WILMOTT
Aug 2001

Ele Ayache relates a tale of

THE DISCRETE AND THE CONTINUOUS

One question that has long preoccupied the mathematical physicists is whether the laws of nature are discrete or continuous. The rejoinder of course is whether the questioner means to say that the laws or nature are discrete or continuous in reality (whatever that means), or only to say that the mathematical equations that we use to reflect them are. The governing laws of fluid dynamics are continuous, but that is only because we have fudged the individual particles into the notion of a continuous medium. The laws of motion are continuous, but again that is only because macroscopic objects are stable, identifiable entities, quite remote from the ghostly effects of cloud chambers. It is all relative to the scale one is looking at. But it also relative to one's point of view and interests. Anyone involved in applied mathematics and the numerical analysis of partial differential equations will at once incline towards the continuity of the laws of nature. What would be the point of his discretising efforts otherwise? A theoretician, on the other hand, will be more critical of the mathematical model, and will really look for reductions. He will go to great lengths to establish that continuous phenomenological laws as we know them, for instance the laws of thermodynamics, emerge as average laws from the discrete population models of statistical physics, or that (continuous) Brownian motion is but the observable effect of a million hidden accidents and discrete encounters. Pushing this line of thinking to the ultimate grounds, he will soon be back to a vision of the world as quantum mechanical, that is to say essentially discrete.

This is not the end of the story, however. Much as the higher level physical laws, such as the equations of fluid dynamics, or Newtonian motion, or thermodynamics, were continuous on no other account than the macroscopic approximation they issued from, and did not satisfy us, for that matter, as being the real, ultimate, laws of the physical universe, the one equation governing quantum mechanics is in fact exactly the opposite. The Schrödinger equation is at the same time genuinely continuous and ultimately valid. What looks discrete in quantum mechanics, and has earned this theory an undeserved reputation in this respect, are just the observable facts: screen spots, energy jumps, decay events. But then these facts are discrete and isolated just as any facts are supposed to be. Only they do not pertain to recognisable objects or their properties anymore, so to enable us to say that some particle has travelled from source to screen and did not follow a traceable trajectory, or that some other has jumped between two energy levels and did not go trough intermediary values. For the claimed discrete character of quantum laws makes sense only against a continuous background. It presupposes that the bearer of the discrete properties is itself stable and continuously identifiable. Discrete facts and events, as recorded in quantum laboratories, pertain only to quantum measurement or to quantum experimental set-ups, and the only "physical" object that there is left in quantum mechanics - the only ontology that may substantialise lawful evolution - is the wave function (Bitbol, 1996). And this strictly obeys the continuous Schrödinger equation.

So it seems we are no better off than when we started, trying to determine whether the laws of nature are at bottom discrete or continuous. Either they look continuous, but then we find they are not genuine and ultimate, or they are really continuous, but then our whole notion of a physical object is disrupted. A game theorist I am used to meeting recurrently at my favourite Parisian café has said to me once that the infinite and the continuous are just manners of speech. They are our best approximation of the finite and the discrete. And this brings me really to the topic I wish to discuss in this paper. Numerical analysis has traditionally been viewed as a lesser form of mathematics, standing no comparison with the beauty and eternity of pure mathematical theory. It strikes the young aspiring student as a poor practice reserved to the mortals among us who have humbly accepted their limitation. Truncation error has currency in this dark land instead of perfect and impeccable proof. Surely enough, the detractors tend to forget that numerical analysis has its own body of pure theory and doctrine. Still, the whole terrestrial endeavour of numerical analysts is to try to get as close as they can to the theoretical solution of the PDE. The continuous PDE still shines as the perfect and inaccessible model above the plethora of numerical schemes they devise everyday to trick it into an approximate solution. I want to argue against this Manichean view. What I wish to show in this paper is that the discrete numerical scheme is no more an inferior or a subordinate to the PDE than the other way round. Truly, the continuous limit can itself sometimes appear as a proxy for an otherwise irreducibly discrete situation, and the best way to look at a numerical scheme is not always as the discretisation of something eminently continuous and eminently right. It will appear that efficiency and robustness - and these are objective values that are more urgent to my mind than the metaphysical question we started off with - can sometimes be gained by recasting some theoretical argument in discrete terms through and through, whereas its more natural derivation in the continuous realm would only add unnecessary trouble and interference.

