The Equity-to-Credit Problem

(Or the story of calibration, co-calibration and re-calibration).

Elie AYACHE*

February 6, 2004

1 This essay was initially intended as the elaboration of the presentation I gave at the Quantitative Finance Review in London, last November. The original title of the talk was: 'The Equity-to-Credit Problem, or How to optimally hedge your credit risk exposure with equity, equity options and credit default swaps.' As will be soon apparent from my line of arguing, I will indeed tackle, in this essay, what has become an urgent issue in those fields where credit volatility and equity volatility intermingle, typically the pricing and hedging of convertible bonds. Originally perceived as equity derivatives, the hybrid securities known as convertible bonds are steadily drifting towards the class of credit derivatives. Indeed the last stories of default have demonstrated the frailty of what used to be the rock-bottom value of the convertible bond and a bedrock notion in its literature and analysis, the "bond floor." It is probably more appropriate today to say that convertible bonds are derivative both on the equity state and the default / no default state of the issuer as it is no longer sufficient, for the purposes of quantitative analysis, to only specify the payoff of the convertible bond in case of conversion into the underlying share. Indeed, a state-of-the-art pricing model almost certainly requires, in addition, the specification of the payoff of the convertible bond in case of default. I refer the reader to [3, 4] where all these insights are quantitatively fleshed out.

Convertible bond pricing is perhaps the first derivative pricing problem to have raised the question of the explicit relation between the credit spread and the share price. The hybrid nature of the instrument is not the only reason, for the delta-neutral convertible bond volatility arbitrageurs have long been worrying about the adjustment of the equity delta implied by such a relation. I shall concern myself with this question, which epitomizes the equity-to-credit problem. Let it be noted, in passing, that traders and arbitrageurs of non convertible corporate debt, and even in some cases, of pure credit derivative instruments such as credit default swaps, are awakening to the notion of equity delta too. Hedging the credit exposure with the traded equity of the issuer is

*I am greatly indebted to Philippe Henrotte whose thought and brilliant mind lie at the heart of the argument developed in this essay. Philippe's continuing advice and tireless conversation have been my most reliable guides, not only through my present writing process but also in my most general thinking.

another name for the equity-to-credit problem, and it reaches far beyond the confines of convertible bonds.

2 Although a proper exposition of the equity-to-credit problem is supposed to follow the particular order of moving from the credit problem as such to the bearing of the equity process on it, I will follow the reverse order in this essay. Instead of asking what the equity can bring to the credit problem, I will ask what the credit risk, or in other words the likelihood of default of the issuer, can bring to an outstanding problem in the equity derivatives field, namely the equity volatility smile. Convertible bonds are a bad case for distinguishing between subtleties of orders of exposition like the ones I am pointing. Obviously the convertible bond quantitative analyst must concern himself *both* with credit spread term-structures and implied volatility smiles, as witnessed in [1]. Let it be noted however, that convertible bond pricing has complexities of its own that have recently engaged, and are still absorbing, the specialists. Beside the complexity inherent in the structure itself (the non linear interplay of the multiple embedded options: the option to convert the bond, to redeem it earlier than maturity, to put it back at a fixed price, etc.) and the growing popularity of exotic features (make-wholes premiums, contingent conversions, variable conversion ratios, etc.), it is the proper treatment of default risk that has been the greatest subject of concern recently. Again, I refer the reader to [3, 4]. Hardly will one want to worry about stochastic volatility and volatility smiles on top of all these problems! The analyst is usually content to specify a certain constant volatility parameter in his convertible bond pricing engine, inferring it roughly from the implied volatility of equity options of similar strike and maturity, and turns to what is the most pressing problem to his mind, that of calibrating a suitable default intensity term-structure, and in the most demanding cases, of worrying whether default intensity should not be made a function of the underlying equity as well.

It has been the rule, for all those reasons, that the equity volatility smile problem did not pose itself *per se* in convertible bond pricing. This is despite the fact that convertible bonds are highly exotic derivative instruments! With the issuer's option of early redemption acting as a knock-out and contingent conversion acting as a knock-in and their triggers being located far away from the conversion price, can we feel comfortable using a single implied volatility number in our pricing engine? Not mentioning that early puts, and likelihood of default, can also considerably shorten the effective life of the bond. Have we not learnt from the independent smile literature that exotic option pricing is irreducibly entangled with smiles and smile dynamics [5]? Yet convertible bond pricing models continue to lead their lives separately from smile models. A very tempting proposition could be to make use of local volatility in convertible bond models. State-of-the-art convertible bond pricing engines, such as those produced by my own company, are equipped to deal with local default intensity, or hazard rate, anyway; they rely on unconditionally stable finite difference schemes and adaptive stepping as a prerequisite for solving what may become a very though numerical problem when local hazard surfaces need to be calibrated to credit default swap market data, and convertible bond deltas and gammas computed off them; so why not simply overlay the numerical solver with a local volatility surface?

Tempting as the proposition may be – as a matter of fact it is starting to receive some attention [1] – I am resentful of the very way it suggests itself to us. I will not dwell on the well-rehearsed and very deep arguments that go against local volatility, precisely when exotic option pricing is of concern. The general philosophy of my company, in this regard, has been given expression in [5]. What is noticeable, and I think most dangerous, here, is the way the complexity – or shall I say the perversity? – compounds itself. Since the model already accommodates local hazard rate surfaces, so the argument goes, let us add local volatility surfaces. For how can we sustain the pressure exerted on the single volatility number in our pricing engine any longer? Can we reasonably become the specialists of local hazard rate surfaces and their calibration routines, can we reasonably propose a tool which has exhausted one side of the problem and allows the calibration of full non parametric hazard rate surfaces to full surfaces of credit spread data (as against maturity and equity level), and not set volatility free on the other side? What meaning can a single volatility number retain, and how can it sit tight as the only remaining hinge, in the midst of such a complex system? Shouldn't the two surfaces be considered simultaneously and the two problems solved hand in hand? Can we solve for local hazard rate and not jointly solve for local volatility? Can we calibrate to surfaces of credit default swap spread data and not jointly calibrate to implied volatility surfaces?

3 This is our problem, precisely. This is precisely the equity-to-credit problem. Since default (or its probability) and the subsequent drop of the underlying share are among the greatest creators of implied volatility skew and term structure, there is no way we could calibrate the default process without worrying about the option data. The equity-to-credit problem is essentially a smile problem. As a matter of fact, it is even worse than a smile problem. It speaks of a more complex underlying process (jump-diffusion) and compounds two ill-posed inverse problems: calibration of volatility instruments and calibration of credit instruments. (There already exists a thick literature on the subject of smoothing, interpolating and extrapolating an arbitrage-free local volatility surface, so try to imagine the result of adding the worry about a hazard rate surface!)

Again, notice how the historical build-up of the problem and how the historical succession of its proposed solutions (among which those proposed by my own company) have got us into trouble. Because the convertible bond pricing problem is such a complex problem to begin with (hybrid nature, embedded options, endless exotic features), trying to formulate it outside a strict diffusion framework is the last thing we want to do! It is not as though the pricing of the convertible bond had become difficult, or perhaps even impossible, on account of the underlying diffusion process and this gave us sufficient reasons to reconsider that process and contemplate jump-diffusion or stochastic volatility instead! The pricing of convertible bonds has become difficult – and in some cases of outdated software, even impossible – for reasons having strictly to do with the convertible, not the process. Only in simpler cases, where simpler exotic structures are liquidly traded, for instance FX barrier options, has it prominently appeared that the smile problem should be treated as such and solved as such. We have so many other things to worry about in convertible bond pricing, before we get down to the smile problem! Enough for now to have taken that first step outside diffusion, which consisted in adding a hazard rate and a jump to zero of the underlying share! (How many people, by the way, recognize the fact that this "simple" jump-diffusion process is already posing a full smile problem, and that the implied volatility number they are importing into their convertible bond pricing system from the prices of equivalent equity options, no longer means the same?).

Against this conservative-sounding kind of argument, I say this is precisely the time when, on the contrary, a break with history should occur. Who said we should ever consider local hazard rates in the first place? Because of the wrong sense of complexity that the convertible bond may convey, the convertible bond problem is probably a wrong place to pose the equity-to-credit problem afresh, and start looking for radical alternatives for solving it. Like we said, convertible bond pricing specialists have been distracted by the wrong kind of difficulty – and I sometimes fear we might be among them – and for this reason perhaps, lack the freshness of the eye. This is why I have set out to pose the equity-to-credit problem as an equity smile problem rather than a hybrid problem, and consider my task today to be the continuation of the work on smiles and smile dynamics pioneered in [5] as well as to find out whether default risk, and its traded instruments such as credit default swaps, cannot help us frame the smile problem better, by any chance. Only when the equity-to-credit problem is addressed and solved as a smile problem proper, with the same critical spirit and independence of vision as we have exercised in [5], will we go back to convertible bonds and realize what we should have known from the start, namely, that their complexity ought to be the driving motive for wanting to base their valuation on a sound theoretical ground – be it at a big intellectual and cultural cost – not an excuse for evading it.

