of the representation, particularly the arbitrariness of the behaviour of local volatility at the boundaries of the domain. Proponents of such approaches are always at pains trying to justify their favourite representation of the local volatility surface on grounds of its intuitive appeal or physical realism or what have you. It is not uncommon that they maximize some entropy or some regularity criterion while minimizing their loss function, the underlying idea being that nature somehow favours smoothness and regularity. In a word, they look for the "most natural local volatility function" matching the option prices. One wonders what that means.

2.3 Arbitrage-free interpolators

Jump-diffusion and stochastic volatility models, by contrast, lend themselves naturally to the routine of fitting the option prices by minimization of a loss function, as they are "naturally parameterized" by the coefficients of the process (for instance the intensity of jumps and the parameters of the jump size distribution in the Merton model (1996); the volatility of volatility, its mean reversion, its correlation with the underlying in Heston (1993), etc.). As research on local volatility models was getting more and more entangled in issues purely computational (finding the smoothest arbitrage-free interpolation, maximizing the right regularity criterion, etc. (Andersen, Brotherton-Ratcliffe (1998), Avellaneda, Carelli, Stella (2000), Bodurtha, Jermakyan (1999), Coleman, Li, Verma (1999), Jackson, Sueli, Howison (1998), Kahale (2003), Lagnado, Osher (1997), Li (2001)), and was drifting farther and farther away from the "physics" of the problem, it so happened one day that our computational expert asked our financial theorist what to his mind the "most natural local volatility function" could be, suited for a given smile. Undecided between many attractive numerical alternatives, our man was seeking guidance from the underlying "physics." Not surprisingly, the financial theorist suggested he looked at local volatility surfaces "such as might have been produced by models of jumps in the underlying, or stochastic volatility, etc." In other words, the suggestion was that the best solution to the *numerical problem* of inferring the smoothest, most regular, and arbitrage-free local volatility surface was to pretend that the option prices were generated by a jump-diffusion, stochastic volatility model! If you are so keen on local volatility, then indeed jump-diffusion/stochastic volatility models can be sold to you as "financially meaningful, arbitrage-free, super-interpolators." This is just the rehearsal of Gatheral's point. Only the question now becomes: If you go this far, why bother with local volatility any longer? For market completeness perhaps?

2.4 "Local" everything?

More to the point: Why hasn't anybody ever tried to fit a *non parametric* jump-diffusion or stochastic volatility model to option data? Why is everybody busy searching for constant (or perhaps only time-dependent) parameters in Heston, Merton, SABR (Hagan, Kumar, Lesniewski, Woodward (2002)), and nobody has proposed that both the diffusion coefficient and the jump coefficients, or both the volatility of volatility and the correlation coefficient, may become non parametric functions of time and space? One possible answer is that the model would very rapidly become computationally infeasible. With the implication that the reason why non parametric inference is actually done in the pure diffusion model and in no other model (or, in other words, the reason why local volatility models simply exist) is that it *can* be done. Hardly a proud conclusion. It means that local volatility models are just a temporary diversion outside the tracks of true progress. Another possible answer is that the continuum of vanilla call prices $C(K, T)$ will no longer be sufficient for calibration purposes when more than one parameter of the pricing equation are made a function of time and space. One would require an additional continuum of market prices, not redundant with the vanillas. Why not add, for instance, the continuum of prices of American one-touches $OT(B, T)$ of different barrier levels and maturity dates? As it happens, this might ensure agreement with the market prices of barrier options, an urgent problem for all exotic options trading desks.

We will have a lot more to say later about additional market information that we may require in the calibration phase. Enough to observe for the moment that the literature is not treating the showdown between local volatility and the other smile models properly. Like we said, local volatility is not a model, it is the tweaking of Black-Scholes. And the tweaking could equally be applied to Heston, or Merton, or any alternative smile model, if only we had the computational guts to do so. It seems the literature is standing at a methodological crossroads between the tough computational decision to involve additional instruments in the calibration—no matter the specific model or its parametric/non parametric status—and the temptation to develop specific models just for their own sake and the sake of an original name, then to go check whether they predict the right exotic option prices, or the right smile dynamics. At any rate, it is unfortunate that external issues, such as tractability, solvability, elegance of formulation, etc., should be the ultimate guides of scientific research. **We motivate our paper by situating it precisely at this crossroads**.

As a matter of fact, an attempt could be made at the calibration of a jump-diffusion model with local diffusion component and local jump intensity. Indeed, a natural extension of the Black-Scholes diffusion model in the equity world is to include the risk of default in the pricing problem of equity derivatives subject to credit risk, like convertible bonds. This introduces the hazard rate function $\lambda(S,t)$ in the usual partial differential equation:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2(S,t)S^2\frac{\partial^2 V}{\partial S^2} + (r(t) + \lambda(S,t))\,S\frac{\partial V}{\partial S} = r(t)V + \lambda(S,t)X$$

where X is the loss given default, and means we would have to calibrate the hazard rate function, on top of the volatility function, to available market data. The obvious candidates are the continuum of vanilla option prices $C(K,T)$ and the continuum of credit default swap spreads as a function of present stock price and maturity $CDS(S,T)$. See Andersen, Buffum (2002) for an example of such joint calibration. Note, however, that Andersen's procedure is parametric in that he proposes simple parametric representations of $\sigma(S,t)$ and $\lambda(S,t)$.But nothing stops us, in theory, from extending the forward induction argument of Dupire, or the Fokker-Planck equation approach of Klopfer and Tavella (2001), to the case where the probability density diffuses under the Brownian component as usual and "leaks" into the state of default through the Poisson intensity of the default jump process, and from inferring $\sigma(S,t)$ and $\lambda(S,t)$ non parametrically.

2.5 The mirage of the vanillas

The conclusion we draw from our first bash at local volatility models is twofold. First, local volatility is not a model. It is the "corruption" of a model[2] and the corruption, for that matter, can spread over to all the other models. At best, local volatility can be seen as a shorthand or an interpretation: it is the local expected variance of some deeper and more realistic dynamics. (Think of Ehrenfest's theorem which interprets the classical mechanical variables as expectations of the "true" quantum mechanical observables). Second, when thinking about the other models (jump-diffusion, stochastic volatility, etc.), one should keep in mind that they can be made "local" too. For once one recognizes that vanilla option prices will not be sufficient for calibration in that case, one realizes that there is nothing special about the vanillas anyway. The only reason why authors of jump-diffusion, stochastic volatility, or universal volatility models insist on fitting them to the vanillas is that they followed in the steps of the local volatility approach and vanillas were the obvious calibration candidates there.

We also fear the real reason might be that vanillas alone admit of analytical solutions in the models they propose, or even worse, that they have precisely grabbed the models which offered analytical solutions for the vanillas to begin with. We would love to see some of these authors calibrate their jump-diffusion, stochastic and universal volatility models, to a handful of options of *significantly different payoff structures*: vanillas, barriers, cliquets, credit default swaps, etc. As a matter of fact, vanilla options can be the poorest candidate for encapsulating the information about the stochastic process, when processes more general than a diffusion are considered. That our problem is called the "smile problem" is no reason why the calibration of the model, or even its whole intention, should revolve around the vanillas. And that vanilla option trading is the ancestor of exotic option trading, or that traders are accustomed to envision alternative stochastic processes in terms of the vanilla smiles they generate, is an even worse excuse. But again, SABR would not be SABR if it did not allow the expansion of the Black-Scholes implied volatility (in other words the vanilla smile) in terms of the parameters of the process, and Heston would not be Heston, or Hull and White (1988) Hull and White, if . . .

3 Formulation of the smile problem

3.1 The real smile problem

Not only can we argue, *on a priori grounds or from a purely methodological point of view*, that the local volatility model is not a model, but it also demonstrably fails *as a model* of option smiles. Indeed the real smile problem is not how to fit the vanillas or how to price them! Straightforward spline interpolation does that very nicely. The real smile problem is the pricing of exotic options and more generally the hedging of all kinds of options, including the vanillas, under dynamic assumptions at variance with the Black-Scholes model. As noted by almost everybody, the local volatility model fails miserably on both counts. Both the barrier option price structure and the dynamic behaviour of the smile predicted by a vanilla-calibrated local volatility model diverge from empirical observation (Lipton, McGhee (2002), Hagan, Kumar. Lesniewski, Woodward (2002)). "The failure of the local volatility model, writes Hagan, means that we cannot use a Markovian model based on a single Brownian motion to manage our smile risk." We need to assume an independent process for volatility. This opens the door to stochastic volatility models, and more generally, to all kinds of alternative dynamics that have been proposed over time as a replacement of Black-Scholes.

Perhaps the most important aspect of the smile problem today is to find a way of discriminating between all the alternative proposals to solve it. This is the symptom of a science in crisis, not just the symptom of a problem. Definitely the accurate pricing of exotics and the soundness of the hedging strategy are good selection criteria. To put it in Lipton's words (2002):

> "We describe a series of increasingly complex models that can be used to price and hedge vanilla options consistently with the market. **We emphasize that, although all these models can be successfully calibrated to the market, they produce very different hedging strategies.** [. . .] A number of models have been proposed in the literature: the local volatility models of Dupire (1994), Derman & Kani (1994) and Rubinstein (1994); a jump-diffusion model

of Merton (1976); stochastic volatility models of Hull and White (1988), Heston (1993) and others; mixed stochastic jump-diffusion models of Bates (1996) and others; universal volatility models of Dupire (1996), JP Morgan (1999), Lipton & McGhee (2001), Britten-Jones & Neuberger (2000), Blacher (2001) and others; regime switching models, etc. [. . .] Too often, these models are chosen ad hoc, for instance, on the grounds of their tractability and solvability. **However, the right criterion, as advocated by a number of practitioners and academics, is to choose a model that produces hedging strategies for both vanilla and exotic options resulting in profit and loss distributions that are sharply peaked at zero."**

This is the most cogent formulation of the smile problem we know of.