Interestingly enough, it is in a field which is man's own creation that the discrete and the continuous will get the equal share of respectability we were talking about. This probably gives us the right leeway in answer to our question about the discrete or continuous character of the laws of nature, for the matter will be completely out of nature's hands, and no ultimate, mind-independent reality will be held accountable for the final answer. The interplay of the discrete and the continuous in the field of Finance will be my topic, and I will look into the PDEs and numerical discretisation schemes specific to the new "physics" known today as "Quantitative Finance." For one thing, "Continuous-time Finance" has been a hard-earned title, that may have produced indeed some very continuous equations, but it hasn't distracted us from the belief that Finance is in point of fact utterly discrete. Underlying stock trades in reality at discrete points of time and value, and traders rebalance their positions in reality in a discrete fashion. Transaction costs is an authentic financial problem, hinging upon the fact that transactions are flesh and blood operations, hence is not best described as a market "imperfection." To put the picture into more relief, maybe the best thing to do is to compare the natural sciences. The only cost in trying to make the continuous models of the natural sciences take into account friction effects, dispersion effects, or any other imperfections that were neglected at first, is that some equation may turn non linear. But that is not a killer. The solving technique will just have to be a bit more involved. Trying to inject transaction costs in the continuous Black-Scholes model, on the other hand, will make it blow up. So it seems the discretised version of a financial PDE can have another purpose in life than converging to its continuous limit. The discrete scheme can take a right turn and try to support transaction costs for a while. Or it can settle for a time step that reflects exactly the re-hedging frequency of the volatility arbitrageur and a space step no smaller than the traded tick. In a word, the numerical scheme can forget about theory and converged solutions for a while, and try to reflect reality for a change! Now, think what a disruption of the usual methodological link between continuous PDE and numerical scheme this means. Discretising a continuous financial model now appears to be no longer the same as financially interpreting a discrete scheme. For there are two things the time step earns us when it shrinks to zero in the Black-Scholes argument, and they are quite different in methodological stature: the continuous PDE of course - whose theoretical solution our numerical scheme is after - but more importantly, market completeness! In other words, truncation error is not the only thing to change when you change the mesh size. For all you know, the whole supporting theory can suddenly shift if you take the time step anywhere above zero! It is okay to think of a finite difference grid as a numerical means to a theoretical end, and the finite difference solution will no doubt converge to Black-Scholes if the scheme is consistent, but do not try to blow in it a financial life of its own, for the meaning of option prices will then require from you, right then and right there, a construct far more elaborate than Black-Scholes: an option pricing theory in incomplete markets! Now who wants to draw this big a picture, only to let it "converge" later to the uneventful simplicity of complete markets? To repeat, there is more going on between the discrete and the continuous in Finance than just numerical analysis.