4 Anyone familiar with our general philosophy of derivative pricing should be guessing my point at this juncture. If one thing is really dear to my heart when it comes to the smile problem, it is the defence of homogeneous models. The case for homogeneity has been extensively argued in [5]. There we showed, among other things, that the pricing of exotics and the hedging of both vanillas and exotics are crucially dependent on the smile dynamics, and that the need to discriminate between the empirical smile dynamics (with sticky-strike and sticky-delta sitting at the two extremes) need not be answered by inhomogeneous models. We argued, on the contrary, that a sufficiently comprehensive homogeneous model – a model we christened 'Nobody' – can perfectly address the exotic pricing and hedging issues under all kinds of smile dynamics, provided a) the hedging problem is aptly formulated in the incomplete markets framework and b) the model is calibrated to those *quoted* market prices which are the empirical reflection of the projected smile dynamics, the prices of the one-touches and the forward starting options.

This is how we dispensed with the local volatility model and all its cross-breeds (universal volatility, etc.). In yet another paper [7], we argued that even in fields where inhomogeneous models are so deeply entrenched as to go unnoticed, for instance the modeling of yield curves and credit spread curves, homogeneous models can have right of way. Term structure is the direct consequence of stochastic character, after all, not of a dissymmetry inherent in time, and a parsimonious time-homogeneous stochastic interest rate model, or stochastic hazard rate model, or stochastic volatility model, can demonstrably reproduce any shape of interest rate, or credit spread, or volatility term-structure.

Insisting that the equity-to-credit problem shall be posed and solved as a smile problem will therefore strike our reader as an overt directive against inhomogeneous models. Unpacking the claim, this would mean that volatility and hazard rate have to be modelled as two independent, time and space homogeneous, stochastic processes, instead of being deterministic functions of time and the underlying. Smile dynamics and credit dynamics would then be accounted for by a suitable correlation with the equity process.

This sounds as a daring claim indeed, all the more so when structural models of the firm (Merton, KMV, CreditGrades) seem to have entrenched the view that the triggering of default is a deterministic function of the underlying equity price. The myth of the bankruptcy threshold has transformed the very liberal notion of probability of default into the very concrete vision of a "distance to default," and it is commonplace nowadays to speak of credit spreads that explode to infinity with a falling share, and of bond floors that strictly collapse to zero. As if somebody had ever managed to measure such infinite spreads, or truly held bonds whose value vanished strictly prior to default and not directly after!

Rather, this picture strikes me as the result of the confusion of possibility and actuality, and it is unfortunately imposed on us by the particular mathematical representation. Take, for instance, the so-called "reduced-form" models which are supposed to disconnect the structural link between default and the equity price. The underlying equity, however, remains the state variable governing the probability of default and this particular choice naturally leads us to sampling equity values which are very close to zero (if only because the PDE numerical schemes require suitable boundary conditions). And now the additional twist brought by default is that a vanishing equity price – which may be vanishing only formally, just for the sake of writing the boundary condition – cannot not *get mixed up* with the question of default. This causal confusion stems from the fact that it all looks as if the equity price were the *sole* state variable, therefore the cause of everything. As a matter of fact, since default is supposed to be triggered independently by a Poisson process, a second state variable should be recognized here: the Boolean admitting of "default" or "no default" as values.

Now of course a natural and most comprehensible recommendation in such a setting is to suggest that the intensity of that process (or in another words the probability of its triggering) may indeed increase with a falling equity price (see Appendix). But this need not imply in any way that the probability of default should become equal to one exactly, or the intensity reach infinity, at the limit where the share price is equal to zero. There is no reason why a phenomenon (default), which normally originates from an independent cause (or a complex set of causes summarized by the Poisson process), should know of no other cause, at the limit, but the share price. The problem with equity-to-credit, when forced to fit inside the inhomogeneous, deterministic representation of the relation, is that it forces all kinds of boundary conditions on us, not just the mathematical, formal one but also a material, causal one. Just because the intensity of the default process is driven by the equity, we are forced to assign a value to the probability of default for all equity prices, particularly when the equity is worth zero, and our state of confusion about the true causal origin of the phenomenon (whether it is default which causes the zero share price or the zero share price which causes default) leaves us no choice but to assign infinity at the boundary and hope to have settled the issue with that. Any number other than infinity would indeed seem unsatisfactory, or call for an even bigger problem. Why should the credit spread be worth 20% rather than 30% at the limit of a worthless share? We certainly do not wish to get involved into this extreme sort of corporate finance and, like a game theorist friend of mine once said, "Infinity, in some cases, is the best approximation of an otherwise arbitrary finite number!"

5 Think, by contrast, of a situation where the default intensity is an independent stochastic process. Think, for instance, of a simple two-factor model such as Nobody, where the second factor is taken care of by a discrete number of "regimes" (depending on the particular problem, this may be volatility, or the hazard rate – see [7]). Each regime is characterized by a given, constant, hazard rate, while the share price can formally vary from zero to infinity in each. Regimes can be interpreted as different ratings of the issuing firm and transitions between regimes are governed by a probability transition matrix. We must consider correlation between the share price process and the regime transitions as it is indeed very unlikely that the share should start falling dramatically and the firm not switch over to a regime of higher default intensity, or indeed to the default regime! (I forgot to say that default receives an interpretation, in such a framework, which is of a piece with the rest. In our discrete regime approach, default is a regime like any other, only a very special one where the hazard rate, or the *probability* of default, is no longer defined. See Appendix.)

Whatever the realistic interpretation of the model may be, the point is that the model does not impose on us restrictions, or philosophical boundary conditions, such as imposed by the inhomogeneous model above. When the firm is alive, it subsists in one of our discrete regimes, with a given finite hazard rate. As the regime representation is only a discretization of the credit state space, the general objection that could be levelled against us, according to which the firm may, as a matter fact, not fall exactly in any of our discrete regimes but somewhere in between, is answered by the fact that the regimes are probabilistically connected and it is only the *probabilistic averages* (i.e. prices) that matter. So long as default is recognized as the extreme regime in our discrete collection, what specific hazard rate numbers, or credit spreads, the other regimes get, no longer matters. And that each individual regime may formally allow, as its boundary condition, that the share may approach zero and the hazard rate remain finite, even constant, will not matter either.

The material process and the formal process are separated in our model. Although there formally exist states of the world where the share is close to zero and the credit spread no larger than a given finite number, these scenarios will materially get a very low probability because of the correlation between the share and the regime shifts. Moreover, correlation can itself depend on the regime and increase as the firm shifts to regimes of higher default intensity. In other words, to answer the question: "What happens when the share collapses?" all you have to do is follow its *material* price process as it really unfolds in real time, and observe that the firm will jump between regimes, and eventually default. Yet you ask: "Will the hazard rate ever reach infinity?" The answer is: "Not in any of the live regimes, no, and not in the default regime either, for then it is too late." Yet again: "Can we not imagine that the firm is actually climbing to regimes of explosively high default intensity as the share approaches zero, only those regimes are invisible to our discrete representation?"

Why not indeed, but then it is precisely the advantage of indeterminism over determinism (of correlation over deterministic functions) not to force on us a determinate answer to that question. It is precisely the advantage of having separated the hazard rate and the equity price into two distinct state variables, and of having distinguished between formal cause and material cause, to allow us to solve (or shall we say, dissolve) both problems simultaneously. "Why have 20% credit spread instead of 30% when the share price is equal to zero?" and the answer is: "Let us have both. As a matter of fact, let us have other values as well, and the question of the boundary condition will not matter anymore because it is only a formal device, required in each regime." "What is the real boundary condition then?" and the answer is: "It disappeared in the probabilistic "interspace." While the inhomogeneous model cannot avoid mixing the bottom of the share price spectrum with the certainty of default, all that our homogeneous model offers, by contrast, is a collection of regimes where the bottom of the share price is in each case immaterial to default, and a *separate* default regime[1]."

6 Taking the suggestion that default is a regime like any other more seriously, we now realize that it has been with us all along and implicit in everything we have ever said about default, even in the traditional inhomogeneous framework. Surely enough, the hazard rate was a function of the share price in that framework and default became certain as soon as the share hit its lowest boundary, but the jump into default itself and its consequence for both the derivative instrument and the underlying share were no different from what we are contemplating today for the homogeneous case. As a matter of fact, we have already suggested in the previous convertible bond article [3]

[1]It won't matter for the prices (which are probabilistic averages) whether the model is inhomogeneous or homogenous. But wouldn't it matter for the hedging? More specifically, would we be hedging with a heavier delta in the homogeneous model, as the share goes to zero, the same way as in the inhomogeneous model (see Table 2)? The answer is yes, and it revolves entirely around the question of incomplete markets and optimal hedging, as in [5]. See Appendix.

that "a softer appellation of the state of default could be "distress regime"" and that "it would certainly make sense to imagine a continuation of life after default." What makes life end at default, after all, is just the assumption that the share drops to zero and the derivative instrument to its recovery value and that the game is over. But what if the share did not drop to zero but to some recovery value as well, then resumed its trading? What if the event of default triggered a restructuring of the issuing firm and of its outstanding debt, and the holders of the bonds were offered the opportunity to hold on their assets and stand by their positions until further notice? What if the holder of the convertible bond found it more optimal to postpone until later his right to convert into the underlying share – provided he still owned a convertible bond after the restructuring – rather than opt out of the game at the moment of default and take away the recovery value of the bond?

Theoretical and impractical as these questions may sound, they have at least the merit of making us think how to properly extend our framework if need be. The suggestion in [3] was that we would end up solving a pair of coupled PDEs, one describing the pre-default regime and the other the default, or distress, regime. The regimes are coupled through the Poisson jump to default, and the transition is supposed to be irreversible in the sense that the firm cannot recover from the distress regime, back to the normal regime. As the share starts its journey, in the distress regime, with the recovery value it has hit after default, it was even suggested that its subsequent volatility might be different from the volatility in the normal, pre-default regime.