3.2 Indeterminateness of the conditionals

We shall quickly review the smile models which are most representative of today's smile literature, but let us first investigate the reason why smile models of different stochastic structure may not agree on exotic option pricing or the option hedging strategies (a.k.a. "smile dynamics") even when calibrated to the same vanilla smile. The picture becomes clear when we have a look at the way the calibration is carried out. Denoting $A^{i,j}_{i_0,j_0}$ the price at state i_0 and time j_0 of a security paying off \$1 at state i and future time j (a.k.a. Arrow-Debreu security), it can be related to the vanilla call option prices in the following way:

$$A^{i,j}_{i_0,j_0} = \frac{C(K_{i+1},T_j) - 2C(K_i,T_j) + C(K_{i-1},T_j)}{\Delta K^2} \tag{1}$$

In continuous time and space this is expressed by

$$p(S,t;K,T)e^{-\int_t^T r(s)ds} = \frac{\partial^2 C(S,t;K,T)}{\partial K^2}$$

where $p(S,t;K,T)$ is the transition probability density from initial state and time (S,t) to (K,T). Introducing the vector notation:

$$A^{j}_{i_0,j_0} = \begin{bmatrix} A^{1,j}_{i_0,j_0} \\ A^{2,j}_{i_0,j_0} \\ \vdots \\ A^{N,j}_{i_0,j_0} \end{bmatrix} \tag{2}$$

and the matrix notation:

$$A^{j+1}_{j} = \begin{bmatrix} A^{1,j+1}_{1,j} & A^{2,j+1}_{1,j} & \cdots & A^{N,j+1}_{1,j} \\ A^{1,j+1}_{2,j} & A^{2,j+1}_{2,j} & & A^{N,j+1}_{2,j} \\ \vdots & \vdots & \ddots & \vdots \\ A^{1,j+1}_{N,j} & A^{2,j+1}_{N,j} & \cdots & A^{N,j+1}_{N,j} \end{bmatrix} \tag{3}$$

Up to a discounting factor, this is the matrix of conditional transition probabilities from states at date j to states at date $j+1$. (Crucially, the assumption here is that states of the world are just states of the underlying).

The conditional probability rule yields the following equation:

$$\left(A^{j+1}_{i_0,j_0}\right)^T = \left(A^{j}_{i_0,j_0}\right)^T A^{j+1}_{j} \tag{4}$$

Without any further information about the structure of the stochastic process, this is the only constraint that the prices of vanilla options today impose on the matrix of conditional probabilities. Infinitely many matrices solve that equation of course. In a continuous diffusion framework this forward equation becomes

$$\frac{\partial p}{\partial T} + \frac{\partial(rKp)}{\partial K} - \frac{1}{2}\frac{\partial^2(\sigma^2K^2p)}{\partial K^2} = 0 \tag{5}$$

and shows why the knowledge of the prices of Arrow-Debreu securities maps the diffusion process $\sigma(K,T)$ completely.

3.3 Smile dynamics and model-dependence

To repeat, the only information contained in the set of vanilla option prices $C(K,T)$ of different strikes and different maturities, independently of any model, is the map of transition probabilities from present day and present spot to whatever future time and future spot we are looking at. This says nothing about the conditional transition probabilities from a future date to a farther future date. Additional information is needed to help determine those conditionals. In theory, we would need the knowledge of all "forward smiles," in other words, the future prices of all vanilla options as seen from all possible states of the world, *not mentioning that the underlying stock price may not be the only state variable* (in stochastic volatility models, typically).

Choosing a particular model for the underlying dynamics definitely adds some structure. It is a form of parametrization of this totally non parametric picture. The only "structure" that the local volatility model adds consists in removing the need for market information beyond the vanilla option prices in the fully non parametric case. The "matrix" of conditionals is fully determined in that case, and there is no spatial state variable other than the underlying. Alternative models such as jump-diffusion, or stochastic volatility, or universal volatility models, also dramatically reduce the degrees of freedom in the choice of the conditionals, particularly so when the coefficients of the given process are constant, or time-dependent, or assume some parametric form. Now think how different the structure of conditionals that they imply can be, compared to the pure diffusion case (e.g. the possibility of jumping and hitting a barrier in between future dates, the addition of another state variable indexing the forward smiles, etc.), yet their authors calibrate them to the vanillas just the same! In a sense, the local volatility model is more honest than the other models with regard to the conditionals. You just know there is nothing you can do. In the other models, by contrast, you calibrate a bunch of *constant parameters* in what seems to be a legitimate calibration move—typically you calibrate them to the vanillas—and this sets for you all the conditional structure. Hardly can a result be more model-dependent!

3.4 Our preferred model

The reason why the local volatility model, the jump-diffusion models, the stochastic volatility models, or more generally the "universal volatility models," may agree or not agree among each other or with the market on the prices of barrier options or forward starting options, is that each model imposes a specific smile dynamics, or structure of conditionals. We claim that this smile dynamics should not be imposed by the model, but inferred from the market. However, we have to pick a certain framework.

Calibration, pricing and dynamic hedging cannot be totally model-independent, even though model-independence should always act as a "regulative ideal" in our research program. **We shall pick the framework with the features that everybody knows today are essential for explaining the smiles.** We know we need jumps (if only to account for shorter dated smiles and default risk) and we know we need stochastic volatility (to account for longer dated smiles and to acknowledge the very *raison d'être* of option markets and market-makers). Our discussion of local volatility and Henrotte's powerful statement[3] should steer us away from inhomogeneous models. **The coefficients of our stochastic process shall be constant**. However, we have learnt from the unhappy story of the conditionals that market option data, other than the vanillas, must be included in the calibration procedure. Under no circumstance shall we be prevented from doing so by what Henrotte describes, in other people's cases, as "a very somber agenda": the need to produce closed form or quasi closed form pricing solutions. **Our pricing equations shall be solved by numerical algorithms**. For all these reasons, chiefly the fact that model names have traditionally been associated with the discovery of analytical solutions, our model shall bear no particular name. **We shall call it "Nobody's model."**

3.5 Including exotics in the calibration

On the calibration side, we have noted that the value of barrier options is sensitive to the flux of probability across the barrier (jumps, and volatility dynamics up to the barrier). The value of forward starting options, on the other hand, is directly linked to the conditional transition probabilities, or forward smiles. In other words, both depend on what extra structure the matrix of conditional transition probabilities may have, on top of the constraint given by the spot vanilla smile. This designates simple barrier options like the one-touch or American digital, and the forward starting options as the natural candidates for extending our calibration set and helping determine the smile dynamics[4]. Traders accustomed to Derman's (1999) classification of smile dynamics in terms of "sticky-strike" or "sticky-delta volatility regimes" know that the delta of the vanillas is very much dependent on the type of volatility regime the market is in. Derman's study produces evidence that both kinds of regimes have obtained over time within a single market. Depending on the regime you think the market is in, you make the following adjustment to your Black-Scholes hedge.

When $\sigma_{imp}(S,t,K,T)$ is the implied volatility for a European style option we have :

$$C(S,t,K,T) = C_{BS}\left(S,t,K,T,\sigma_{imp}(S,t,K,T)\right) \tag{6}$$

The delta-hedge becomes a combination of Black-Scholes delta and a correction term due to the regime of movement of the smile with a moving underlying:

$$\Delta = \frac{\partial C}{\partial S} = \frac{\partial C_{BS}}{\partial S} + \frac{\partial C_{BS}}{\partial \sigma_{imp}} \cdot \frac{\partial \sigma_{imp}}{\partial S} \tag{7}$$

We claim that nobody should be in a position to decide which particular smile dynamics will prevail. It is really like guessing a price (as Marco Avellaneda once rightly observed in a financial workshop at NYU). Only the market can provide such information. We are saying that *your wrong guess about the smile dynamics can generate an immediate arbitrage opportunity against you, if somebody picks the right security to trade against you*. As a matter of fact, all FX option traders are aware of the existence of such a security! It is the barrier option, the simplest instance of which is the one-touch.

Different projected evolutions of the vanilla smile lead to different spot prices of barrier options in the FX traders' minds, because they think of the future cost of unwinding the vanilla static hedge that they have set up against the barrier option. This insight can be further refined and made rigorous in a fully dynamic hedging picture. (Indeed the vanilla static hedge that those FX exotic option traders have in mind is not always consistent with the smile dynamics they project. For instance they immunize the vega, the vanna and the volga of the barrier option with a static combination of vanillas, yet they derive their hedging ratios from the Black-Scholes model wich assumes constant volatility[5]).

The price structure of the one-touches contains implicit information about the smile dynamics, therefore about the delta you should be using to hedge the vanilla options! So does the price structure of the forward starting options. This is why the one-touches and the forward starting options must be included in the calibration.

In conclusion, the exotic option pricing problem and the problem of smile dynamics are intimately linked, and the pricing/hedging model cannot dispense with including exotic options in the calibration.

4 A quick review of representative smile models

4.1 Stochastic volatility

In stochastic volatility models (Heston (1993), Hull & White (1998)), volatility is itself stochastic and follows some mean reverting process with its own volatility and correlation with the underlying share. The stochastic volatility models can be seen as modelling the option price as an average of the Black-Scholes prices with respect to volatility. This model is essential for the pricing of longer-dated options which are most sensitive to volatility changes. It avoids the scale effect observed in long-term local

>

volatilities. Least square fit is used to search for model parameters to match observed market prices.

The problem with stochastic volatility models is that the derivative instrument is exposed to volatility risk on top of market risk, and the underlying cannot hedge both Brownian motions.

The Heston model is, for instance, given by the following risk-neutral process:

$$\frac{dS}{S} = rdt + \sqrt{v}dW$$

$$dv = \kappa\,(\theta - v)\,dt + \varepsilon\sqrt{v}dZ$$

where the volatility process and the underlying process are correlated through a correlation coefficient ρ. And the pricing equation is given by:

$$\frac{\partial V}{\partial t} + \frac{1}{2}v\left(S^2\frac{\partial^2 V}{\partial S^2} + 2\rho\varepsilon S\frac{\partial^2 V}{\partial S\partial v} + \varepsilon^2\frac{\partial^2 V}{\partial v^2}\right) + rS\frac{\partial V}{\partial S} + \kappa\,(\theta - v)\,\frac{\partial V}{\partial v} = rV$$

The calibration of the model consists in finding parameters of the volatility process: κ (mean reversion), θ (long term volatility), ϵ (volatility of volatility), ρ (correlation between the volatility process and the underlying process) as well as initial volatility state v_0, such that option market data is fitted.