So let us leave the continuous Black-Scholes PDE for a while and invest the discrete numerical schemes with an interest of their own. There was something fishy in the way the Black-Scholes argument proceeded to the continuous limit anyway, and it wouldn't be the first time that the Black-Scholes model strikes us as different in character from the models of the physical sciences (Ayache, 2001). The way the physical PDEs are usually derived, say in the case of fluid dynamics, is that a certain infinitesimal region is considered over a certain infinitesimal period of time, and the laws of conservation are written for this "control volume." Then the limit is taken, and the continuous equation is obtained. Taking limits is a mathematical operation quite separate from the underlying physical theory. Bracketing the metaphysical question of the continuous or discrete character of the laws of physics, we can even say that the continuous mathematical limit serves no other purpose than elegance of formulation. To quote from the introduction of Karatzas and Shreve's most recent book, Methods of Mathematical Finance: "The continuity of the time parameter and the accompanying capacity for continuous trading permit an elegance of formulation and analysis not unlike that obtained when passing form difference to differential equations." Surely the authors must have had in mind the elegance of the physical PDEs, and we certainly approve of financial models formulated as elegantly. But let us not rush the issue as far as mathematical, or "physical," finance is concerned, and let us go back to straight physical theory. Elegant as it may be, the step to the continuous limit certainly needs formal grounding in its own right, even in the case of traditional physics. But this has always been an independent matter. Physicists have traditionally left it to pure mathematicians or logicians to come back and "rigorize" the mathematical tools they were using, long after their physical insight and vivacity had blazed the trail, and the physical theories that came out of it had encountered the wide success that they did. See how Bolzano rigorized the calculus and finally gave precise meaning to Leibniz's "ghosts of departed quantities" or Newton's "nascent increments," and how Laurent Schwartz finally spelled out the exact mathematical nature of Dirac's function. What I mean to say at any rate is that physical theory is itself "continuous" through the process of jumping to the limit, because it is independent from it. Should the formal tool prove to be inadequate, or the mathematical intuition of the physicist not to be the most rigorous one, then a new formal tool would be needed. But the physical theory would remain unscathed. Things are different with Black-Scholes, however. We do need to hit the limit in the Black-Scholes argument, and we do need continuous Brownian motion, because it is the only way the market can become complete, and the theory can find its foundation. To put it in martingale parlance, general stochastic variables can be represented as stochastic integrals - or in other words, general payoffs can be replicated by the proceeds of self-financing dynamic strategies - only if the underlying process converges to continuous Brownian motion. A discrete Black-Scholes is Black-Scholes no more. Or if you insist, let us really try to find a discrete numerical scheme that not only discretises the Black-Scholes PDE but the whole physical argument as well! The answer of course is Cox's binomial tree. And the reason it is binomial is not that it is a numerical scheme so simple as to be understood by traders all round or by financial quantitative analysts freshly introduced to the field of numerical analysis. The binomial tree is binomial - and so persistent in the financial textbooks - because two states of the world are the only way we can regain market completeness in a discrete setting. Seldom has a physical model enjoyed recognition as wide as Black-Scholes, yet depended on a coincidence as improbable as Brownian motion or a binomial world. Think by contrast that the control volume of fluid dynamics need not be cubical, hexagonal or spherical in any way. Since we believe above all in the law of conservation of philosophical interest, it is all pointing to the idea that the significance of Black-Scholes must lie elsewhere, and to that we will turn later. What is needed in the meantime for Continuous-time Finance to really establish itself as a physics worthy of that name, is that it frees itself from the gripping coincidence lying at its heart. For instance, financial instrument prices are averages over states of the world, not pinpoint values, and moments of probability distributions are the truly observable statistics, not the distributions themselves. Therefore what is needed for Finance to find its own principles of symmetry and invariance, and for its laws to be covariant with changes of "co-ordinates" the same way the laws of physics are, is that the pricing theory no longer be dependent on particular states of the world, or for that matter on any particular structure of uncertainty such as Brownian motion. Prices themselves should on the contrary be the state variables (so that new "states of the world" are actually added when new instruments are introduced, such as options), and arbitrage relations should be retained only when they result from the description of the assets and their payoffs (what Philippe Henrotte has called intrinsic arbitrage relations (Henrotte, 1996)), not when "they depend critically on the description of the states of the world ," let alone their number. Philippe Henrotte has made this issue plain in his technical paper, and I understand that he will soon be writing - hopefully in the present series - about its methodological consequences for Finance.

Going back to our investigation of the methodological link between the continuous and the discrete in Finance, we are now in a better position to map it out. As we said, the binomial is the most "comprehensive" numerical scheme: it is the discrete version both of the pricing PDE and the physical argument supporting it. (That it should be numerically the least efficient could be a direct consequence of this, due to a methodological trade-off imposed on us by the physical indigence of Black-Scholes theory - but let us not press this point further). The trinomial goes one step further towards separating the two issues. It is numerically more flexible, but is not reducible to the complete market framework anymore. It retains a "realist" outlook, however, in that it starts out directly as a discretisation of the underlying stochastic process. Hence, it still speaks the language of states of the world, of transition probabilities, and of expected values. Now the fact that trinomial probabilities may sometimes turn negative under node manipulation, and that the financial engineers should not mind it so long as the numerical solution is right, is to our mind the first indication of a total change of picture. For we are here witnessing the slow and timid shift of the financial engineer towards what should have struck him, in the first place, as the separate issue of numerical analysis and discretisation of the Black-Scholes PDE proper, had his imagination not been captured by the financial interpretation of the binomial tree for so long. Wilmott et al. (1998) is probably the first to have comprehensively formulated the derivative pricing problem in terms of PDEs, and paved the way for their independent numerical treatment. Cultural mutation is usually slow to come about, however, especially when the people's numerical scheme, the binomial, is so well entrenched in the ideology itself. Under the increasing demand for numerical efficiency and the increasing complexity of the derivative instruments (barrier options, convertible bonds, etc.), the combination of mathematical craft and loyalty to the tree could not but beget the most baffling hybrids. The Brownian bridge adjustment of tree probabilities, allowing for the possibility that the underlying hits a given barrier between nodes, is not an uncommon trick nowadays. But what, if we may ask, becomes of the notion of states of the world in a mixture such as this? What becomes of the original conception of the tree as the discrete proxy of the stochastic process? And if worse comes to the worse, and the tree adjuster urges us that his only business in this unnatural mating of the discrete and the continuous is numerical efficiency, then we would be happy to take him seriously. There is virtually no limit to the variety and ingenuity of the numerical schemes one can deploy to crack a PDE. So why not interpret what our man is doing as just the particular choice of one? We would be much obliged if our bridge engineer could abandon the tree picture altogether – where the barrier subsists in a continuous limbo and the (discrete) world is not enough – and just make his numerical scheme manifest.