To summarize: If the assumption should ever be made that the share might not drop to zero upon default but to some recovery value and then resume its trading life, and if we should ever consider holding on our instrument for reasons such as restructuring or rescheduling, then the proper way to value the instrument *as of today*, would be to solve a system of two coupled PDEs, possibly with different diffusion coefficients, perhaps even with different payout conditions written on the instrument (for instance, the conversion ratio may change after restructuring, or coupons may no longer be paid, etc.).

7 As we stand now, we are just one step away from the full, homogeneous, equity-to-credit model that we've been hinting at. If the thread of default is capable alone of leading us to a regime-switching model, even in the classical inhomogeneous case, what is to keep us from spreading the idea over to the no-default side of the picture? If volatility is allowed to be different in the default regime, why wouldn't it be different in each of the pre-default regimes which corresponded, in our homogeneous model, to different ratings of the issuing firm? In other words, the suggestion here is this: Let *both* a stochastic volatility process and a stochastic hazard rate process be taken care of by the regime representation. The regime is not a traditional state variable, after all, in the sense that volatility is one such, and the hazard rate is another, in the traditional stochastic volatility or stochastic hazard rate models that people usually have in mind. A regime can be identified, not just by one hazard rate number λ as proposed in [7] or one volatility number σ as proposed in [5], but by the mathematical pair (σ, λ). (See

Appendix.) As a matter of fact it can generally be identified by the n-tuple (σ, λ, r, s,...) where r,s,... might be other parameters of the pricing equation that we want to turn stochastic, such as the short term interest rate, etc.

We are touching here on an interesting, and I think, very deep, idea. This is the idea that a pricing problem might be multi-factor yet we are able to handle it with the same unchanged regime representation. I am being very cautious here in picking the right words. Notice that I did not speak of a "multi-dimensional" pricing problem. People are used to think that anything stochastic has somehow to be diffusing, therefore to be mathematically represented by a full, continuous, spatial dimension. A three-factor pricing problem would mean solving a three-space-dimensional PDE, etc. Whereas I contend that a regime-switching model with a few regimes, say three or four, where each regime is characterized by a different triplet (σ, λ, r) and the underlying share diffuses in each regime, can in theory handle the pricing of, say, convertible bonds under stochastic volatility, hazard rate and interest rate. And this is achieved at no extra computational cost other than solving a system of three or four coupled one-dimensional PDEs. You can see now why I am at a loss for words, trying to frame the *nature* of the regimes. Surely enough, the regime is our extra state variable, but I must refrain from calling it a "dimension" as it is able to embed *multiple* dimensions, or rather, multiple factors, and I hesitate to call the individual factors, σ, λ, r "dimensions," as they do not represent separate state variables that live independently in their own individual, continuous spaces. A multidimensional solid need not be the tensor product of the individual dimensions, after all. I am not even sure we can put a name on the regime. A regime is a state variable, surely enough. But what is it a state of? Volatility, hazard rate, interest rate? Or is it an abstract entity with no name, a sort of container which may contain the name of volatility, of hazard rate and interest rate – or any other collections of names depending on the particular pricing problem – and will assign to these names the particular numbers that they get depending on the particular calibration to market? It seems we have found one more reason why our model should be called "Nobody," the model with no name.

It remains to deal with the general objection that the market may, as a matter of fact, not fall in any of the n-tuple regimes we are considering, all the more certainly so that our distinguished regimes now require that volatility should be the particular number that the particular regime says it is and, *simultaneously*, that the hazard rate should be the particular number, *and* the short rate the particular number, that the particular regime says they are. Even worse, the real pair-wise correlation between these three factors may be such that no transition between any pair of the n-tuple regimes can reflect it. As before, we reply that the real volatility, the real hazard rate, the real interest rate, and the real correlation are not observable. All that really matters are the observable prices of traded instruments. And these prices are probabilistic averages over the regimes. It is up to us to calibrate the regime-switching model to the market prices of the volatility, credit risk, interest rate, and correlation instruments that we think are most representative. Maybe we should increase the number of regimes – always an open question that the relative ease of calibration can answer alone.

As the values that the parameters get inside the regimes and the intensities of transitions between the regimes are determined by calibration only, the hope is that the calibration procedure will settle on a certain solution – consequently our model on a certain specification – which will serve no other purpose, in the end, but to price other instruments relative to the initial ones. The general philosophical idea being here that, just as we could not put a *name* on the regime, even less so will we be able to *read* in the regime some value of volatility, or hazard rate, or interest rate, that we believe shall obtain in reality. This is another way of saying that the philosophical doctrine known as instrumentalism is perfectly acceptable as an alternative to metaphysical, or even semantic, realism.

8 Let me turn now to what is probably the deepest, and by far the most original, insight about the regime-switching representation. (Notice that I am no longer calling it "regime-switching model" as it is, I think, more general than a model.) The argument is rather subtle so please bear with me. Earlier I said that a regime could be characterized by an n-tuple rather than the single value of a single parameter, and I gave the combination of the hazard rate, volatility, and the interest rate, as an example. The underlying share followed a diffusion process in each one of the regimes anyway, so let us not worry about the underlying share for the moment. (As a matter of fact, I can refer to [5] where it is suggested that the process of the share *inside* a regime might as well be a full jump-diffusion, or that the share might undergo one, or indeed several, jumps in value, as it switched over between regimes. But let us leave it at that.) Now suppose, for a moment, that the stochastic processes that we wish to worry about are not exactly the volatility process and the process of some other financial variable, such as the hazard rate. Suppose volatility is stochastic alright and the coefficients of its *own* stochastic process are stochastic too. Since writing is such an endless process, all I am proposing here is to take it one step further. Just as a diffusive, mean-reverting process was once written for the volatility of the underlying Brownian diffusion, and gave us the Heston model or the Hull and White model, I propose today to write a further process for the *volatility of volatility* – why not another diffusion? – or indeed for the mean-reversion coefficient of volatility, or for the correlation coefficient between the underlying share and volatility, etc. More specifically, let our tentative model be:

$$\begin{aligned}
&dS = rSdt + \sqrt{v}SdZ_1 \\
&dv = \kappa\left(\theta - v\right)dt + \varepsilon\sqrt{v}dZ_2 \\
&d\varepsilon^2 = \xi dt + \varphi dZ_3 \\
&\ldots \\
&\ldots
\end{aligned} \tag{1}$$

In theory, the writing process should never stop, for this is the essence of trading and the result of submitting the theoretical model to the market [6]. Just as the trading of the vanilla options turned the coefficient of the Black-Scholes formula, implied volatility, into a stochastic variable and created the need for stochastic volatility models such as Heston (or volatility smile models more generally), all we are noting here is that daily calibration of the given smile model, as well as trading the derivative instruments which are higher up in the hierarchy and specifically sensitive to the smile (for instance, the barrier options), will in turn create the need for a higher-level model such as the one we are proposing. Just as implied volatility once became a widely talked-about and a liquidly traded commodity, we are now talking of the next evolutionary stage where smiles become a commodity and are traded in turn. (Necessity of re-calibration of the smile models and its counterpart, this open-endedness of the trading / writing process, are perhaps the greatest challenges facing any theory of smiles today. In my sense, they properly belong to the *metatheory* of smiles, or in other words, the philosophy of derivative pricing.)

Taking my cue from what I said earlier, namely that a three-factor pricing model need not be confused with a three-dimensional pricing problem, what I propose next is to apply the regime-switching idea to the third process. Instead of assuming a diffusion process for the volatility of volatility ε, why not consider a three-regime representation, $[\varepsilon_1, \varepsilon_2, \varepsilon_3]$, and appropriate transitions between the regimes? $[\varepsilon_1, \varepsilon_2, \varepsilon_3]$ will be our regime-switching model of stochastic volatility of volatility, and each individual regime ε_i will act as a super-container containing a full Heston model with a different volatility of volatility parameter. And now the regime-switching idea can be further invoked and the question asked how the Heston model inside the super-container can be itself replaced with a regime-switching model of stochastic volatility. Most probably the answer will be: put containers inside the super-container, with sub-regimes of volatility, $[\sigma_{i1}, \sigma_{i2}, \sigma_{i3}]$, now occurring inside each super-regime ε_i.

The problem is, ε was meant initially as the diffusion coefficient of the *Heston volatility process*. As soon as the Heston model is replaced by a *regime-switching model* inside the super-containers, ε becomes meaningless. The stuff inside the container blows up the super-container. To keep out of trouble, we should turn the problem on its head. We should really start with the Heston model, replace it with a regime-switching model of stochastic volatility, then find an appropriate meaning for the super-container supposed to make the latter model stochastic. As apparent from [5], the regime-switching model of stochastic volatility $[\sigma_{i1}, \sigma_{i2}, \sigma_{i3}]$ is characterized by a bunch of parameters, (the individual values σ_{ij}, the intensities of transitions between sub-regimes j and the sizes of the simultaneous jumps of the underlying), and not just three like Heston. There is no diffusion coefficient, or mean-reversion coefficient, or long volatility coefficient in the regime-switching model of stochastic volatility, at least not explicitly, but a conjunction of regime and regime-switching parameters which can only act *together* to reproduce any of these features. So the way to make our regime-switching model of stochastic volatility become stochastic in its turn is to assign different conjunctions of parameters to each super-container and not just a different single-valued parameter ε_i. Each one of the super-regimes of the model of stochastic volatility of volatility becomes, so to speak, a full regime-switching model of stochastic volatility.