4.2 Jump-diffusion

Jump-diffusion models (Merton (1996)) add jumps and crashes to the standard diffusion process of the underlying. They intend to reproduce the underlying dynamics more realistically and to capture the strong smile exhibited by short-dated options. The underlying share price follows a risk-neutral process governed by the following equation:

$$\frac{dS}{S} = (r - \lambda m)\,dt + \sigma\,dW + \left(e_j - 1\right)dN$$

where N is a Poisson process with frequency λ, W is a wiener process independent of N, j is a random logarithmic jump size with pdf $\phi(j)$ and m is the expected value of $e_j - 1$.

The problem again is that the Black-Scholes continuous hedging argument breaks down in the presence of jumps.

Some other models lay jumps on top of stochastic volatility models (Bates (1996)).

4.3 Universal volatility

4.3.1 Blacher

The universal volatility model of Blacher is described by the following risk-neutral process:

$$\frac{dS}{S} = rdt + \sigma\left(1 + \alpha(S - S_0) + \beta(S - S_0)^2\right)dW$$

$$d\sigma = \kappa\,(\theta - \sigma)\,dt + \varepsilon\sigma\,dZ$$

The volatility σ follows a mean reverting process to level θ, correlated with the underlying process via ρ.

It is worthy of note that Blacher motivates his universal volatility model for reasons almost opposite Hagan, et al (2002). Like Hagan, he speaks for stochastic volatility models. However, he notes that although the "smile is stochastic, simple stochastic volatility models [such as Heston's] do not predict a systematic move of the relative smile when the spot changes." "Not what we observe in the market," he says. "This means hedging discrepancies, starting with a wrong delta." In other words, Blacher is noting that space homogeneous models like Heston's follow the sticky-delta rule. The "relative smile" they imply, i.e. the smile with respect to moneyness or delta of the option, is unchanged when the underlying spot changes. Yet Blacher wishes that the vanilla smile may not always move coincidentally with the underlying. He claims control over the smile dynamics. In order to achieve this, he has no choice but to re-introduce inhomogeneity in the spot homogeneous stochastic model.

He writes: "α, the slope of the deterministic part, creates skew and governs the change of ATM implied vol with respect to change of underlying. β, the curvature of the deterministic part, creates smile curvature and governs the change of the slope of the smile curve with respect to change of underlying."

Note that SABR also breaks the homogeneity of degree 1 by allowing values for β different from 1, in the risk-neutral process:

$$dF = \alpha F^{\beta}\,dW_1$$
$$d\alpha = v\alpha dW_2$$

F is the forward price, α its volatility, v the volatility of volatility, and dW_1 and dW_2 are Wiener processes correlated through:

$$\langle dW_1, dW_2\rangle = \rho\cdot dt$$

4.3.2 Lipton

Lipton (2002), on the other hand, argues for his universal volatility model on grounds of its adequacy for pricing barrier options. He writes:

> "A properly calibrated universal model matches the market [of barrier options] much closer than either local or stochastic volatility models, which tend to sandwich the market. [. . .] While both local and stochastic volatility models produce price corrections [for barrier options] in qualitative agreement with the market, only a universal volatility model is capable of matching the market properly. In our experience, this conclusion is valid for almost all path-dependent options."

By "properly calibrated universal model" Lipton means "calibrated to the vanillas." On the specific topic of calibration he otherwise notes: "Because of its complexity, the universal volatility model can be solved explicitly only in exceptional cases (which are of limited practical interest). [. . .] The model calibration, of course, is a different matter."

Lipton's risk-neutral stochastic process is given by:

$$\frac{dS}{S} = (r - \lambda m)dt + \sqrt{v}\sigma_L(t,S)dW + \left(e^j - 1\right)dN$$

$$dv = \kappa\,(\theta - v)\,dt + \varepsilon\sqrt{v}dZ$$

And the pricing equation is given by:

$$\frac{\partial V}{\partial t} + \frac{1}{2}v\left(\sigma_L^2(t,S)S^2\frac{\partial^2 V}{\partial S^2} + 2\rho\varepsilon\sigma_L^2(t,S)S\frac{\partial^2 V}{\partial S\partial v} + \varepsilon^2\frac{\partial^2 V}{\partial v^2}\right)$$
$$+ (r - \lambda m)S\frac{\partial V}{\partial S} + \kappa\,(\theta - v)\,\frac{\partial V}{\partial v} + \lambda\int_{-\infty}^{+\infty} V\left(e^j S\right)\phi(j)dj = (r + \lambda)V$$

where $\sigma_L(t,S)$ is the local volatility part, κ the mean reversion of volatility, θ the long term volatility, ϵ the volatility of volatility, ρ the correlation between the volatility process and the underlying process, λ the intensity of the Poisson jump process, $j > 0$ the random logarithmic jump size with PDF $\phi(j)$), and m the expected value of $e^j - 1$.

4.4 Conclusion

In conclusion of our review of existing smile models, let us retain the following fact. The local volatility model and the stochastic volatility model stand at opposite extremes. The first is inhomogeneous, the second is homogeneous. Neither one predicts the right smile dynamics or produces the right barrier options prices. Only the universal volatility model, which allows explicit control over the smile dynamics (by re-introducing inhomogeneity and by mixing local volatility behaviour with stochastic volatility behaviour), manages to fit the smile dynamics (Blacher) and at the same time to fit the barrier option prices (Lipton, McGhee (2002)).

Let us then solemnly pose the question: "Is *the recourse to inhomogeneity really indispensable?*" Or again: "*Given our plea for inclusion of the exotics in the calibration and our credo in homogeneous models, can we also claim control over the smile dynamics?*"

5 Numerical illustrations of the smile problem

We will try to answer that big question by way of practical examples rather than fundamental theorizing. The examples will also serve the purpose of illustrating the smile problem, namely that models of different stochastic structure may very well agree on the vanilla smile yet completely disagree on the exotics and smile dynamics. Instead of solving Heston's model, or Dupire's model, or Lipton's model, we will build up our series of examples from a simple instance of the "model with no name," the model we have called "Nobody's model."

5.1 The calibration issue

5.1.1 Baby examples

First, we consider a simple jump-diffusion model where the underlying diffuses with a constant Brownian volatility and may incur two jumps of fixed size and constant Poisson intensity. We call this simple stochastic structure "Baby1."

For illustration, we consider a Brownian volatility component of $v = 7\%$, an upward jump of size $y_1 = 10\%$ and intensity $\lambda_1 = 0.40$ and a downward jump of size $y_2 = -25\%$ and intensity $\lambda_2 = 0.2$. Table 1 summarizes the parameters of Baby1.

The probabilities of jump are given in the risk-neutral measure. Consequently, we can compute the vanilla option prices generated by this process and re-express them in Black-Scholes implied volatility numbers (see Table2), thus producing the smile. The interest rate is $r = 2\%$ and the underlying spot is $S = 100$.

Note that the smile is steepest for shorter dated options, and tends to flatten out for longer terms (see Figure 3). We can see this simple model as a discretization of the "traditional" jump-diffusion models (e.g. Merton) with a probability distribution of jump sizes.

Volatility smiles can alternatively be represented as a function of the option delta and maturity rather than its strike and maturity . This is the origin of the appellation "sticky-strike" and "sticky-delta." Smiles

TABLE 1: BABY1 PARAMETERS

Brownian Diffusion	7.00%
Jump size	**Jump intensity**
−25%	0.2
10%	0.4

TABLE 2: VOLATILITY NUMBERS IMPLIED BY BABY1

Strike \ Maturity (years)	0.16	0.49	1
80	30.67%	22.20%	18.97%
85	27.41%	20.97%	18.33%
90	22.12%	18.47%	17.19%
95	15.47%	15.32%	15.70%
100	10.90%	12.96%	14.32%
105	11.69%	12.12%	13.37%
110	13.67%	12.16%	12.83%
115	14.48%	12.42%	12.58%
120	15.79%	12.73%	12.49%
130	17.37%	13.44%	12.56%
140	18.74%	14.08%	12.77%

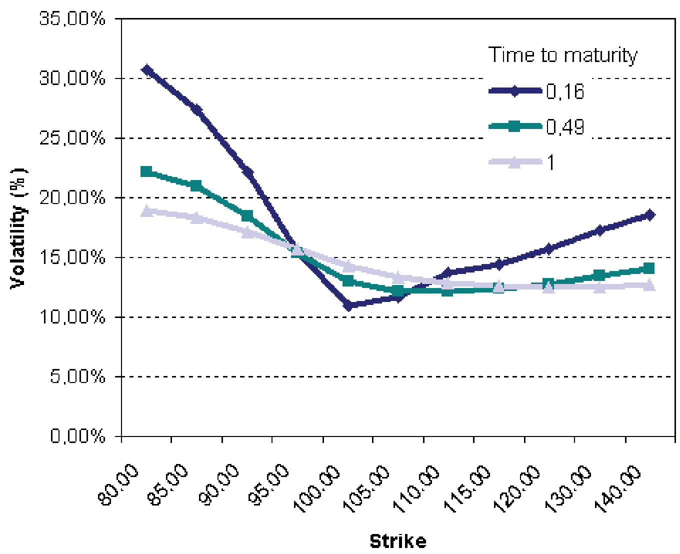

Figure 3: Volatility smile generated by Baby1 against strike price for three different expirations and underlying spot price of 100

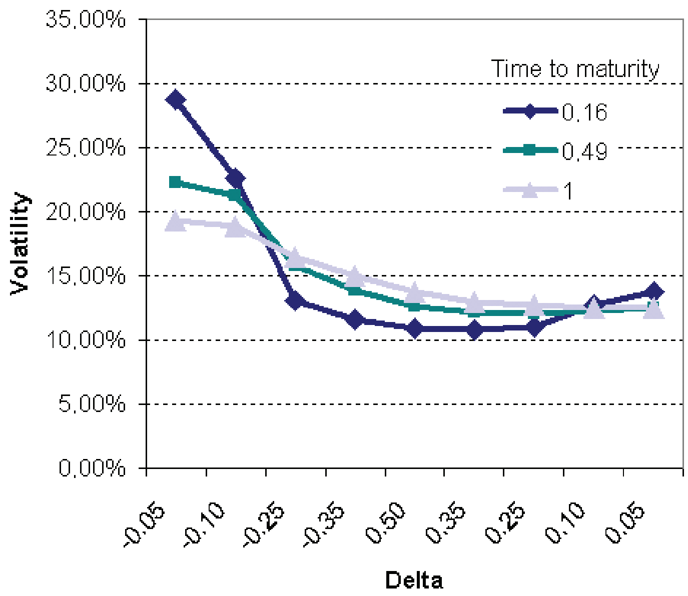

Figure 4: Volatility smile generated by Baby1 against delta for three different expirations and underlying spot price of 100

that are a function of the moneyness of the option are sticky-delta. Their representation in the delta/maturity metric is invariant when the underlying moves. Figure 4 shows the alternative graph of our smile in that metric.