Two paragraphs back, we suggested that continuous PDEs were not the last word in financial modelling, and that some insight could be independently gained from a rehabilitation of the discrete scheme as such. Are we now saying that PDEs are the only way of rigorously grounding the pricing problem? Are we implying that trees should recede into natural history, and that unless you begin by clearly writing the PDE you wish to solve, your whole pricing enterprise will be fraught with trouble and ambiguity? The answer is clearly yes, and it will finally bring us to the true, methodological role that we think the continuous PDE should play in the pricing problems, and to the real significance of the Black-Scholes model, that we said must lie somewhere else than expected. What we like about the Black-Scholes PDE is its "physical" flavour, and the way in which the argument leading to its derivation looks like the theoretical arguments leading to traditional PDEs in physics. What we do not like about it is the Black-Scholes argument itself when taken seriously, and its physical pretence. Let me elaborate the bright side now, having amply discussed the dark one. It's been our experience that every time some particular pricing complication (stemming either from the derivative instrument itself or from the financial model we wished to apply to it) had made it unclear how we work out the maths exactly, it is the unhesitating regress to the origin of the PDE, i.e. Itô's lemma, the continuous-hedge argument, and the no-arbitrage argument, which has helped straighten out the difficulty. In pricing the cross-currency convertible bond, for instance, one is tempted to take as coefficients for the first partial derivatives in the two dimensional PDE, the foreign interest rate and the difference between domestic and foreign interest rate respectively for the stock component and the exchange rate component. A quick look at the way Wilmott deals with the quanto pricing problem (Wilmott, 1998) however, clearly proves this intuition to be wrong. "We could write down the differential equation directly, writes Wilmott, assuming that the underlyings satisfy lognormal random walks with correlation r. But we will build up the problem from first principles to demonstrate what hedging must take place." Wilmott then goes on to construct a portfolio consisting of the quanto in question, hedged with the foreign currency and the foreign stock. He then rehearses the entire continuous hedge argument, taking care that the Itô differential is correctly written for the composite portfolio when the latter is measured in the domestic currency, and that the holding in the foreign currency earns the appropriate foreign rate. He ends up with the following PDE (we refer the reader to Wilmott's book for the detailed derivation):

$$\frac{\partial V}{\partial t}+\frac{1}{2}\sigma_X^2X^2\frac{\partial^V}{\partial X^2}+\rho\sigma_X\sigma_SXS\frac{\partial^2V}{\partial S\partial X}+\frac{1}{2}\sigma_S^2S^2\frac{\partial^2V}{\partial S^2}$$

$$+X\frac{\partial V}{\partial X}(r-r_j)+S\frac{\partial V}{\partial S}(r_j-\rho\sigma_X\sigma_S)-rV$$

=0

where $V(S,X,t)$ is the derivative instrument, X the exchange rate, and S the foreign stock. Notice the unexpected addition to the convection term in the stock direction.