Since transitions between super-regimes are modelled the same way as the transitions between sub-regimes, i.e. through Poisson processes of given intensity (possibly with simultaneous jump in the underlying), and given the associative character of the operation of grouping the sub-regimes into super-regimes, *our regime-switching super-model of stochastic volatility of volatility is in the end indistinguishable from a regime-switching model of stochastic volatility with a large number of regimes.* To fix the ideas, if we are talking about three super-regimes supposed to represent the stochastic character of the volatility of volatility and three sub-regimes, occurring inside each super-regime, supposed to represent the stochastic character of volatility, then the resulting construction will be indistinguishable from a regime-switching model of stochastic volatility, with nine regimes[2].

9 And now we are ready for the last step of the argument. Remember that calibration to the market prices of derivative instruments is all that matters in our derivative pricing philosophy. It is the only key to unlocking the relative value of other instruments (and locking their hedging strategies [5]). There is nothing to tell us whether our model should have three, four or nine different volatility regimes other than the number and the variety of the instruments we are calibrating against, and our satisfaction with the calibration results. Given the large number of parameters implied by a nine-regime-switching model, chances are that a more parsimonious regime-switching model will fit the market prices just as well, and perhaps even exhibit more robust and more stable calibration behaviour. Perhaps a three-regime-switching model will do the job!

The most extreme form of the thought here is this: our volatility regime-switching representation, Nobody, is not just a model of stochastic volatility; *it is also a model of stochastic volatility of volatility, and a model of stochastic volatility of volatility of volatility, etc.* It contains at once the whole endless model-writing chain. Or rather, it is open like the writing process is open. So the question becomes: What could ever determine the particular hierarchical level at which the particular instance of Nobody shall land? Like I said, the answer lies in the particular nature and the particular prices of the derivative instruments we are calibrating against in the particular instance. Imagine that the option prices are not given by the market but artificially produced by a "first-level" stochastic volatility model such as Heston. Then our regime-switching model will match the corresponding vanilla smile and, for all practical purposes, mimic the behaviour of a first-level stochastic volatility model. Now suppose the instruments of concern are not plain vanilla but exotic structures, whose value, we know, is dependent on a level of complexity higher up than the given static smile, e.g. smile dynamics. For instance, we know from [5] that first-level stochastic volatility models such as Heston may not simultaneously match the prices of the vanillas and the cliquets, for they imply a certain smile dynamics which may not accord with the cliquet prices. Higher-level models are called for in this case, with volatility dynamics more complex than Heston,

[2] I must qualify the latter statement a little bit. If the individual values of volatility characterizing the different sub-regimes inside a given super-regime i, $[\sigma_{i1}, \sigma_{i2}, \sigma_{i3}]$, are equal respectively to those occurring in a different super-regime k, $[\sigma_{k1}, \sigma_{k2}, \sigma_{k3}]$, then the result will be indistinguishable from a three-regime-switching model of stochastic volatility, only the transitions between the regimes will occur in many more different ways.

for instance "universal volatility models" or indeed models even more general, such as Nobody. I refer the reader to [5] for the details. Once Nobody is calibrated to the vanillas and the cliquets, it will behave like a higher-level stochastic volatility model. Finally imagine a situation where the instruments are very complex structures with no definite sense of the particular hierarchical level at which their sensitivity stops, and that their prices are just given by the market. Then hopefully Nobody will match those prices, and only the market "will know," in that case, at which level we landed.

Yet you complain: "Surely there must be something out there to help us distinguish between models of stochastic volatility and models of stochastic volatility of volatility (or at least, between models of significantly different hierarchical levels). The probability distributions must come out different, and there surely must come a stage where Nobody is ruled out *a priori*. There must be some probability distribution that gets generated at some level yet cannot be reached by the regime-switching representation, no matter the number of regimes or the value of the parameters. For how could a two-factor model, such as the particular instance of Nobody that we are considering (where volatility is the only identifier marking the regimes), ever reproduce the richness of a multi-factor model such as stochastic volatility of volatility of volatility...?"

My reply to you is that you first try to realize what you are saying. Your perplexity relates back to the confusion we have already mentioned, between number of factors and number of dimensions. Recall the nine-regime-switching stochastic volatility model $[\sigma_{ij}]$ that we had obtained after unpacking the sub-regimes j of stochastic volatility occurring inside the super-regimes i of stochastic volatility of volatility. The double-index notation reflected the three-factor nature of the model, so everything looked OK. Are you now saying that as soon as associativity is invoked and we realize that all we have on our hands is a nine-regime-switching model $[\sigma_k]$, we lose the third factor? How could a change of notation, or in other words, a mere *change of name*, have such deep consequences?

Underlying your worry about Nobody failing to account for a multi-factor situation is in fact just a worry about different *names* that a thing can be called. And we had warned you that Nobody was precisely the model with no name, and the regime precisely the state variable with no particular label and no particular dimension attached to it! To put it differently, a regime-switching model which has the name of a *single* factor – for instance volatility – written on each regime, is not necessarily a one-factor, or a two-factor, or a three-factor model of stochastic volatility. It can be anything[3]. What it is really will largely depend on the "depth" and the variety of the derivative instruments we are calibrating against. Let me put it this way. If Nobody calibrates successfully to prices of vanillas, cliquets, as well as higher-order structures which may be sensitive to the volatility of volatility of volatility, and none of the other models, commonly known as one-factor or two-factor, does, then Nobody may be said

[3]Of course the picture would have been different if, instead of considering a *nesting* of stochastic processes written one off another, such as stochastic volatility, stochastic volatility of volatility, etc., we had considered a radically different second factor, such as the hazard rate. For in that case, the regimes could not be labelled by anything short of two factors.

to be of a "level higher than two-factor." Pushing the argument a little further, we may even wonder: Why should what we have to say about Nobody depend on what other models can do? For all we know, the other models may have never existed. Is it not our purpose to break with tradition anyway? All that we can hope to say, then, is that Nobody *was* able to calibrate to the prices of certain derivatives instruments traded in the market, full stop.

10 And by the way, isn't the whole language of "number of factors" just a heritage of the tradition? At least this much is sure: the regime-switching representation is not reducible to a classification in terms of number of factors. This, we have shown with two complementary arguments: a) the argument that the regime can be identified by an n-tuple of names relating to different financial categories (volatility, hazard rate, interest rate) and b) the argument that it can alternatively be identified by a single name (for instance, volatility) yet the picture be richer than the traditional two-factor framework. There is no question that talking of multiple factors is legitimate in the first case. When we contemplate stochastic volatility, stochastic hazard rate and stochastic interest rate, the situation is different from the one where the second process concerns the volatility of the first, and the third the volatility of the second. Typically, derivative instruments can have radically different underlyings in the first case. The underlying that a credit default swap is written on (the state of default or no default of certain issuer) is very different from the underlying of an equity option, which is in turn very different from the underlying of an interest rate option. Derivative instruments such as the vanilla equity option and the cap and floored cliquet, by contrast, are written on the *same* underlying, although we did refer to them, a paragraph back, as instruments of different "depth." Even a variance swap is written on the same underlying as the vanilla option. The only difference is that its payoff explicitly depends on a variable, realized variance, which can only be measured over the whole path of the underlying.

The point of this distinction is to suggest that talking of "multiple factors" in the second case might be misleading and might have imposed itself on us for no other reason than the written tradition and the tradition of *writing* – for no other reason indeed than that a model, say, of stochastic volatility of volatility, has traditionally been *written* in three lines, as shown above. The volatility process has a diffusion coefficient ε and this coefficient diffuses in turn, etc. A new line is written every time and this suggests that a new kind of derivative instrument, specific to the new kind of factor, can be written every time. Vanilla option writing is specific to stochastic underlying; variance swap writing is specific to stochastic volatility, etc.

When you think about it, however, all that is really meant by the three lines written above – or any additional number of lines for that matter – is jut a complex stochastic process the underlying is supposed to be following, and the corresponding complex probability distribution. Again, compare the situation where some truly different processes are involved: a hazard rate process, an interest rate process.

The other real things are the derivative instruments. True, they may be differently styled, and may admit of different levels of complexity, but they all fall back, in the end, on the underlying they are written on. Our writer has created himself a fiction (that the Brownian diffusion coefficient might be diffusing) and he is now growing a new fiction inside the fiction (that the diffusion coefficient of the coefficient of the Brownian diffusion might itself be diffusing, etc.). Think, by contrast, how the un-writable, un-nameable, dimension-less, story of regimes manages to describe the world just as well, or even describe it better (calibrate to the traded prices of derivatives, propose optimal hedging strategies), yet offers no predetermined format, no predetermined number of lines, to fit the story in.

Another way of looking at things, and seeing how our writer entraps himself in his own fiction, is to compare the *a priori* attitudes of the traditional representation and the regime-switching representation, when both are brought face-to-face with the market. By committing himself to two or three lines of writing, the traditional writer faces a market which can be completed *a priori* with the help of two or three traded instruments. Markets can be completed *a priori*, under Heston or any other "first-level" stochastic volatility model, by trading an option together with the underlying as a dynamic multi-hedging strategy, and they can be completed *a priori*, under "second-level" models such as the one written above, by adding one further instrument, etc. Our writer faces a market that he knows he can complete *a priori* with the help of a given number of instruments, regardless of the variety, the depth, and the price structure of the market he will face in effect. The regime-switching representation, by contrast, does not impose such strictures on the future story. It cannot tell, in advance, the degree of incompleteness of the market. It cannot tell *a priori* at which particular "writing level" it will fall. Only the market and the result of calibration can. Philippe Henrotte, the head of theory in my company, and the father of Nobody, summarized the point brilliantly: "The reason why the regime-switching representation cannot be completed *a priori* is precisely that the regimes bear no particular name!" As Jacques Derrida, the leading figure of French theory, would put it, we've been held captive by a long tradition of writing and logocentrism. Indeed the noticeable consequence of writing and naming is to make present and actual for us what could very well be different and have to be *deferred*, what could only *be* later.