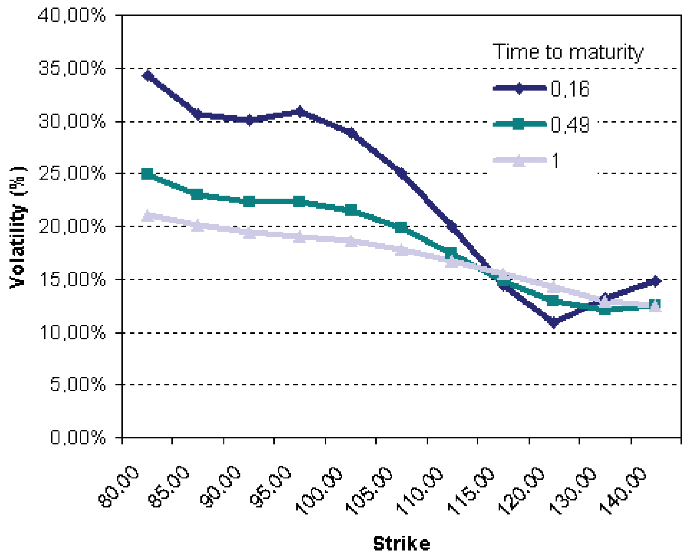

Figure 5: Smile produced by Baby1 against strike price for three different expirations and underlying spot price of 120

We re-compute our smile for $S = 120$ (Figure 5). As our jump-diffusion model is homogeneous and volatility and jump sizes relate to proportional changes of the underlying, the resulting smile surface is sticky-delta. It is unchanged in the delta/maturity metric, and it moves along with the underlying in the strike/maturity metric.

Next, we consider another simple stochastic structure that we call "Baby2." The volatility of the Brownian component is now stochastic and can assume two states, or regimes. The transitions, or jumps, between the two volatility states are caused by Poisson processes of constant intensity. At least two Poisson processes are needed to secure the transition from Regime 1 to Regime 2 and back. As Brownian volatility jumps between regimes, the underlying may simultaneously incur a jump of fixed size. This builds in correlation between jumps in the underlying (or return jumps) and volatility jumps. By convention, Regime 1 is the present regime. You can think of Baby2 as a simplification of stochastic volatility models with correlated return jumps and volatility jumps.

We then propose the following. We shall use Baby2 to try to fit the vanilla smile generated by Baby1. Note that Baby1 admits of five free parameters (the Brownian diffusion coefficient, the two jump sizes and the two jump intensities) and Baby2 of six (the diffusion coefficients in the two regimes, the two inter-regime jump sizes and the two jump intensities).

Calibration of Baby2 is achieved by searching for the six parameters by least squares fitting of the option prices produced by Baby1. The calibration results are shown in Figures 6 and 7 and the set of parameters is summarized in Table 3. Then we see how Baby1 and Baby2 price a given barrier option.

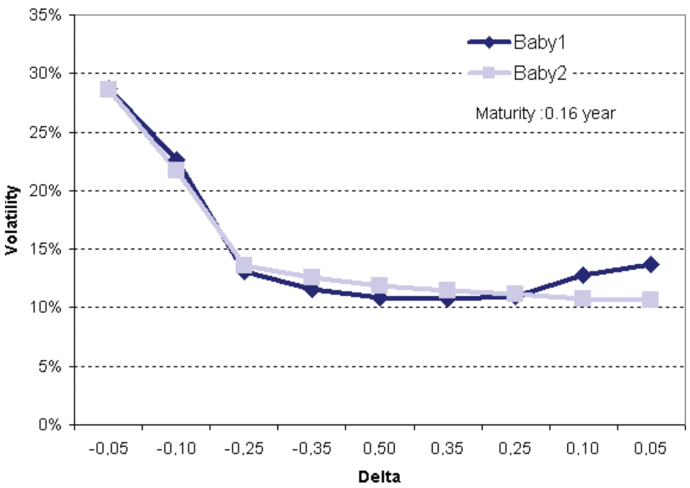

Figure 6: Comparison of the implied volatility curves of Baby2 and Baby1 for 0.16 year maturity

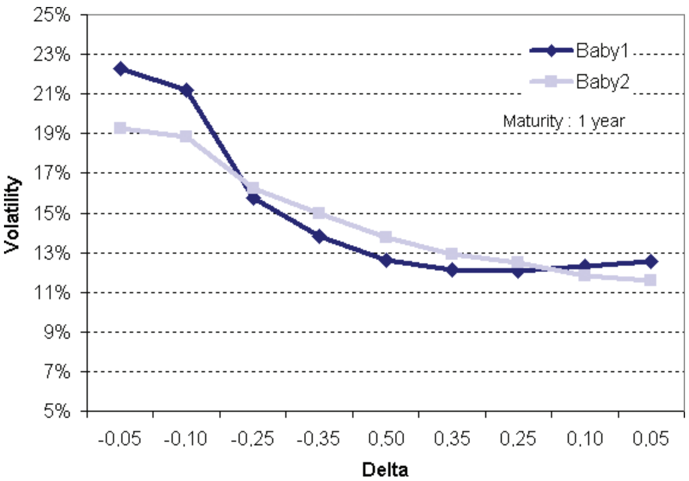

Figure 7: Comparison of the implied volatility curves of Baby2 and Baby1 for a maturity of 1 year

TABLE 3: BABY2 PARAMETERS WHICH BEST FIT THE VANILLA SMILE GENERATED BY BABY1 (TABLE2)

	Brownian Diffusion	
Regime 1	10.02%	
Regime 2	8.44%	
	Jump size	**Jump intensity**
Regime 1→Regime 2	-28.07%	0.1395
Regime 2→Regime 1	0.24%	0.3947

TABLE 4. COMPARISON OF THE PRICES GENERATED BY BABY1 AND BABY2 FOR DIFFERENT 6-MONTH-MATURITY OPTIONS

		Call 100	Call 107	Put 93
Baby1	**Price**	4.12	1.28	1.58
	Implied volatility	12.96 %	12.07%	16.58%
Baby2	**Price**	4.22	1.25	1.51
	Implied volatility	13.31%	11.93%	16.24%

TABLE 5. CALL 100 UP & OUT AT 107, OF MATURITY SIX MONTHS PRICED BY BABY1 AND BABY2

	Price
Baby1	0.74
Baby2	0.49

TABLE 6. TOTAL VOLATILITY IN THE REGIMES OF BABY1 AND BABY2

		Total volatility
Baby1	1	14.63%
Baby2	Regime1	14.50%
	Regime 2	8.44%

As seen in Table 4 and Table 5, Baby 1 and Baby 2 seem to be in agreement on the prices of the vanilla options and yet in disagreement on the price of the call 100 up & out at 107. You may think the discrepancy between the barrier option prices is due to the fact that Baby2 has not exactly matched the vanilla smile generated by Baby1. Indeed, Baby2 is structurally different from Baby1 in that it can only pick up a *single* return jump, when it starts in Regime 1. This jump takes it to Regime 2, and it is only then that it may incur a jump of a different nature. Notice how Baby2 has managed to decipher Baby1's downward jump (it finds a jump of size −28% and intensity 0.14 to account for the jump of size −25% and intensity 0.20), and how it has fudged Baby1's 7% Brownian and upward jump into a Brownian component of 10.02%.

However, total volatility in Regime 1 of Baby2 is very close to total volatility[6] in Baby1 (see Table 6). As a result, Baby2 performs better at fitting the out-of-the-money put skew of Baby1 than the out-of-the-money call skew. Still, it may look surprising that the difference between the barrier

>

option prices produced by the two models should be so big, especially so when the prices of the calls of strike 100 and 107 are not that different.

5.1.2 Body examples

To clear any remaining doubt, we move to the next stage and consider a more evolved model. The underlying can now find itself in *three* different regimes of Brownian volatility. Transition between the regimes is still carried out by a Markovian matrix of six inter-regime Poisson jumps. The model now involves 15 free parameters (three Brownian diffusion coefficients, six jump sizes and six jump intensities). We call this new stochastic structure "Body."

Baby1 and Baby2 now appear as special cases of Body. Baby2 corresponds to Body with the transitions to Regime 3 disabled. And Baby1 corresponds to Body with the three diffusion coefficients set equal to 7% and the two Poisson jumps from any of the three regimes to any other set equal to Baby1's Poisson jumps.

We then propose the following. We shall calibrate Body twice to a full vanilla smile, each time with a different initial guess on the 15 process

TABLE 7: COMPARISON OF THE IMPLIED VOLATILITY SURFACES GENERATED BY BODY1 AND BODY2 WITH THE ONE INFERRED FROM VANILLA MARKET PRICES. THE SPOT PRICE IS 100.