I am not saying that recasting the pricing problem in terms of a PDE is the only way to get round a difficulty as subtle and insidious as the transformation of drift across currencies. As a matter of fact, Hull and White (1994) offer an alternative argument, based on how risk-neutrality, as viewed by the foreign investor, transforms into risk-neutrality as viewed by the domestic investor, and on a implicit reference to the "work of Cox, Ingersoll, and Ross (1985) and others" (whose main bearing on the problem in hand, mind you, is the abysmal "market price of risk"), to show that the drift of the foreign asset transforms the same way as Wilmott has predicted. Hull and White's derivation is no doubt as rigorous as Wilmott's. What I am saying is that unless you are aware of the profundity of risk-neutral pricing and the correlative change of probability measure, there is a good chance you will miss the drift adjustment when building your tree in practice. The problem with martingale theory and its numerical translation into Monte-Carlo or trees, is that theory and mathematical algorithm are separate, hence increase your chances of losing the thread of your argument. The tree, remember, is built initially for the underlying process. Pricing of the derivative instrument is then superposed on it, and appeal is made to risk-neutrality and what not, to enable you to use the same drifts and probabilities . The PDE, by contrast, is the pricing algorithm of the derivative instrument right from the start. If you take the trouble of deriving it seriously every time a new pricing problem arises, if, that is, you carefully go through all the stages of the physical argument, as is done in physics, then the derivation will lead you, as a matter of course, to the right coefficients . Complex or cross-currency pricing problems are not, for that matter, the only case where the way of the PDE is the most beneficial course of action. Speaking of risk-neutrality, it even turns out that pondering the Black-Scholes PDE is the best way to dispel the confusion that surrounds this "misused, misnamed, and misunderstood" notion (to put it in Joseph Pimbley's words),even in the simplest pricing situations. In a talk he gave at the Risk 2001 conference in Paris , Pimbley dared to ask what was the meaning of the "risk-free rate" that is supposed to govern the growth of the underlying stock in risk-neutral pricing. "The theoretical PDE having shown that option value does not depend on the true stock [risky] appreciation rate," argues Pimbley, it looks as if the risk-free rate, as given by government Treasuries, or interbank lending rate, is the drift to be fed into the pricing algorithm, be it a tree, Monte-Carlo, or the PDE itself. On close examination of the Black-Scholes derivation, however, Pimbley finds that risk-neutrality (and by that we now mean the consequence of the preliminary continuous hedge argument) requires a funding rate, not a risk-free rate! This is the rate at which the trader is supposed in theory to fund his continuous hedge position, and this can of course vary, in reality, according to the trader's particular access to funding. While all this insight is readily packed inside the Black-Scholes PDE, risk-neutrality in Cox-Ross looks more like an external notion. "Cox-Ross discovered 'risk-neutrality', argues Pimbley, and argued [separately] that this principle applies only when a hedge for the derivative trade exists." Given, "by argument," that the composite continuously hedged portfolio should grow at a deterministic rate, and that the derivative instrument should grow (by martingale theory, martingale representation theorem, market completeness, and what have you) at the risk-free rate, then the most "convenient" rate for the remaining component of the portfolio to grow at, says Pimbley, is the risk-free rate too. Risk-neutrality strikes us as a mathematical trick in Cox-Ross, and that is only because the algorithm makes it look as if the whole idea were to model the evolution of the underlying stock and as if the tree were our only end.

So far we have seen two ways in which the Black-Scholes PDE was methodologically operative. One consisted in taking up the complex and making it look simpler (the cross-currency problem), and the other in taking up what looked simple and cutting with it very deep (Pimbley's rebellion). Soon, we will be diverging from the Black-Scholes PDE strictly speaking, as we try to figure out how to incorporate default risk in the pricing problem (convertible bond pricing, typically), only we will retain its main lesson, that of "building up the problem from first principles" in order to see what exact equation follows. But we find that the Black-Scholes equation has also opened up a whole new direction of inquiry, that of investigating the exact meaning of its own ingredients, and we realise that this new domain of investigation – let us call it the philosophy behind, or beyond, Black-Scholes – is far from being settled, however complete and settled the Black-Scholes PDE may be in mathematical form. This, the penetrating criticism of Pimbley has already shown. The further philosophical idea we would like to touch upon, however, is again how different from the PDEs of the physical sciences it all makes Black-Scholes's. Surely enough, philosophers of science are still debating today the meaning of Schrödinger's equation, or even Maxwell's equations, but hardly would you find anybody questioning their parameters the same way Pimbley has questioned the risk-free rate and suggested we used the funding rate instead. The reason why our methodological doubt and philosophical astonishment could very well but increase as we go with the Black-Scholes PDE, may be that it depends, in the end, on human contingencies and practicalities to a degree unparalleled by its physical cousins. The speed of light in continuous mediums is not a matter left for us to determine the same way we ought to determine every trader's funding rate and transaction costs. Perhaps the most fascinating thing about the Black-Scholes equation is how unquestionably we've been using it all this time, and how endlessly we could question it if we wished, down to each ingredient. Peter Carr has once pointed to another "massive detail" that usually goes unnoticed in the mainstream derivation of the Black-Scholes PDE. Once the continuous hedge portfolio has been written as:

$$\Pi = C + \Delta S$$

thus setting the stage for the mathematical differentiation, we should be very careful, argues Carr, not to call the next step of the derivation "the expression of the *differential* of the portfolio." What we write are rather the infinitesimal *gains* of the portfolio:

$$\delta\Pi = dC + \Delta dS$$

In this expression, dC and dS are the Itô differentials of the contingent claim and the underlying stock alright, but the crucial point here is that no such differential was written for the amount Δ of stock we are holding. Indeed, stochastic differentiation would have produced the following different expression:

$$d\Pi = dC + \Delta dS + d\Delta S$$

The reason why Δ is held constant over dt is not some first order approximation, for it is here that the self-financing constraint is crucially condensed. Therefore, we must not go about "mechanically" with the Black-Scholes argument, trusting, as it were, the rules of mathematical differentiation to do the job, or leaving it *a priori* to the laws of nature to determine whatever "nascent increments" should follow, while our task would only reduce to writing the mathematical transcript. The Black-Scholes PDE is no physical PDE. This, we have stressed enough. But it is not a PDE from the economical sciences either, as our reference to a human element of choice, either in transaction costs or in self-financing strategies, may have suggested. There certainly exists an interpretation of the Black-Scholes model from the overarching point of view of general equilibrium theory. Our immediate point, however, is that as far as option science and option *pricing* theory are concerned (and it was precisely Black and Scholes's great achievement to have earned option pricing its independence from investors' preferences and utility functions), people have been using the Black-Scholes formula as if its authors had discovered, in 1973, a new law of physics – a matter that we definitely resent –, but it is equally our concern, on the other hand, that what relates specifically to "human affairs" in the Black-Scholes model (the fact that the theoretical risk-free rate should be equal to the actual funding rate, or that the infinitesimal gain should not be identified with the mathematical differential) be given real *physical* prominence and no longer necessitate the penetration of a Pimbley or of a Carr to attract our attention. Continuous-time Finance is usually perceived as a special application of stochastic calculus, and Itô certainly did not predict that his special way of defining an integral over the paths of Brownian motion was going to become such a universal tool in such a special science. No wonder that the self-financing constraint should look as an oddity against such a general mathematical background. We are often reminded that the creation of the mathematical tool can sometimes precede the physical theory for which it is tailor-made. Still, we do not believe the relation of Finance to Itô's integral to be as maximal and solemn as the relation of Quantum Mechanics to Hamiltonian formalism. Rather, we think that Finance owes its mathematical apparatus to a fortunate coincidence, and that its debt is now long overdue. What is needed for Finance to possess a calculus of its own is on the contrary that the self-financing condition be recognised as its mathematical trademark, and that a stochastic integral be defined afresh with this condition built-in. This is the task Philippe Henrotte has set out to achieve, and we can already refer the reader to what he calls "a stochastic integral for Finance" (Henrotte, 2001).

Other than unravelling the cross-currency problem (a positive, productive task) or ushering us into the deep criticism of its own building blocks (a "morbid," philosophical task), what is the real significance of the Black-Scholes PDE? Given the vast scientific space that Finance is waiting to conquer in its own name then to fully colonise, either physical (symmetry principles and invariance) or mathematical (a tailor-made calculus), you may wonder how a tight rope such as Black-Scholes has supported so much traffic for so long. As we said, its significance must lie elsewhere. Its great contribution to the future history of sciences is to have provided the first lead to that space. Black, Scholes and Merton have truly invented rational option pricing. Brownian motion didn't fool anybody as a realistic model of stock evolution, and econometricians knew from the start that their job was different to the derivative "cracker's." For objectivity is a different matter to metaphysical truth, and our best scientific theories make no claim about reality in itself. Brownian motion may be a very improbable basis, and the Itô integral may be a very fortunate coincidence, the one lasting consequence of the Black-Scholes PDE is its scientific character (what he have called earlier its "physical flavour"). To our mind, its most valuable result is to have established a logical, mathematical link between some real, albeit very stylised, action the trader can perform on the underlying, an intuitive, albeit very hard to estimate, parameter (volatility), and the value of the option. In a word, option value has become a distinct object whose claim to independence and subsistence is all the better supported that it enters, thanks to Black-Scholes and Merton, into a definite functional relation with the other variables. Perhaps an extreme example can best illustrate my point. What better evidence of the objectivity of the Black-Scholes analysis, and of its disconnection from the question of metaphysical truth, than the fact that it may extend to a case where it is blatantly know not to be true? Take the case of default risk for instance. Convertible bonds present us with the perfect occasion to link the Black and Scholes world with a notion completely foreign to it, more readily imported from the traditional field of fixed-income and corporate debt, the notion of credit spread. The science of convertible bond pricing has apparently not yet detached itself entirely from the old fixed-income framework, and not yet fully embraced modern option pricing theory as pioneered by Black-Scholes, for most of the convertible pricing models widely used today still mix theoretic and heuristic insight, and convertible bond "manuals" such as Connolly's are quite disappointing when it comes to exploring credit risk (to our mind, the one big pricing problem concerning convertibles), or indeed viewing the convertible bond as something less simplistic than the combination of a bond and a call option or a tree with special clauses.