11 The point is of importance because the real test of a smile model is to see how robust it is to a reality that may contradict its assumptions. (This is the whole point of the *metatheory* of smiles.) Complete markets are the least robust notion. You perturb them a little bit and they become incomplete. Surely enough, Nobody will fall on what we have called a "particular writing level" once it is calibrated; the market it describes will assume a certain degree of incompleteness, and it will be susceptible of completion by using the appropriate number of hedging instruments in the dynamic hedging strategies. The point is that such degree and such level are dictated by the reality of the instruments we calibrated against. They are determined *a posteriori*, not *a priori*. Should a new instrument become traded the next day, and its price fall outside the range of prices that were attainable by the completed market of the day before, then new calibration to that new instrument will open new levels to Nobody,

and new degrees of incompleteness. By contrast, you cannot but throw away your Heston, when such a situation occurs. This also tells us that we should always leave the door open, in Nobody, for such new possibilities and such self-upgradings. We should not strive to complete the market, at any level of writing that we may be standing. To keep our hedging robust, we should always keep it optimal, and never try to make it perfect. HERO is the measure of residual risk borne by the optimal hedging strategy[4]. What is interesting is how HERO decreases when additional instruments are used in the hedging portfolio, or in other words, the hedging opportunity that the additional instruments may offer. But HERO should never be driven down to zero. What is interesting is the way we approach completeness, not completeness as such. For surely we lose the sense and the measure of all that when HERO is equal to zero!

What we are really saying is that we might be offered a chance, with the regime-switching representation, which was not available in any smile model before: the invaluable possibility that Nobody might just be equipped to deal with *model risk* on top of the risk which is the normal, *contained* subject matter of the models of risk. Somehow, Nobody might be "aware" of its own metatheory. Recall that Nobody is not just a first order stochastic volatility model, but can potentially instantiate any of the higher order models corresponding to the higher levels of writing. This level-invariance is due to the fact that a regime-switching model *made stochastic*, and consequently calling for "regimes of regimes," is in the end just another regime-switching model. The regime-switching representation does not iterate or pile up in the same fashion as the traditional lines of writing. Only the variety of the derivative instruments we are calibrating against and the richness of their price structure can set the level of writing for us. For instance, a newly traded in instrument and a new range of prices can precipitate an "internal" transition from a certain level to a higher level. Our stochastic volatility model suddenly becomes a stochastic volatility of volatility model, yet with no visible change. We will still be looking at the same old regime-switching representation, only we will be calling it different names. As a matter of fact, this can work both ways. Why should the change be invisible in the one way, and not in the other? What is to stop us indeed from thinking that the stochastic volatility of volatility model was already available to us the day before? True, we may have not calibrated to the *relevant* instruments the day before, or the relevant instruments may have not been available the day before, but is this reason not to think that the *thought* was available to us the day before?

Put differently, Nobody may have already opened itself to the self-upgrading the day before, only the visible derivative instrument, specifically crystallizing the upgrading, was simply not available the day before! Maybe the upgrading was *in part* contained in the set of derivative instruments of the day before, and was just missing one last instrument to express itself in full. Things may have already been "in the

[4]Our hedging strategies are optimal in the sense that you break even on average and the standard deviation of the P&L of the hedged portfolio is minimal. HERO is this minimal standard deviation. (cf. [5].)

cards," as the saying goes, or in other words, the market may have already anticipated its own upgrading the day before. Or we may argue the other way round. Although the new instrument takes us up one level, there might be no reason to believe that its introduction will automatically bring a mutation in the market. At least not the first day. Chances are, on the contrary, that the newly created instrument – I am thinking, for instance, of a new complex structure, something like a complex cap and floored cliquet, which, although written on the same underlying, is sensitive to higher order distortions, such as volatility of volatility of volatility – chances are that it will begin its journey right on the tracks from the day before, as the market participants will no doubt start pricing it with state-of-the-art models not yet aware of the next level. It is only when the instrument comes alive and starts leading a trading life of its own that the true change in the market will begin.

I guess my whole point is that changes of "writing levels," or degrees of incompleteness, can take place smoothly within the regime-switching representation, and can only cause breaks and fractures within the *writing tradition*. Nobody can thus enjoy continuity of life and allow us, for the first time, to really address the question of *history*. Indeed a big question, perhaps the biggest, is whether the given smile model should be calibrated to the instant prices of derivative instruments or to their history. This is the story of re-calibration looming again. People are sooner or later led to back-test their model and they become very excited when they notice that the parameters of the model are stable over successive re-calibrations. They think they have hit upon some deeply significant invariant. Whether they admit it or not, everybody is striving towards this end, and calibration to the history of prices is one way of making it explicit. The main objection, however, is that any time series of any given length will end up revealing some invariant or other when submitted to a model of sufficient complexity. This is another way of objecting that the "true" data generating process may in fact admit of no finite moments and may require time series of infinite length before any parameter is stably estimated. As a matter of fact, being able to *name* an invariant may be the worst thing that could ever happen to the searchers engaged in that kind of quest. For what is then to stop us from writing an extra line, and turning the name of the invariant into the name of a new stochastic process?

12 "Model risk," "necessity of re-calibration," "endlessness of the writing process," are all different names for the same big problem. Perhaps *the* big foundational problem which puts into question the very possibility of quantitative finance as a science [6, 2]. We all know that the option pricing tools that we are using are forward-looking. Precisely for that reason, we all know that they are subject to the necessity of re-calibration and to the threat of model change, or in other words, model risk. Yet we lack the means to address that problem almost by definition. A model can do anything except look into its own assumptions. ("What by definition can hurt you most is what you expect the least" – Taleb[8].) This is why we almost inevitably turn to history as the only way out of our predicament. That we may not have the faintest idea how to address the problem of re-calibration makes historical calibration look useful, to say

the least, in comparison. Although a backward-looking procedure can be the last thing we want to consider when addressing a forward-looking question, it is unfortunately the only thing available. We have no choice but to entertain the momentary hope that history might repeat itself.

Now think again of the essentially nameless character of our regime-switching representation. Since our regimes bear no particular name (the name of volatility, or volatility of volatility, etc.), the temptation is simply *not there* to look back at the recent history of calibration of Nobody, and try to identify a stable parameter. First of all, it is not even clear what level of writing is getting instantiated everyday! Alternatively, we can look back at the series and read into it any story we want. We can argue, for instance, that the full richness and the full incompleteness of the market were right there with us from the start – just as Nobody was with us from the start! – only the hidden variables became manifest, and Nobody was able to calibrate to them, from that day when the relevant complex instrument became alive and started leading a trading life of its own. So I guess the second part of the recommendation never to complete the market and always to leave the door open for future upgrading is the recommendation never to trust that Nobody *has been perfectly calibrated to the market*. "Calibration is just another word for completion," as Philippe Henrotte says, for an instrument whose price process we could not attain with a suitable self-financing strategy involving the existing calibration instruments would mean that a new level is up and that the new instrument adds richness to our existing information set. Therefore it should independently be included in the calibration procedure, and our past calibration was not right.

Whatever interpretation it may be that we care to put on Nobody's recent history of calibration, the fact remains... that only the facts remain. Whether the richness was all there from the start and Nobody was simply not "perfectly calibrated to the market" or whether the richness all emerged when the new instrument first diverged from the known tracks and opened the door for an actual upgrading of Nobody, is in the end a purely nominalistic issue and just a question of how we want to *call* the story. The absence of writing inherent in Nobody relieves us completely of the necessity either to read the past into the future or to read the future into the past. Only the fact of re-calibration remains and the hope, like I said, is that the HERO shall pick up a component of model risk (or meta-risk) on top of the standard risk that it is picking up at the object level, thus allowing smoothness of upgrading. Although theoretically unjustified – for the computation of HERO presupposes that the level of writing is fixed and the model is final –, this is just the hope that successive re-calibrations and successive re-hedging operations shall prove robust *as a matter of fact*. Since a "stochasticized" Nobody is just an instance of Nobody, the hope is that the two instances shall bear a few similarities other than just by "name"!

In the end, the reason why I think that Nobody may just be offering us a chance to *finally* address the question of history, is that this is the question whose answer we should expect to be the least naïve of all, or in other words, the least *expected* of all (for any answer falling within our range of expectations will become history, therefore will be overtaken by history), and that Nobody, whose inherent non-writing is essentially a non-answer, seems to offer just the right kind of divergence and the right kind of... digression. By remaining open to the future and by coming out virtually unchanged through the future upgradings, Nobody in fact *postpones* any temptation we may have to look back at the past in order to figure out the future. (Like I said, the temptation may be due to a lack of choice.) Nobody takes over the task that we thought was reserved for history, the task of softening up the future for us, and teaching us patience in matters strictly relating to the future (as if history was the only thing we could read in patience while waiting for the future). As it is *both* forward-looking and capable of self-upgrading, Nobody allows us at last to deal with the future *seriously* (when we thought history alone could achieve that purpose). In fact, Nobody reinstates the balance of power in favour of the *present* (rather than the past or the future). Nobody is the perfect tool in the hands of the trader (as I shall argue in a philosophical column in Wilmott magazine) and the pair that the trader and Nobody will thus constitute is, in the last instance, essentially *present* (like the living trader is present) and utterly self-upgrading (or futuristic).