	Strike										
Maturity(years)	**80**	**85**	**90**	**95**	**100**	**105**	**110**	**115**	**120**	**130**	**140**
	Market	19.00%	16.80%	13.30%	11.30%	10.20%	9.70%				
0.18	**Body1**	19.22%	16.38%	13.35%	11.69%	10.38%	10.29%				
	Body2	19.11%	17.14%	13.91%	10.93%	10.76%	10.00%				
	Market	17.70%	15.50%	13.80%	12.50%	10.90%	10.30%	10.00%	11.40%		
0.43	**Body1**	17.56%	15.85%	13.97%	12.43%	11.14%	10.08%	10.07%	11.53%		
	Body2	17.49%	15.89%	14.11%	12.22%	11.29%	10.35%	9.82%	10.30%		
	Market	17.20%	15.70%	14.40%	13.30%	11.80%	10.40%	10.00%	10.10%		
0.70	**Body1**	17.34%	15.90%	14.37%	13.00%	11.85%	10.87%	10.11%	10.20%		
	Body2	17.15%	15.86%	14.50%	12.96%	11.91%	10.95%	10.36%	10.37%		
	Market	17.10%	15.90%	14.90%	13.70%	12.70%	11.30%	10.60%	10.30%	10.00%	
0.94	**Body1**	17.22%	15.93%	14.60%	13.39%	12.36%	11.47%	10.69%	10.23%	11.04%	
	Body2	17.05%	15.94%	14.77%	13.42%	12.39%	11.44%	10.81%	10.64%	10.74	
	Market	17.10%	15.90%	15.00%	13.80%	12.80%	11.50%	10.70%	10.30%	9.90%	
1.00	**Body1**	17.19%	15.93%	14.65%	13.48%	12.46%	11.60%	10.83%	10.32%	10.86%	
	Body2	17.04%	15.96%	14.82%	13.52%	12.50%	11.55%	10.91%	10.71%	10.74	
	Market	16.90%	16.00%	15.10%	14.20%	13.30%	12.40%	11.90%	11.30%	10.70%	10.20%
1.50	**Body1**	16.99%	15.98%	14.97%	14.03%	13.19%	12.46%	11.80%	11.24%	10.56%	10.89%
	Body2	16.95%	16.08%	15.17%	14.13%	13.24%	12.38%	11.71%	11.34%	10.96%	10.96%
	Market	16.90%	16.10%	15.30%	14.50%	13.70%	13.00%	12.60%	11.90%	11.50%	11.10%
2.00	**Body1**	16.87%	16.03%	15.20%	14.42%	13.71%	13.07%	12.48%	11.98%	11.17%	10.76%
	Body2	16.86%	16.13%	15.38%	14.53%	13.78%	13.02%	12.38%	11.94%	11.35%	11.11%
	Market	16.80%	16.10%	15.50%	14.90%	14.30%	13.70%	13.30%	12.80%	12.40%	12.30%
3.00	**Body1**	16.74%	16.12%	15.52%	14.94%	14.40%	13.89%	13.42%	12.99%	12.26%	11.67%
	Body2	16.70%	16.16%	15.61%	15.02%	14.47%	13.90%	13.37%	12.93%	12.21	11.73%
	Market	16.80%	16.20%	15.70%	15.20%	14.80%	14.30%	13.90%	13.50%	13.00%	12.80%
4.00	**Body1**	16.68%	16.19%	15.72%	15.26%	14.83%	14.42%	14.03%	13.67%	13.03%	12.48%
	Body2	16.58%	16.15%	15.74%	15.29%	14.87%	14.44%	14.01%	13.64%	12.96%	12.41%
	Market	16.80%	16.40%	15.90%	15.40%	15.10%	14.80%	14.40%	14.00%	13.60%	13.20%
5.00	**Body1**	16.63%	16.24%	15.85%	15.48%	15.12%	14.78%	14.45%	14.14%	13.58%	13.09%
	Body2	16.49%	16.14%	15.81%	15.45%	15.12%	14.78%	14.44%	14.13%	13.53%	13.01%

parameters. And we shall pick a real vanilla smile this time (the one in Figure 1 that gave us the local volatility surface in the first section), not an artificially created one. Then we shall turn to the pricing of barrier options. The results of calibration are shown in Table 7 and the corresponding sets of parameters are shown in Tables 8 and 9.

Notice that two calibration instances, Body1 and Body2, match the given market vanilla smile fairly closely (see Table 7 and Figures 8, 9 and, 10). Also note that we manage to fit a *whole surface* of options prices, with different strikes and different tenors, with one set of constant parameters, when other smile models typically require that the parameters become functions of time.[7] True, the reason for that may be that our parameters are many (15) and our "Body" model not so parsimonious after all. This also explains why the calibration procedure may produce multiple solutions and the loss function admit of several local minima. As far as barrier options are concerned, we first look at the one-touches. In market practice, one-touches are identified and quoted relative to Black-scholes. The "30% one-touch" conventionally refers to the American

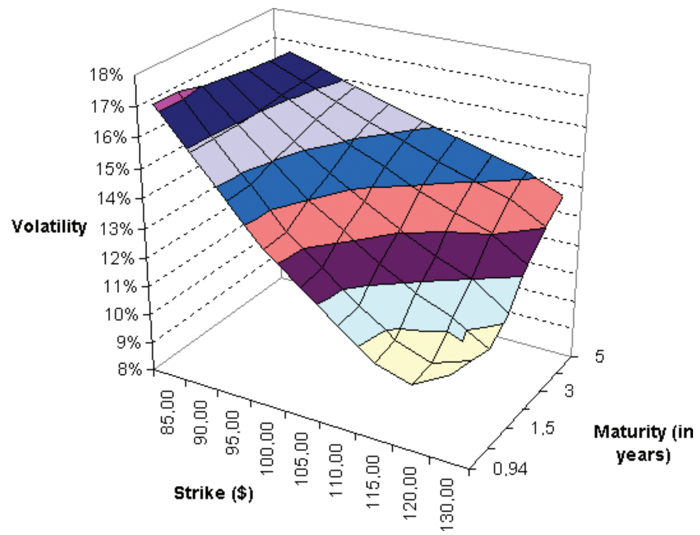

Figure 8: Body1 implied volatility surface

TABLE 8. BODY1 PARAMETERS

	Brownian diffusion	Total volatility
Regime 1	9.57%	11.67%
Regime 2	6.24%	32.23%
Regime 3	2.25%	11.88%
	Jump size	**Jump intensity**
Regime 1 → Regime 2	−9.07%	0.2370
Regime 2 → Regime 1	62.67%	0.0855
Regime 1 → Regime 3	2.72%	3.3951
Regime 3 → Regime 1	−3.17%	2.9777
Regime 2 → Regime 3	24.63%	1.0944
Regime 3 → Regime 2	−22.66%	0.2040

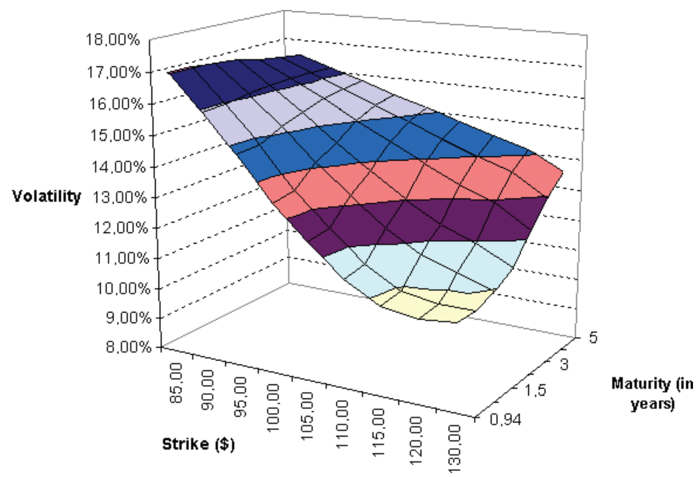

Figure 9: Body2 implied volatility surface

TABLE 9. BODY2 PARAMETERS

	Brownian diffusion	Total volatility
Regime 1	7.77%	11.63%
Regime 2	19.11%	25.08%
Regime 3	3.98%	7.45%
	Jump size	**Jump intensity**
Regime 1 → Regime 2	−9.02%	0.6254
Regime 2 → Regime 1	15.85%	0.5124
Regime 1 → Regime 3	5.24%	0.8750
Regime 3 → Regime 1	2.19%	0.7163
Regime 2 → Regime 3	17.17%	0.4589
Regime 3 → Regime 2	−11.20%	0,2891

digital option, paying out \$1 as soon as the barrier is hit from below, *that would be worth 30 cents in the Black-Scholes world*, when priced with the ATM implied volatility of corresponding maturity. ("−30% one-touch" conventionally means that the barrier is hit from above). A market quote of −4.88% for that one-touch means that it is actually worth (30% − 4.88%) = 25.12% in the present market, or smile, conditions.

Table 10 describes the one-touch price structures given by Body1 and Body2. The differences are considerable. As a result, standard barrier options will also be priced very differently by the two models (see Table

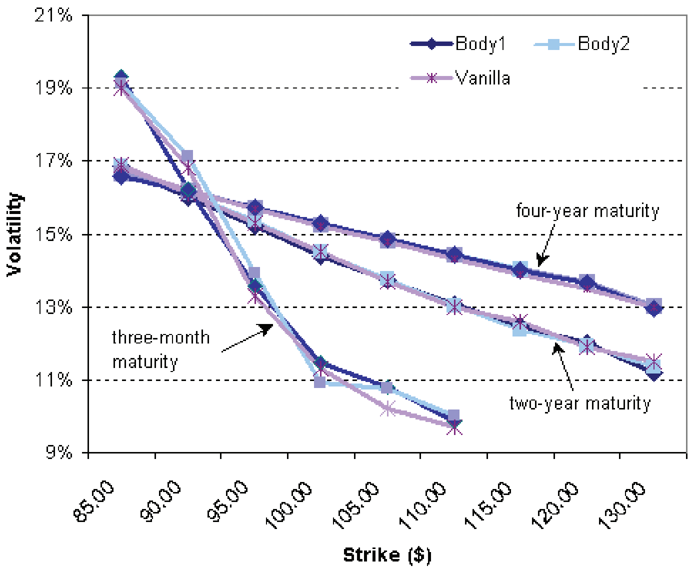

Figure 10: Cross-sections of the implied volatility surfaces shown in Figures 8 and 9 at three different maturities

11). Notice that it is the *same* model (Body) that is producing agreement on the vanillas and total disagreement on the barriers between two calibration instances. The situation is different from the case of agreement/disagreement between two *different* models, such as local volatility and stochastic volatility, or jump-diffusion. Those simpler models simply disagree with each other because of a big difference in what otherwise qualifies as *simple* stochastic structure. It is not even guaranteed that they can fit a complete vanilla smile surface. Their case is somewhat comparable to the agreement/disagreement we found between Baby1 and Baby2. When the stochastic structures become complex, however, and start combining stochastic volatility and correlated return jumps and volatility jumps (in models such as Body, or universal volatility, which seem to be imposed on us anyway by the natural course of events and by the evolution of the smile problem), we shall expect to witness increasingly frequent cases where a certain vanilla smile is perfectly matched, yet certain exotic options are very badly mispriced, or priced just by pure luck. In other words, we are way past the old debate on whether local volatility is better, or jump-diffusion is better, or stochastic volatility is better, on whether they agree or disagree on the exotics, and whether universal volatility should come and replace them all. Definitely universal volatility is the answer and Lipton's model has somewhat outgrown Lipton's article. As universal volatility models or SVJ models (stochastic volatility + jumps) seem unavoidable, the preoccupying issue today is how to avoid a dilemma, occurring within the *same* universal volatility model, such as embodied by Body1 and Body2.