Indeed, convertible bond quantitative analysts have learnt this much from Black-Scholes that equity should be "discounted" at the risk-free rate in the pricing tree, and have remembered this much from their fixed income classes that a corporate bond should be discounted at the risky rate. Hence they find it quite appropriate to split the convertible bond into an equity component and a bond component, roll back each component separately in the tree, making sure, when early conversion is optimal at a node, that the equity component goes to full equity and the risky bond component goes to zero. In those upper regions of the tree where the convertible is most likely to be converted, so the argument goes, it is only fair that it should behave like equity, and in the lower regions where it is more likely to mature as pure debt, it is fair enough that a credit spread should apply to it. The mixed influence of these two different behaviours is supposed to propagate back to the present spot and to the present day through the simple mechanics of the tree probabilities, and it is this way, argue our analysts, that this great mixture of credit exposure and volatility exposure that we call the convertible bond gets reflected in the present price. To our mind, such a pricing methodology is among the worst abuses one can commit in the field, and we will once again blame the tree for it, together with its probabilistic intuition. Of course equity should be discounted risk-free (Pimbley permitting) and risky bond risk-bound. But what has all this got to do with convertible bond pricing? A converted convertible bond behaves like equity alright, but then it is a convertible no longer, and should at once leave the computational scene! And an unconverted convertible bond behaves like risky bond alright, but that is only because it still is a bond, and subject to default risk as well as the bond! The statement we are making here is that the convertible bond should, under no circumstance, be split into two components. Mortality does not split man into a living being and a dead person. (Only immortality does, into body and soul, but then it is not a contingent feature like the conversion of the convertible bond). Early conversion, early redemption, or even default are different events happening to the one and same derivative instrument. So we must look for a single evolution law ruling its existence. Information flowing into the present price of the convertible bond should admit as only source the infinitesimal “control volume” surrounding it, and not the conflicting dialogue of two separate components that owe their whole existence to the remotest regions of the tree. What I am trying to make felt here is the need for a PDE that would really establish the functional relation between the value of the convertible and the possibility of default in a whole and consistent manner, regardless of any forward guess about the fate of the instrument.

Assuming that the default event is triggered by a Poisson process whose intensity p may depend on time and share, and following Wilmott, we can model the infinitesimal fate of the convertible bond in this fashion:

- with probability $1-pdt$, no default occurs, hence the usual Black-Scholes continuous hedge argument applies, our hedged portfolio (convertible and stock) is immune to market risk, hence grows at a deterministic rate:

$$\delta\Pi = \left(\frac{\partial C}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2}\right) dt$$

- with probability pdt , default occurs, and changes in the portfolio value are swamped by the loss of the defaulted part:

$$\delta\Pi = -X$$

Now, we need to take expectations under probability of default (something we can always do), AND we want the expected change in the portfolio to equal the risk-free rate, in order to get the same right hand side as in the Black-Scholes PDE. The problem here is that the portfolio is not completely immune to risk because we cannot hedge against default. Hence, the usual no-arbitrage argument does not apply. Unless we reformulate the entire problem in the incomplete market framework, where real probability of default and real probability distribution of the underlying stock are needed (not mentioning the revision of the whole pricing theory), the only course of action that is open to us is a leap of faith, right in the middle of the risk-neutral world. The intensity of the default Poisson process has then to be given in that world, and an argument from Merton (1976), also cited by Wilmott (1998, p.330), to the effect that “the risk in the discontinuity should not be priced into the option when the jump [or the default] is uncorrelated with the market as a whole,” or in other words, that “diversifiable risk should not be rewarded,” allows us to write:

$$E(\delta\Pi) = r\Pi dt$$

Neglecting second order terms, this lead to the following PDE:

$$\frac{\partial C}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} + rS\frac{\partial C}{\partial S} = rC + pX$$

All we need now is fill in the defaultable component X , and we’re done.