13 So far, my strategy has consisted in arguing that default lent itself naturally to the discrete regime-switching representation on account of the two-valued nature of the state variable (default / no default) and that the regimes could be further extended to the no default side of the picture – naturally so as concerns the hazard rate (as different regimes got interpreted as different ratings), and not so naturally as concerns volatility (as people traditionally had in mind continuous diffusion processes of volatility). As a matter of fact, the second part of my argument very quickly developed into a formidable digression on the regime-switching representation, whose main purpose was to establish the originality, and I dare say, the *uniqueness* of the proposed solution. Beside the advantages specific to the regime-switching representation (its economy, its computational efficiency), and beyond the equity-to-credit problem, I have attempted to show that the regimes might in fact provide just about everything everybody ever wanted in a smile model and could not get before: robustness of calibration and its correlate, the absence of presupposition with regard to the degree of incompleteness of the market (and adaptation, instead, to the effective degree of incompleteness through effective calibration), as well as the beginning of an answer to the abysmal question of re-calibration. Incidentally, the argument earned us a criticism of the traditional representation, the tradition of *writing*, and debunked its self-created myths and mythicized names. Philosophically, this meant that our regime-switching representation, Nobody, had seized mastery of the subject and taken precedence over the other models, as it was able both to propose a working solution and to take an unprecedented metatheoretical stand.

It remains, however, to clear one last obstacle before moving on to the problem at hand, the equity-to-credit problem. I would like to call this obstacle the "continuous-discrete fallacy." In a word, this is the general worry that the regime-switching representation might just not be suitable for derivative pricing *because* of its discrete character. All is well when the variables assume two, or perhaps only a few, discrete values (such as credit ratings or the default state), but is a regime discretization of the credit spread able, for instance, to handle the pricing of options *on the credit spread*? Is a three-regime-switching model of stochastic volatility able to price volatility swaps? Also, how could such a formidable campaign ever be launched in the field of smiles, and such a systematic attack ever be mounted against the traditional "continuous" smile models, from such a frail and discrete base? Isn't the regime-switching model "negligible," and almost degenerate, in the space of the smile models?

The proximate answer is that the convertible bond is an equity derivative, not a credit spread derivative, and the equity is still getting modelled as a continuous diffusion process, spanning all values from zero to infinity, inside each one of our regimes. As shown by our numerical examples (see Appendix), recognizing two or three hazard rate regimes (actually four, if we count the default regime) is amply sufficient to capture the impact of default on the pricing and hedging of the convertible, or indeed to fully explain the term-structure of credit spreads. Two or three volatility regimes are also sufficient to capture volatility risk and the relevance of vega hedging. Actually, a three-regime stochastic volatility model or a three-regime stochastic hazard rate model are much richer than you think. The intensities of transitions between regimes participate fully in the specification of the model. Also, let us not forget the whole orientation of my essay. My main topic is the equity smile, not credit risk, and my task is the continuation of the work achieved in [5]. Like I said, I am willing to consider default only to the extent that it is a component of the equity smile and that the credit default swap prices can help us calibrate and hedge our equity smiles better.

The ultimate answer, on the other hand, is that the regime-switching model is not negligible after all. Nothing stops us in theory (or in practice) from multiplying the number of hazard rate or volatility regimes up to the point where they become indistinguishable from the discretization of a "real," continuous space. Nothing stops us from assuming as many regimes (σ_i, λ_j) as there are pairs in the tensor product of the full (discretized) continuous volatility space and the full (discretized) continuous hazard rate space. Solving a diffusion PDE for the underlying equity in each one of those regimes, and coupling this incredibly large number of one-dimensional PDEs through the usual transitions between regimes, will then turn out to be numerically indistinguishable from discretizing a full three-dimensional PDE in the $(\mathrm{S}, \sigma, \lambda)$ variables, and solving it with the usual techniques. As a matter of fact, the argument can work both ways. For the equity diffusion process that we had assumed, so far, was taking place in each one of our parsimonious regimes, and its numerical treatment by the usual PDE discretization techniques, can now in turn be interpreted within the regime-switching representation. Discretized Brownian motion is just a regime-switching model with a large number of tinily spaced regimes. Numerically speaking, it is all then but a massive regime-switching operation. All I am saying here is that, ultimately, everything becomes discrete. (Not mentioning that everything, actually, is *initially* discrete, as stocks trade by the tick and hedging takes place in discrete time.)

In conclusion, the regime-switching representation is the one and all-pervasive representation. Continuous path processes and continuous time Finance are but useful fictions that were invented to *summarize* the incredibly disarticulate picture left over by the regimes. The name of "volatility" (Black-Scholes), and following it, the name of "stochastic volatility" (Heston), etc., are *given names* that merely indicate that the wiring between the regimes has been laid out in a certain way and not another. Two or three names, two or three coefficients (volatility, mean-reversion, correlation, volatility of volatility, etc.), throw order and rigidity on an incredibly rich and multifarious picture.

This immediately poses the following question: Among the two extreme situations we have described, which one is in fact the poorest? Is it the situation where we have only a few volatility and hazard rate regimes and have no particular names attached to the model, or the situation where we have infinitely many regimes but only a couple of names? Might not the order in question be a restraining order? The main objection against Nobody was that the discrete regime representation might be missing something important that the continuous representation could provide. Shouldn't we be worrying about the opposite by now? Shouldn't we be worrying that the continuous representation might be giving us much more than we actually need, and charging us a very high price for it – imposing names on us that we might not even need? Let us not forget indeed that the tradition of writing and naming creates strata that are impermeable to each other (volatility, volatility of volatility, etc.), when the nameless regimes manage at once to *see through* the whole thing. As the numerical fate common to both representations demonstrates their equivalence in the limit, and shows that it is all but a question of representation, perhaps we could now step back from the limit and try to get the best of both worlds. Not wanting an infinite number of regimes and not wanting the names either, perhaps we could find, in the midfield, the compromise which is best adapted to each particular situation. At least we want the freedom to do so.

Again, recall the freedom to be in any of the possible "writing levels," that Nobody is allowed by the absence of names. And imagine, for instance, a situation where calibration to a newly introduced derivative instrument has exhausted all the possibilities offered by three regimes of volatility and still could not be achieved to our satisfaction. This is typically the situation where we contemplate adding a regime. Now adding a fourth regime achieves much more in effect than adding a fourth state of volatility. It opens for us whole new levels of writing and new degrees of incompleteness not previously available. On the other hand, it can have richer consequences than adding a line of writing as in the classical writing tradition. The writing levels of the classical tradition are stratified. Below the name of "volatility of volatility of volatility" there just lies the immutable name of "volatility of volatility." Whereas the three regimes that lie "below" the fourth can now react to calibration in ways we could not even dream of when all we had was three regimes. Actually, the three regimes are not *below* the fourth, but at the same level. There is a sense of wholeness, of richness, and an overall feeling of economy in the regime-switching representation that is unavailable to the stratified writing tradition.

In the end, I do not exactly place the interesting debate between the *discrete* and the *continuous*. This is a wrong divide, and the two are in the end equivalent. Rather, our "digression" into the regime-switching representation will have served the following philosophical purpose (beside providing us with an extraordinary calibration, pricing and hedging tool). It has shown us that the real difference is between worshipping the name and breaking the name, between iconolatry and iconoclasm, if you will.

I have nothing against the Black-Scholes option pricing model, or the Heston option pricing model. Numerically, that is to say, ultimately, they are but instances of Nobody. I guess what I have against you is that you may be tempted to build up your world like a venerable writer, not like an engineer. You may be tempted to recognize in Black-Scholes or in Heston nothing but the names and the strata that will allow you to append your own model and write your own name.

Analytical pricing formulae are very hard to come by nowadays because of the complexity of the derivative instruments and the complexity of the underlying processes. Computation power, on the other hand, is allowing numerical speed-ups which compare with the analytical formulae of yesterday. So what could be the point of insisting that the model should be cast in terms of elegant continuous processes – so remote indeed from our nameless, shameless, mess of a regime-switching representation! – other than the facility of writing the model *on paper*, and producing a nice mathematical paper? True, continuous path Brownian motion is what afforded the perfect replication in Black-Scholes and created the myth of the complete markets. But the complete markets are precisely the thin leaf that we should try to escape at all costs, or only admire from a distance as an elegant argument written on a nice piece of paper! Complete markets are enmeshed with the guilty "tradition of the name" anyway. Options are redundant in the Black-Scholes world; they do not truly exist and are only the diminutive *name* of a particular dynamic trading strategy involving the underlying alone.

Derivative pricing science has been taken hostage by mathematicians when it should be handed back to the financial theorists – How many working solutions mention incomplete markets and real hedging? – and to the engineers. But how come, you may wonder, a philosophical argument as general as this – the argument against the naming and the writing – happens to occur precisely in the field of Finance, and more particularly so, the field of derivative pricing? Doesn't the same massive attack equally apply in other engineering fields? What is so specific about derivative pricing that allows me today to draw such deep conclusions about the fallacy of naming and the lure of writing?