You can easily imagine what the obvious trap would be. "How shall we distinguish between multiple local minima, such as Body1 and Body2, and pick the right one?" and you may be tempted to answer: "Let us pick the solution that fits the vanillas best, down to the last penny!" This is what a well-known analytics vendor seem to be proposing. Their way out of the dilemma is that a simulated annealing algorithm shall find *the global minimum of the loss function involving the vanillas only!* Has anyone worried where that would leave the exotics? We live in a very dangerous world indeed.

We know what the right proposal should be. *Include the one-touches, or other relevant exotic options, in the calibration procedure.* As a matter of fact, calibrating to the one-touches together with the vanillas transforms the ill-posed problem into a well-posed one. We will no longer try to reach for the global minimum among many local minima, but for a unique global minimum, full stop.

To illustrate that, we calibrate Body to the vanilla smile and to *the whole collection of one-touches produced by Body1* (Table 10), yet we select as initial guess of parameters *the solution produced by Body2* (Table 9}). This way we can see whether the one-touches will pull us out of what used to be the wrong local minimum. The calibration result is summarized in Table 12. We call it "Body1Double," and check it against Body1. Our minimization routine is a standard Newton method.

Notice the following interesting phenomenon. Within an acceptable numerical tolerance, Body1Double and Body1 seem to agree on the

TABLE 10. ONE-TOUCH PRICES INFERRED BY BODY1 AND BODY2

		One-Touches									
Maturity (year fraction)		**−5%**	**−10%**	**−20%**	**−30%**	**−50%**	**50%**	**30%**	**20%**	**10%**	**5%**
0.175	**Body1**	0.51%	−1.26%	−3.81%	−5.37%	−6.44%	−6.13%	−7.81%	−8,36%	−6.08%	−3.58%
	Body2	3.99%	0.51%	−5.80%	−10.45%	−14.78%	−7.72%	−7.01%	−5.91%	−4.18%	−2.66%
1.5	**Body1**	7.15%	6.23%	2.44%	−1.70%	−8.19%	−3.04%	−6.64%	−7.89%	−6.67%	−4.04%
	Body2	8.78%	8.94%	6.63%	3.08%	−4.88%	−3.62%	−7.76%	−8.16%	−5.98%	−3.55%
5	**Body1**	8.12%	8.74%	7.56%	5.17%	−0.87%	0.02%	−2.63%	−4.10%	−4.45%	−3.30%
	Body2	8.06%	9.12%	8.74%	7.10%	2.43%	−0.11%	−3.14%	−4.65%	−4.59%	−3.17%

TABLE 11: PRICING BY BODY1 AND BODY2 OF A PUT 103, KNOCKED OUT AT 95, WITH A 90-DAY MATURITY

	Price
Body1	0.99
Body2	1.29

TABLE 12: COMPARISON OF THE PARAMETERS AND TOTAL VOLATILITY NUMBERS OF BODY1DOUBLE AND BODY1

	Brownian Diffusion		Total volatility	
	Body1Double	**Body1**	**Body1Double**	**Body1**
Regime 1	9.55%	9.57%	11.69%	11.67%
Regime 2	6.44%	6.24%	32.23%	32.50%
Regime 3	2.41%	2.253%	11.88%	11.76%
	Jump size		**Jump intensity**	
	Body1Double	**Body1**	**Body1Double**	**Body1**
Regime 1 → Regime 2	−9.05%	−9.07%	0.2405	0.2370
Regime 2 → Regime 1	25.02%	62.67%	1.1279	0.0855
Regime 1 → Regime 3	2.79%	2.72%	3.3208	3.3951
Regime 3 → Regime 1	−3.07%	−3.17%	2.9882	2.9777
Regime 2 → Regime 3	65.12%	24.63%	0.0729	1.0944
Regime 3 → Regime 2	−22.68%	−22.66%	0.2025	0.2040

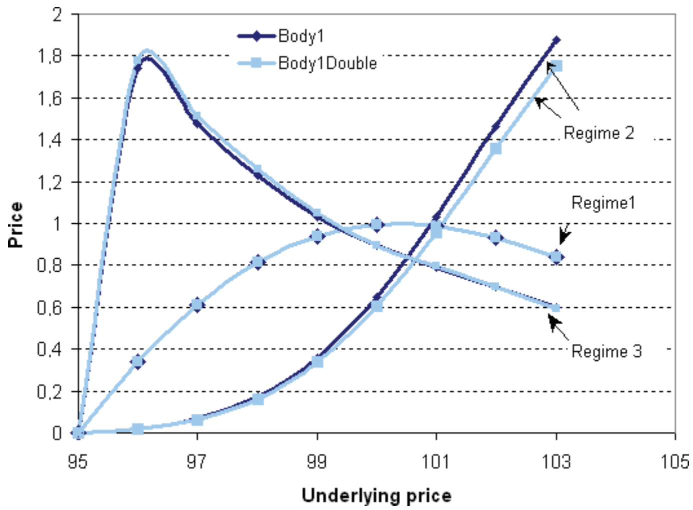

Figure 11: Price of the Put 103 down-and-out at 95 against the underlying price using Body1 and Body1Double parameters, in all three regimes

Brownian diffusion in all three regimes and on the Poisson jump sizes and intensities taking us from Regime 1 to Regime 2 and 3. They also agree on the Poisson jumps leading from Regime 3 to Regime 1 and 2. However, Body1Double and Body1 seem to have switched the Poisson jumps leading from Regime 2 to Regimes 1 and 3. The explanation is that total volatility is roughly the same in Regime 1 and Regime 3 (while it is much higher in Regime 2), and that the only things that the underlying can "see," once in Regime 2, are the *total volatility* of the Regime it will visit next and the Poisson jumps of course. While formally different, Body1 and Body1Double are in fact perfectly equivalent solutions (as when you permutate the regimes). As a matter of fact, we can check how well they agree on the pricing of the Put 103 knocked-out at 95, for different spot prices and different regimes (Figure 11).

5.1.3 Full Body, anybody, and nobody

You may wonder what is so special about the stochastic structure of Body. Nothing really, except that it has the minimum features that seem to be required to capture the phenomenology of smile and smile dynamics. As far as we are concerned, this is the only thing that counts. The question whether volatility should be diffusing rather than jumping in between discrete states, whether the Poisson jump distribution should be continuous rather than discrete, is in the last resort an aesthetic question (and often driven by the desire of analytical solutions). And there is just no way we could discriminate between the probability distributions of such models, by looking at the time series of the underlying. Volatility of volatility is hardly measurable. Not mentioning that every continuous model turns "discrete" when solved numerically.

To the aesthetically-minded, however, we can always suggest that Body can be further worked out into a full-bodied version that we call "Full Body." There is no limitation to the number of volatility regimes we may want to consider, so a continuum of regimes is in theory possible. And there is no limitation either to the number of Poisson jumps occurring between regimes or within regimes. As we shift between Regime 1 and Regime 2, it could be a random draw whether the concurrent return jump is positive or negative, and of what size. And Regime 1 could be characterized, not just by a Brownian diffusion, but also by a collection of Poisson jumps occurring within that regime. Body is very flexible and can mimic any given model. Body is really anybody's model. Or it can be everybody's model at the same time (for instance Regime 1 can harbour a full local volatility model, Regime 2 a full Heston model, Regime 3 a full Merton model, etc.). Yet Body will always be the dynamic, perfectly intertemporally consistent, version of such "mixings," by contrast to what has come to be known as the "mixture" or "ensemble" approach (Gatarek

(2003), Johnson, Lee (2003)). We should really be talking of "superpositions of models" in our case rather that "mixtures" (if we may borrow this crucial distinction from quantum mechanics), in order to distance ourselves from the unhappy "ensemble" approach.

Full Body is in fact a general structure, a family of models rather than a model. The way people are used to think about regimes is in temporal succession. A regime of "sticky strike" smile behaviour can follow a regime of "sticky delta," etc. In the limit, we propose that you wake up every day in a state of stochastic superposition of such regimes (yet, we repeat, with total inter-temporal consistency and homogeneity), and that you watch for the market prices (one-touches, forward starting options, etc.) that will best determine the superposition. This may sound as the end of modelling to some people: "Black-Scholes, Merton, Heston, SABR, Bates, sticky-strike, sticky-delta, etc., those are *models*, those are good names!" Indeed so. Our model deserves no name.

5.2 The hedging issue: Optimal hedging

Let us now explore the other side of the smile problem, which we said was intimately linked to the pricing of exotic options, namely the discrepancy that may occur between the hedging strategies of two different models despite their being calibrated to the same vanilla smile. Before we do so, however, we have to introduce a fundamental concept. In all the smile models we've been considering (jump-diffusion, stochastic volatility, universal volatility) markets are incomplete. In other words, contingent claims cannot be replicated with the underlying alone. Indeed the Black-Scholes argument of self-financing, perfect dynamic hedging breaks down in the presence of jumps and/or stochastic volatility. Local volatility smile models try desperately to save the complete market paradigm, but are unrealistic precisely for this reason. They imply, for instance, that a barrier option is perfectly hedgeable with the underlying, no matter the volatility smile.

The other models evade the hedging issue altogether. They lay the stochastic process of the underlying in the risk-neutral world directly, and assume that option value is the discounted expectation of payoff under the risk-neutral measure[8]. While this guarantees that their option prices do not create instant arbitrage opportunities, they offer no guarantee that the option value is "arbitraged" against the process of the underlying, in the Black-Scholes sense of "volatility arbitrage." In other words, you cannot hedge the option with the underlying, and "lock" the option value at the inception of the trade, through subsequent dynamic action on the underlying. All you are offered in terms of hedging is the partial derivative with respect to underlying—never a hedge in the presence of jumps—or some "external" bucketing of the volatility surface, which almost certainly contradicts the assumptions of the model.