The first answer is that derivative pricing is only *starting* to become a proper engineering field. Inelegant and massively computational solutions have been occurring for some time now in fluid mechanics, or structural mechanics, or thermodynamics. The second and most important answer is that we enjoy a freedom in our specific engineering field which is hard to find in other fields. Laws of gravity, laws of mechanics, and more generally laws of nature, compel an Einstein, a Schrödinger or a Navier and Stokes to *write* the models they have written. (Yet I am not so sure that an antirealist will not argue that laws of nature *are not* written in nature after all, and that the physical theories that we have, and their quantitative models, are mere computational tools.) We, by contrast, can "write" or "name" or "program" the model that we want, so long as the model is robustly calibrated and some derivative instruments are robustly priced and hedged relative to some others. (The searcher of the ultimate data generating process can wait all the time he wants, even wait infinitely – but then he *is not* in the business of derivative pricing.) We have all this freedom, yet some writers insist on following the inherited path of research, and the inherited lines of writing! Path-dependency is an even worse case than inhomogeneity...

Above all, I think the real strong argument for wanting to erase the previous lines of writing and having the fresh start that I suggest, is that the *open* regime-switching representation can address the problem of *co-calibration* just as easily as it addressed the problem of calibration or re-calibration. Take the equity-to-credit problem, for instance. (This way, I can circle back to my original topic.) A lot of effort has been spent recently in developing credit risk models quite separately from the smile models. While jump-diffusion, local volatility, stochastic volatility or universal volatility have been suggested on the one side, stochastic interest rates, stochastic hazard rates and stochastic recovery rates have been suggested on the other. My whole argument about the level-invariance of Nobody with respect to the volatility factor can of course be reiterated with respect to the hazard rate factor. Nobody can scale up to any degree of incompleteness that the market of credit derivatives may impose in effect (stochastic hazard rate, stochastic volatility of hazard rate, stochastic volatility of volatility of hazard rate, etc.), just as it scaled up in the volatility case. More interesting, perhaps, is the fact that Nobody can co-calibrate to credit instruments and volatility instruments *with no visible change*. Wouldn't we have to wait, otherwise, for two *distinct* traditions of writers to deliver to us their successive lines of writing?

This isn't just going over the fact that the regimes of our regime-switching representation can be identified by a pair (σ, λ), for I am now implying a deeper phenomenon, namely, that the volatility skew implied in the market prices of out-of-the-money puts can positively give us information on the default process and, reciprocally, that the CDS term-structure of spreads can positively give us information on the value of equity options. The equity-to-credit problem is precisely a smile problem. As a matter of fact, it may even be better-posed than the pure smile problem, as the information from the CDS will have a tendency to help the calibration, and help determine the solution (see Appendix). Provided, of course, our smile model and calibration tool can *handle* co-calibration, and make it a friend, not an enemy. Nobody *loves* the idea that CDS prices can be included in the calibration set, when other (stratified) models most certainly resent it! What is a complication in the traditional representation is a simplification in ours. This is exactly the correlate, in co-calibration, of the idea we have already explored in calibration and re-calibration, according to which Nobody loves the prospect of adding barrier options, or cliquets, or more complex structures still, in the calibration phase.

Appendix

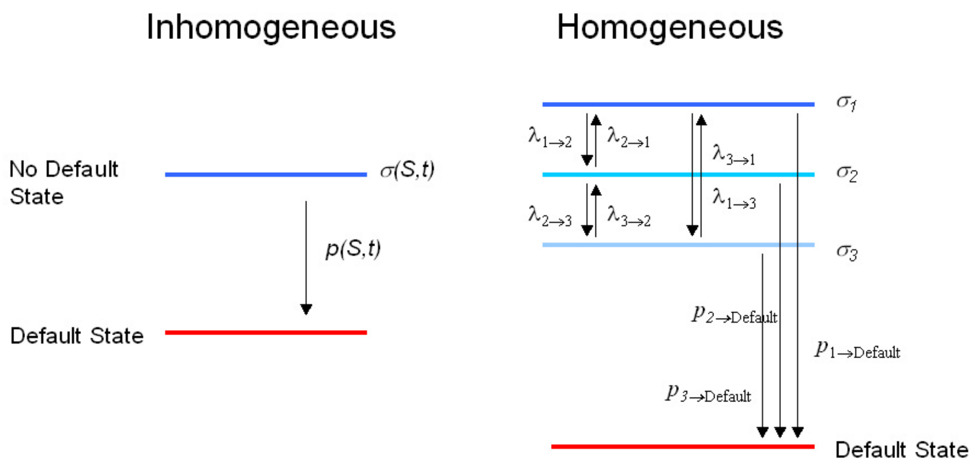

Figure 1: Comparative logic of the inhomogeneous and homogeneous equity-to-credit models. The inhomogeneous model consists of one default regime and one non-default regime. The implied volatility surface and the credit smile surface are explained by a local volatility surface and a local hazard rate surface. In the homogeneous model, volatility and hazard rate are stochastic and switch between the three non-default regimes.

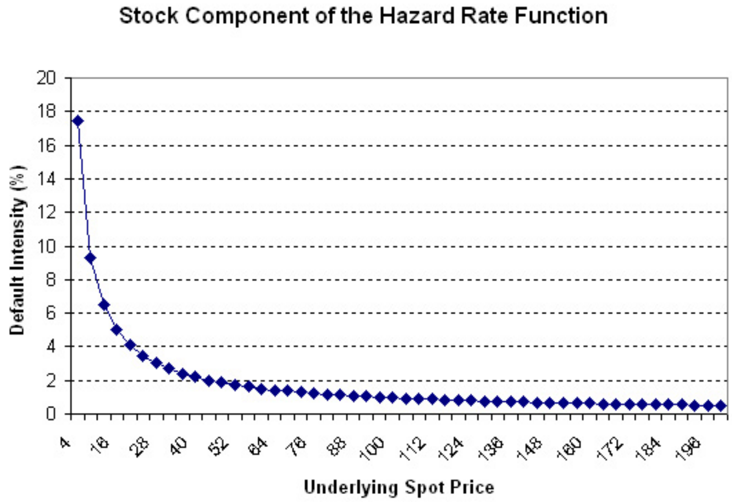

Figure 2: In our inhomogeneous model, the hazard rate function is the sum of a time component $g(t)$ and a space component $f(S)$. The space component accounts for the dynamics of the credit spread curve against the moving underlying stock. The time component ensures that a given credit spread term structure is matched for a given stock level. Above is the plot of the space component, $f(S) = p_0 \left(\frac{S_0}{S}\right)^{\beta}$, $S_0 = \$16$, $p_0 = 5\%$, $\beta = 0.9$.

Maturity Date	15/01/2008
Semi-annual Coupon Rate	2.75%
Nominal	100
Conversion Ratio	4.38
Recovery Rate	0

Table 1: Terms of a Convertible Bond (The Tyco 2.75% 2018 to the first put date).

Figure 3: Theoretical value of the convertible bond described in Table 1 against the underlying equity in a inhomogeneous equity-to-credit model. The pricing date is 03/02/2003. The interest rate is flat 3%. The credit spread term structure is given in Table 5 when the underlying stock is $S = \$16$. The Brownian volatility is 45.3%. The CB is worth \$105.50 against $S = \$16$, its Delta is 3.80 shares, its fixed income component is worth \$81.75.

Stock Price	**Equity-to-Credit Delta**	**Delta under static spread**
16	3.81	2.91
14	3.86	2.66
12	3.96	2.37
10	4.16	2.01
8	4.58	1.59
6	5.38	1.10
4	7.33	0.57
2	10.80	0.13

Table 2: Convertible bond in the inhomogeneous model. We report its Delta against stock levels under the equity-to-credit model where the hazard rate function is given in Figure 2. The equity-to-credit Delta increases on the way down as the bond floor collapses. We also report the delta under static spread for comparison.

	Brownian Diffusion	Total volatility
Regime 1	49.86%	61.18%
Regime 2	27.54%	40.83%

	Jump size	Jump intensity
Regime 1 → Regime 2	4.48%	3.3429
Regime 2 → Regime 1	-58.68%	0.1697
Regime 1 → Default Regime	-100%	0.1190
Regime 2 → Default Regime	-100%	0.0324

Table 3: Calibrated parameters of the homogeneous regime-switching stochastic volatility and stochastic hazard rate model, **Nobody** (two non-default regimes and one default regime). **Nobody** is calibrated to the full implied volatility surface given in Table 4 and the full credit default swap spread term structure given in Table 5. The source of market data is Tyco on 03/02/2003 and the underlying stock is $S = \$16$.

						Strike						
Maturity (years)	**5**	**7.50**	**10**	**12.50**	**15**	**17.50**	**20**	**22.50**	**25**	**30**	**35**	**45**
21/02/2003	**Market**		158.10%	112.70%	79.80%	58.10%	49.40%	56.30%	72.40			
	Model		175.99%	122.30%	76.07%	55.80%	48.39%	48.09%				
21/03/2003	**Market**		122.20%	92.90%	71.20%	56.00%	48.40%	45.40%	53.10%			
	Model		129.28%	93.11%	66.24%	54.82%	49.39%	49.13%	47.17%			
17/04/2003	**Market**	138.20%	108%	82.80%	66.30%	54.60%	47.20%	45.40%	45.70%	53.20%	65.30%	78%
	Model	150%	112.07%	83.38%	63.62%	54.07%	49.06%	48.28%	46.91%	45.78%	44.90%	43.50%
18/07/2003	**Market**	99.10%	87.30%	72.60%	60.50%	52.10%	47.00%	44.70%	43.60%	44.30%		
	Model	113.80%	88.77%	71.43%	59.45%	52.22%	47.78%	46.14%	44.72%	43.87%		
16/01/2004	**Market**	92.40%	75.40%	64.40%	56.50%	51.40%	47.10%	45.00%		42.80%	43.20%	45.20%
	Model	89.87%	73.99%	63.47%	55.90%	50.68%	46.98%	44.94%		42.07%	40.91%	40.49%
21/01/2005	**Market**	73.70%		58.60%		49.40%		46%		42.70%	41.10%	41.10%
	Model	74.50%		58.34%		50.40%		45.88%		43.02%	41.31%	40.19%

Table 4: Quality of fit of a full implied volatility surface with with the equity-to-credit homogeneous regime-switching model. Source: Tyco on 03/02/2003. The underlying stock is $S = \$16$.