What is needed is a theory of option pricing and hedging in incomplete markets. We will introduce the concept of "optimal dynamic hedging." By that we mean a self-financing dynamic portfolio, involving the underlying and the money account, which optimally replicates the derivative instrument, in some sense of "optimality." Our choice of criterion is the minimization of the variance of the P&L of the total portfolio. In other words, we draw on stochastic control theory to propose a self-financing dynamic hedging strategy for the derivative that lets you break-even on average and guarantees that the distribution of your P&L is the most "sharply peaked at zero" that can be. We *then* propose as a definition of "derivative instrument value" the initial cost of the self-financing optimal hedging strategy. And we find that the initial cost of the optimal self-financing replicating portfolio has the property of a pricing operator, therefore behaves like a risk-neutral probability (Henrotte (2002)).

Because our optimal hedging takes place in the real world, and our risk-neutral probability measure is associated with optimal hedging, we are able to link our risk-neutral probability with the real probability. Calibration and pricing can take place in the risk-neutral world. Since our process parameters are inferred from the market prices of options, it is as if we were reverse-engineering the pricing operator from those traded prices, and reapplying it to find the unknown prices of some other options. However, when we start worrying about hedging the option, this can only take place in the real world and necessitates the transformation of the probability measure. This transformation requires an independent input: the market price of risk of the underlying, or its Sharpe ratio.

We also define the variable **HERO** (*Hedging Error at Replicating Optimum*) as the minimized standard deviation of the hedged portfolio. HERO is the measure of market incompleteness with regard to the given instrument. It may be large either because the underlying is "incomplete" (large jumps, stochastic volatility...) or because the payoff is complex (exotics...). In the absence of jumps and stochastic volatility, our optimal hedge would indeed coincide with the Black-Scholes perfect hedge, and HERO would collapse to zero. Alternatively, the HERO of the underlying is trivially zero, no matter the stochastic process.

5.3 The "true" smile dynamics

Let us now go back to our solemn question: "Can we have control over the smile dynamics in homogeneous models?" At first blush, It seems the answer is no. Indeed, in space homogeneous models, Euler's theorem implies the following relation:

$$C = S(\partial C/\partial S) + K(\partial C/\partial K) \qquad (1)$$

where C is the vanilla option price, S the underlying price and K the option strike.

C, S, K and $\partial C/\partial K$ being fixed for a fixed smile surface, this implies $\partial C/\partial S$, or Δ, is fixed. So it seems that two homogeneous models will agree on the option delta when they are calibrated to the same smile, no matter their respective stochastic structures. The Merton model, the Heston model, the Bates model, the SABR model when $\beta = 1$, will all produce the same vanilla option delta. Only space inhomogeneous models (like local volatility or universal volatility which involve an explicit relation between the diffusion coefficient and the underlying), can yield a different delta, because of the corrective term they introduce (see Equation 7).

But we wonder. Is $\Delta = \partial C/\partial S$ the right measure of smile dynamics? The answer is clearly "yes" in the local volatility case where the underlying is the sole driving variable. However, in models involving another state variable, typically in stochastic volatility or universal volatility models, one cannot realistically move the underlying over an infinitesimal time interval and freeze the other variable. As volatility is correlated with the underlying, it is very likely that it moves too. Partial derivatives, such as $\partial C/\partial S$ and $\partial C/\partial \sigma$, capture the smile dynamics only partially. What we really need is the real time dynamics of the option price. In the local volatility case, we were able to apply the chain rule to get the real time delta. The question is, How can we apply the chain rule when volatility is an *indeterministic function* of the underlying, i.e. is correlated with it?

Before we try to answer what seems to be a challenging mathematical question, let us ask why do we need the information on smile dynamics in the first place. Obviously in order to determine the number of underlying shares that should be held against the derivative, or in other words, to hedge. Only in the local volatility model does the notion of hedge coincide with the *mathematical* derivative with respect to underlying. In incomplete market models, there is no mathematically ready, i.e. *non financial*, notion of hedge. We need to form the financial notion of hedge first (for instance optimal hedging in the sense of minimum variance), then work out the mathematics.

We claim that our "optimal hedge" is the substitute of the notion of smile dynamics in incomplete market models. As a matter of fact, the whole notion of "smile dynamics" appears to be muddled once the problem is set in the right frame. It is but a heritage of the local volatility model—the only place where it finds its meaning—and the whole comparison of smile behaviours between local volatility and stochastic volatility models appears to be ill-founded for that matter (you are not comparing apples to apples), if all that is meant is the partial derivative with respect to the underlying. So we might as well drop the whole notion of smile dynamics and get down to the hedge directly. What good is the notion of smile dynamics in jump-diffusion models anyway?

Recall that as the market is incomplete, we can only hedge optimally, and the HERO reflects how imperfect the hedge is. The optimal hedge that we produce already factors in the fact that the underlying may diffuse and jump, and that volatility may be stochastically varying, correlated with the underlying. In other words, it captures precisely the sense of "total derivative" that mathematics alone was unable to give us. What seemed to be a purely mathematical question (How do we generalize the chain rule when the functions are indeterministic?) receives a *financial* answer once the real purpose of the question is recognized (i.e. hedging).

However, if your only interest in smile dynamics is to predict the future shape of the smile surface, and not necessarily to hedge, then your question may admit of a *probabilistic* answer—and a probabilistic answer only—outside the one-factor framework. Conditionally on the underlying trading at some level S at some future date t, you may want to know what the *expected* value of the vanilla options may be at that time, or in other words, what the smile surface may be *expected* to look like. *Expectation* here means probabilistic averaging (either risk-neutral or real) over the possible states of the other state variables(s), conditionally on the underlying being in state S. You should bear in mind, though, that this *expected value of the option* is a different notion to its *future price*, as it is purely mathematical and unrelated to replication.

Therefore the big question really becomes: "Can two homogeneous models agree on the vanilla option prices, yet disagree on their *optimal hedging strategies*?" The answer is a resounding "yes," as will be seen from the same Body examples as before. Recall the two instances of our calibration of Body to a full vanilla smile which had resulted in two different local minima, and consequently, in two different one-touch price structures. We weren't sure at the time whether the two solutions implied *different* smile dynamics, as they agreed on the option delta by homogeneity and by Euler's theorem. That they agree on the option price and delta, yet disagree on the optimal hedge (and HERO) can now be made explicit (see Table 13).

Only when additional information is included in the calibration, that is to say, information constraining the conditional transition probabilities, will the models agree on the "smile dynamics." And this is now meant both in the sense that they will agree on the exotic option pricing and that they will agree on the (optimal) hedging strategy. "How do we gain control over the smile dynamics?" is therefore simply answered by controlling some exotic option price structures, typically the one-touches or forward starting options.

This is a general answer, not just specific to homogeneous models. Indeed, optimal hedging in incomplete markets is a general idea. It is just that the homogeneous models have helped us make our point more sharply, thanks to the "surprising" feature due to Euler's theorem and to

TABLE 13: BODY1 AND BODY2 OUTPUTS FOR A 107 CALL

Sharpe Ratio	**0.1**		**0.5**		**0.9**	
	Body1	**Body2**	**Body1**	**Body2**	**Body1**	**Body2**
Price	1.0131	1.0189	1.0132	1.0189	1.0132	1.0189
Hero	1.4429	1.2609	1.4429	1.2608	1.2811	1.1238
Optimal hedge	**0.2217**	**0.1543**	**0.2177**	**0.1543**	**0.2409**	**0.1803**
Delta	0.2894	0.2774	0.2895	0.2774	0.2894	0.2774
Gamma	0.0531	0.0540	0.0531	0.0540	0.0531	0.0540

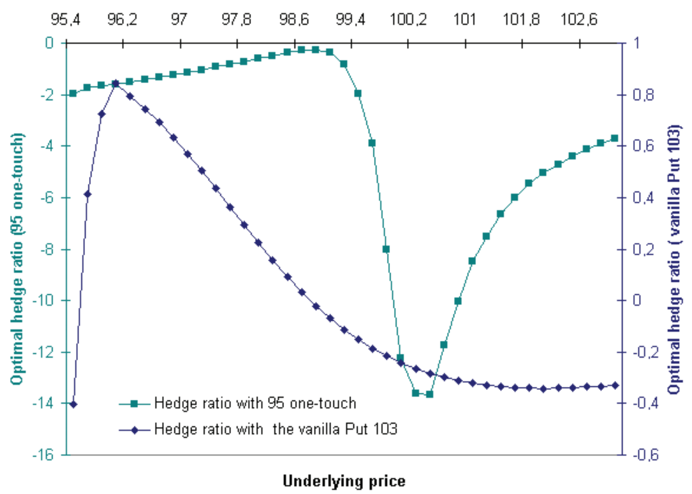

Figure 12: Optimal hedging ratios of the Put 103 KO 95 when either of the 95 one-touch or the vanilla Put 103 are used for dynamic hedging in combination with the underlying. The HERO (for $S = 100$) is 0.96 when no additional hedging instruments are used. It is 0.44 when the one-touch is used and 0.73 when the vanilla Put is used

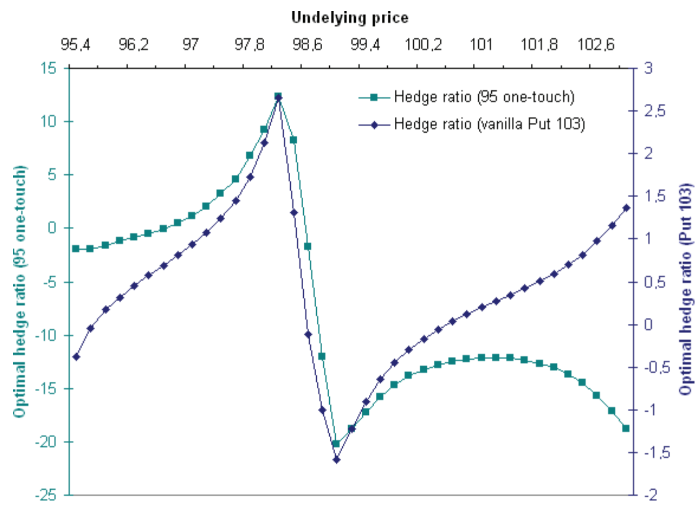

Figure 13: Optimal hedging ratios of the Put 103 KO 95 when both the 95 one-touch and the vanilla Put 103 are used for dynamic hedging in combination with the underlying. The HERO is now nearly zero over the whole range of spot prices

what seemed to be a loss of control over the option deltas. Also recall that Hagan and Blacher, who were arguing for control of the smile dynamics in inhomogeneous models, were not really taking into account what we have called the *true* smile dynamics.