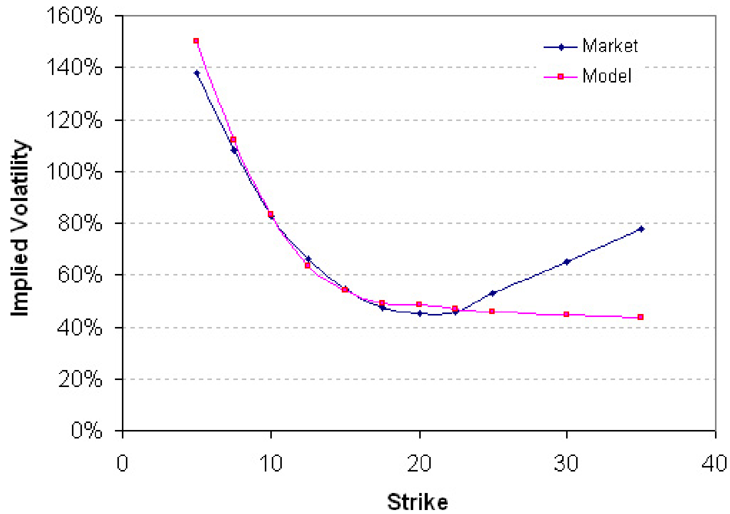

Figure 4: Zooming on the quality of fit of the implied volatility smile of options maturing on 17/04/2003. The model does not do such a good job on the out-of-the-money calls because of the crazy prices reported in the market. Probably those calls are not so liquid.

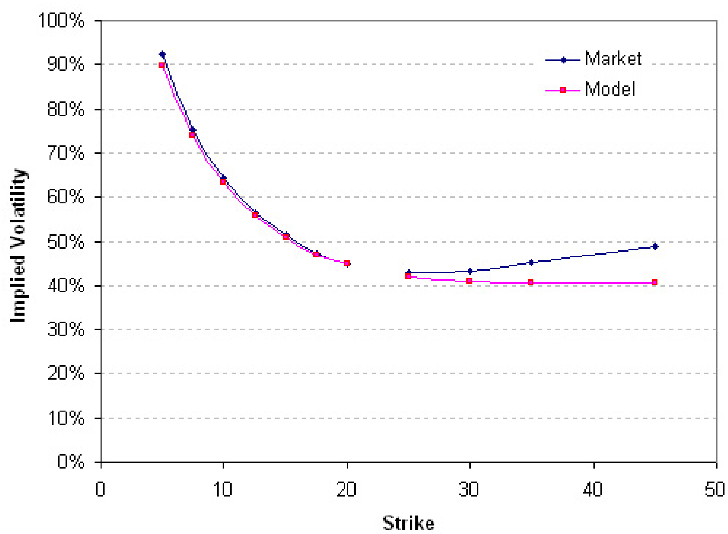

Figure 5: Zooming on the quality of fit of the implied volatility smile of options maturing on 16/04/2004. The model does not do such a good job on the out-of-the-money calls because of the crazy prices reported in the market. Probably those calls are not so liquid.

Maturity (years)	Market premium (quarterly paid coupon)	Model premium
1	1.25%	1.50%
2	1.17%	1.22%
3	1.14%	1.13%
4	1.11%	1.08%
5	1.09%	1.05%
6	1.05%	1.03%
7	1.03%	1.02%
8	1.01%	1.01%
9	0.99%	1.00%
10	0.98%	0.99%

Table 5: Quality of fit of the credit default swap term structure with the equity-to-credit homogeneous regime-switching model. Source: Tyco on 03/02/2003. The underlying stock is $S = \$16$.

Stock Price	Optimal stock hedge ratio	HERO
16	3.79	9.62
14	3.70	10.25
12	3.60	10.93
10	3.52	11.68
8	3.50	12.48
6	3.71	13.29
4	4.64	14.04
2	8.64	14.51

Table 6: The Tyco convertible bond is dynamically optimally hedged, in the homogeneous model, *with the underlying stock alone*. The simulation takes place on 03/02/2003. It uses **Nobody** with the parameters inferred from calibration given in Table 3. The CB is worth $\$105.50$ against $S = \$16$ (same reference point as in the inhomogeneous model). We report the HERO, expressed in dollars, and the optimal dynamic hedging ratio. The hedging ratio is expressed as the equivalent stock position in number of shares (you should sell the shares in order to hedge). Notice that the hedging ratio for $S = \$16$ is 3.79 shares, very close to the delta in the inhomogeneous model against the same stock level (see Table 2). Optimal hedging under default risk is to our mind the real reason why you should go heavy on the delta, not some deterministic function linking the hazard rate and the stock.

Stock Price	**CDS hedge ratio**	**HERO**	**Stock hedge ratio**
16	-55.9	5.00	3.05
14	-59.7	5.10	2.80
12	-63.8	5.14	2.48
10	-68.2	5.06	2.08
8	-72.7	4.76	1.58
6	-76.9	4.11	1.00
4	-79.9	2.89	0.41
2	-81.2	1.06	0.04

Table 7: The Tyco convertible bond is dynamically optimally hedged in the homogeneous model with a *combination of the credit default swap and the underlying stock*. We use the 5-year maturity CDS whose premium is reported in Table 5. We report the residual HERO and the dynamic hedging ratios. The CDS hedging ratio is expressed as equivalent CDS position (you should buy the CDS and short the stock to achieve the optimal hedge). The CDS hedge is in percentage of nominal. As the CDS takes care of the major jump due to default risk, the stock hedges againt the diffusion and other small jumps. Its contribution in the hedge is very similar to the Delta under static spread. (See Table 2.)

Stock Price	**Hedging Ratio (Call option)**	**Hedging Ratio (CDS)**
16	4.18	-80.23
14	4.16	-80.36
12	4.12	-80.46
10	4.05	-80.53
8	3.92	-80.59
6	3.63	-80.67
4	2.62	-80.78
2	-13.17	-80.93

Table 8: The convertible bond is dynamically optimally hedged in the homogeneous model with a *combination of the underlying, the CDS, and a call option* of same maturity as the CB and strike price $K = \$22.50$. We report the dynamic optimal hedging ratios on both the call option and the CDS. The HERO is now almost identically equal to zero as the CDS and the call option cancel the default risk and the volatility risk. Notice the stability of the hedging strategy. The CB is almost perfectly decomposed into a volatility instrument and a credit instrument.

Stock Price	Hedging Ratio (Put option)	HERO
16	-41.59	1.45
14	-42.29	1.69
12	-43.28	1.99
10	-44.79	2.41
8	-47.27	3.01
6	-51.87	3.91
4	-62.87	5.45

Table 9: We now try to take full advantage of our equity-to-credit homogeneous model. We analyze a **non convertible** bond, otherwise identical to the Tyco CB. This is a pure credit play. We use an out-of-the-money put option of same maturity as the bond and strike $K = \$3$, to hedge default risk. We report the dynamic optimal hedging ratio in equivalent put position (you should buy the puts to achieve hedging) and the HERO. The HERO increases as the stock goes down because of the increasing volatility risk borne by the puts. We could have hedged default risk with the underlying alone, but HERO would have been much larger ($\$14.60$ against $S = \$16$).

Stock Price	Hedging Ratio (Put option)	Hedging Ratio (ATM Call option)	HERO
16	-44.54	1.02	1.22
14	-46.25	1.54	1.36
12	-48.82	2.45	1.54
10	-52.98	4.21	1.74
8	-61.27	8.52	1.95
6	-80.97	21.45	2.21
4	-184.26	110.61	2.53

Table 10: To hedge against the residual volatility risk manifested in Table 9 we now add an at-the-money call in the hedged portfolio involving the corporate bond and the out-of-the-money put. We report the dynamic optimal hedging ratios on both options and the HERO. Obviously, we should sell the ATM call.

References

[1] L. Andersen and D. Buffum. Calibration and implementation of convertible bond models. October 2002.

[2] E. Ayache. The philosophy of quantitative finance. Forthcoming in Wilmott.

[3] E. Ayache, P. A. Forsyth, and K. R. Vetzal. Next generation models for convertible bonds with credit risk. *Wilmott*, December 2002.

[4] E. Ayache, P. A. Forsyth, and K. R. Vetzal. Valuation of convertible bonds with credit risk. *The journal of Derivatives*, pages 9–29, Fall 2003.

[5] E. Ayache, P. Henrotte, S. Nassar, and X. Wang. Can anyone solve the smile problem? *Wilmott*, January 2004.

[6] E. Ayache and D. Tudball. One step beyond. *Wilmott*, page 6, January 2004.

[7] P. Henrotte. The case for time homogeneity. *Wilmott*, January 2004.

[8] N. Taleb. 2001 hall of fame. *Derivatives Strategy*, March 2001.