In conclusion, there is no need to re-introduce inhomogeneity just for the sake of fitting a desired smile dynamics or a desired barrier option price structure. Henrotte's principle can thus be reiterated: **Any departure from homogeneity should be the cause of great concerns and should therefore be strongly motivated**.

We also find interesting that the answer to what seemed at first an "innocent" yet very relevant question ("How do I control the smile dynamics in my smile model?") should require *the theory of hedging and pricing in incomplete markets* as indispensable intermediary step. Financially relevant questions can only be answered by relevant *financial* theory. The need to go back to the "basics" is a very welcome conclusion, to say the least, at a time when quantitative finance seems to be wasting itself in sophisticated mathematical exercise, or even worse, in sophistical pseudo-models imported from foreign domains (e.g. the "mixture of models," or "ensemble," approach which cannot even afford an inter-temporal process, let alone a hedging rationale)[9].

6 Conclusion: Generalizing Black-Scholes

We have made the case for the necessity of introducing exotic options in the calibration phase of the smile model, and the necessity of thinking in incomplete markets. *Smile dynamics* is more important than *smiles* as pricing and hedging are essentially *dynamic* concepts, and incomplete markets are omnipresent as smiles are essentially a departure from Black-Scholes. As a matter of fact, the smile problem really *begins* with the question of the smile dynamics and the question of the hedging rationale[10]. These questions had remained hidden from us as long as we remained blind to the degree of model-dependence in the traditional models. Calibration to the exotics not only validates the right guess about the smile dynamics, but it allows us, thanks to an extension of the argument of optimal dynamic hedging in incomplete markets, to further *lock* the implied smile dynamics.

Indeed, stochastic control theory can be invoked again and our optimal dynamic, self-financing, hedging portfolios can be generalized to include other hedging instruments beside the underlying (see Figures 12 and 13). The price processes of the hedging instruments are independently available to us as the initial costs of their respective optimal hedging strategies involving the underlying alone. This guarantees that the price of the hedged derivative instrument can still be defined as the initial cost of the composite hedging portfolio, and be independent of the particular choice of hedging instruments other than the underlying. Dynamic multi-hedging of a derivative instrument allows the resulting HERO to be even smaller and the market to approach completeness.

Typically a barrier option will be dynamically hedged with a combination of the underlying, a vanilla option, and a one-touch. A convertible bond will be hedged with a combination of the underlying, an equity option and a credit default swap. A complex cliquet will be hedged with the underlying and a combination of simple forward starting options.

Calibration should be calibration with a point. It achieves nothing on its own. Treating the vanillas, the one-touches, the forward starting options, or the credit default swaps, as alternative liquid instruments underlying our jump-diffusion/stochastic volatility process, and using them in the dynamic hedging of the given derivative instrument the same way that the underlying stock is traditionally used in Black-Scholes, is the right way to generalize Black-Scholes to the case of smiles. Making sure that the smile model prices the "underlyings" in agreement with the market, and that it is calibrated to their dynamics, is in the end no different from saying that the Black-Scholes model prices *the* underlying in agreement with the market and is calibrated to its Brownian volatility.

When the hedging instruments are appropriately chosen, we expect the hedge ratios to be robust. Our hope is that they may even not depend on the particular model. In the end, a model is just a piece of machinery, "cogs and wheels" that allow us to dynamically glue together the appropriate derivative instruments. If the relevant dynamics is properly captured (in other words, if the model is calibrated to the maximum relevant information), and if the hedging instruments are properly chosen, then the hedging strategy should more or less impose itself naturally. As a matter of fact, we found that it very often corresponded to the trader's, model-independent, intuition.

Thus we conclude with the *disappearance of the model*. If solving the smile problem means finding *the right tool*, then the directions we have suggested are indeed the right directions to pursue. This goes hand in hand with a constant awareness of the perfectibility and relativity of the tool. What we have proposed in this paper is not so much the "definitive smile model" as it is the definitive way to think *critically* about any model.

But if solving the smile problem means finding the absolutely true process and the absolute pricing algorithm, then we can safely declare: "Nobody can solve the smile problem!"

FOOTNOTES & REFERENCES

1. We will later refer to the local volatility model(s) in the singular or the plural depending on whether we mean the theoretical principle or the particular numerical techniques.
2. As it was once argued in one thread of the Wilmott forums—Skew and forward volatilities(http://www.wilmott.com/messageview.cfm?catid=4&threadid=2551&FTVAR_MSG DBTABLE=).
3. See (Henrotte (2004))
4. Of course, one-touches and forward starting options will not, in general, determine the smile dynamics completely (as Peter Carr once objected in a private communication). Think how large the number of degrees of freedom would be in the matrix of conditionals, if the problem were left completely non parametric. Not mentioning the multiplication of that number by the number of spatial state variables. When we say the one-touches and the forward starting options help determine the smile dynamics, we mean it only relatively. Indeed, we, too, will have to depend on our particular choice of model for imposing the missing constraining structure. We need however to strike the right balance between the degree of structure imposed by the model and its ability to match the prices of contingent claims with very different payoff structures. Our solution is original both in the sense that it avoids the trap of non parametric inference and that it is more flexible than the traditional parametric models.
5. See Lipton, McGhee (2002)
6. Total volatility includes the Brownian volatility and the volatility due to jumps, it is expressed by $V_i^2 = v_i^2 + \sum_k \lambda_i^k \left(y_i^k\right)^2 + \sum_j \lambda_{i\to j} \left(y_{i\to j}\right)^2$ where i denotes the regime for which the total volatility is calculated, j denotes the regimes i the underlying can migrate to, k denotes the jumps occuring within regime i. The rest of the notation is self-explanatory.
7. e.g. Dynamic SABR.
8. Typically, Lipton (2002) writes : *"As always, we can evaluate the price of an option as the discounted expectation of its payout under a risk-neutral measure. We set aside many important issues related to the incompleteness of the market in the presence of jumps and stochastic volatility, and use the risk-neutralised dynamics* [. . .] *throughout."*
9. See Piterbarg (2003) for a sweeping criticism of the ensemble approach.
10. See Ayache (2004)

■ L. Andersen and R. Brotherton-Ratcliffe. The equity option volatility smile: an implicit finite-difference approach.*The Journal of Computational Finance*, 1(2), Winter, 1998.
■ L. Andersen and D. Buffum. Calibration and implementation of convertible bond models. October 2002.
■ M. Avellaneda, A. Carelli and F. Stella. A bayesian approach for constructing implied volatility surfaces through neural networks. *The Journal of Computational Finance*, 4(1):83—107, Fall 2000.
■ E. Ayache. A good smile is vital. *FOW*, May 2001.
■ E. Ayache The philosophy of quantitative finance. Forthcoming in Wilmott.
■ G. Bakshi and C. Cao. Risk-neutral jumps, kurtosis, and option pricing. *Working paper*, December 2002.
■ D. Bates. Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *Review of Financial Studies,* 9(1):69, 1996.
■ G. Blacher. A new approach for designing and calibrating stochastic volatility models for optimal delta-vega hedging of exotic options. *Conference presentation at Global Derivatives, Juan-les-Pins*.
■ F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, (81):637, 1973.
■ J. Bodurtha Jr. and M. Jermakyan. Nonparametric estimation of an implied volatility surface. *The Journal of Computational Finance*, 2(4):29—60, Summer 1999.
■ T. Coleman, Y. Li and A. Verma. Reconstructing the unknown local volatility function. *The Journal of Computational Finance*, 2(3):77—102, Spring 1999.
■ E. Derman. Regimes of Volatility. *RISK*, April 1999.
■ E. Derman. I. Kani. The volatility smile and its implied tree. *Quantitative Strategies Research, Goldman Sachs*, 1994.
■ B. Dupire. Pricing with a Smile. *Risk*, 7(1):18, 1994.
■ D. Gatarek. Libor market models with stochastic volatility. March 2003.
■ J. Gatheral. Stochastic volatility and local volatility. *Lecture Notes*, Fall Term 2003.
■ P. Hagan, D. Kumar, A. Lesniewski and D. Woodward. Managing smile risk. *Wilmott*, page 84, 2002.

■ P. Henrotte. Dynamic mean variance analysis. *Working paper, http://www.ito33.com/theory/henrotte_philippe-portfolios.pdf*, July 2002.
■ P. Henrotte. Pricing kernels and dynamic portfolios. Working paper, *http://www.ito33.com/theory/henrotte_philippe-kernels.pdf*, August 2002.
■ P. Henrotte. The Case for Time Homogeneity. *Wilmott*, January 2004.
■ S. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 1993.
■ J. Hull and A. White. An analysis of the bias in option pricing caused by a stochastic volatility. *Advances in Futures and Options Research*, 1988.
■ N. Jackson, E. Sueli and S. Howison. Computation of deterministic volatility surfaces. *The Journal of Computational Finance*, 2(2):5—32, Winter 1998.
■ S. Johnson and H. Lee. Capturing the smile. Risk, March 2003.
■ N. Kahale. An Arbitrage-free interpolation of volatilities. *Working paper*, Hiram Finance May 2003.
■ R. Lagnado and S. Osher. A technique for calibrating derivative security pricing models: numerical solution of an inverse problem. *The Journal of Computational Finance*, 1(1):13, Fall 1997.
■ Y. Li. A new algorithm for constructing implied binomial trees: does the implied model fit any volatility smile? *The Journal of Computational Finance*, 4(2):69—95, Winter 2001.
■ A. Lipton. The volatility smile problem. Risk, February 2002.
■ A. Lipton and W. McGhee. An efficient implementation of the universal volatility model. Risk, May 2002.
■ R. Merton. Options pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3(1), 1996.
■ V. Piterbarg. Mixture of models: A simple recipe for a ...hangover? *Working Paper, Bank of America*, July 2003.
■ M. Rubinstein. Implied binomial trees. *Journal of finance*, (49):771, 1994.
■ D. Tavella and W. Klopfer. Implying local volatilities. *Wilmott*, August 2001.