The most conservative hypothesis is to let the convertible default absolutely like a common risky bond:

$$X = C$$

This corresponds to the usual tree technique mnemonically known as “Grow risk-free, discount risky.” The model can be further worked out to allow for the possibility of partial recovery. In which case:

$$X = C - \alpha C \text{ or } X = C - \alpha B$$

where α is the recovery rate and B the naked bond, depending on whether the holder recovers part of the convertible or part of the coupon stream.

In the exchangeable bond case, the holder has the opportunity to convert on the spot in case of default, hence:

$$X = C - RS$$

where R is the conversion ratio.

We have suggested that the latter model be used even in the case of the convertible bond, when no juridical issues prevent the holder from converting in case of default. The PDE becomes:

$$\frac{\partial C}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} + rS\frac{\partial C}{\partial S} = rC + p(C - RS)$$

As the convertible bond is always greater that the conversion value, the last term in the RHS is always positive. Also note that this term will shrink to zero as the stock moves up and the convertible approaches the conversion region, alternatively that it will become very close to C as the stock collapses. The heuristic split into equity and bond finally finds a theoretical grounding. Since default is modelled as an instantaneous event, only instantaneous actions that you may take on the spot can affect the convertible price. The equity-like and the bond-like behaviour are seen here to emanate from your instantaneous capacity to convert, and not from the behaviour of the convertible in the long run.

The problem with a model like this, however, is that your “instantaneous capacity to convert” is most likely to be worthless in case of default, because the share is most likely to be worthless when default occurs. One convertible bond pricing model, recently publicised by a software company who is “still the leader in convertible analytics,” even goes so far as assuming that both the share and the bond jump to zero in case of default.

Now, do they mean that the share follows a jump-diffusion process, such as described by Wilmott (1998, p. 328) for instance:

$$dS = \mu S dt + \sigma S dX + (j-1) S dq$$

where dq (a Poisson process) equals 1 with probability λdt, in which case S jumps to value JS?

And do they further mean that $J = 0$ in our case, and that the Poisson jump process and the process triggering default are one and the same?

Such a model could be easily accommodated in the PDE framework provided the jump intensity is given. It would result in a PIDE (partial integro-differential equation) whose numerical implementation is a bit harder but perfectly manageable.

As a matter of fact, the still-leading software company suggests that default may occur only "from the lower parts of the tree," i.e. as soon as some critical threshold is reached. This also could be reflected in the previous PIDE trough the somewhat abrupt choice of jump-default intensity:

$$\lambda(S,t) = \delta(S_c(t))$$

and it would give us freedom to set the threshold level anywhere we like, why not infer it from some structural model of the issuing firm, or other.

On close examination, however, it appears that the threshold level $S_c(t)$ must be inferred from the tree of the underlying itself in such way as to match the credit spread curve. This is done by interpreting the threshold as a knock-out barrier on the issuer's zero coupons and by inferring it through forward induction. But it would mean abandoning completely the Poisson process model, for the Brownian probability of hitting a barrier would now act as the "intensity" of the default process.

This model is perfectly alright to price risky convertibles. As a matter of fact, Wilmott proposes one such (1998, p. 580). "We only need to add the condition

$$C(S_c,t) = 0$$

to our favorite (no credit risk) CB model," notes Wilmott.

But, if we may ask, how could such a model ever accommodate the fact that the share itself shall fall to zero on default? If Brownian, or tree, probabilities are all there is to default, then the share has definitely non zero probability to "diffuse" unharmed below the threshold level. While it perfectly makes sense to write:

$$C(S,t) = 0 \text{ for all } S \le S_c$$

as is usually done in barrier pricing, we can't make any sense of :

$$S = 0 \text{ for all } S \le S_c$$

The difference, in the Poisson case, was that we had zero chance of writing such a nonsense, because the jump, triggered at $S_c(t)$, would never allow it.

If I were to recap our predicament in a word, I would again blame it on the tree. Trees tempt us into a vision of derivative instruments as lotteries, where the state variable is the "state of the world": ω, the random draw from the sample space Ω. PDEs, by contrast, have the underlying stock itself as state variable. When somebody says that default causes *both* the share and the bond to go to zero, he must imply that the share and the bond are both *derivative* on an hidden, ultimate, state variable (this could be the value of the firm for instance, or, for that matter, the sanity of its CEO). But then let him write down explicitly the default process in terms of *that* variable, and let him convince us that his tree matches *that* process.

We said that the problem with the last PDE model we wrote was the question how on earth we can convert into a worthless share when default occurs. As the value of the share is our state variable (hence our ultimate cause), the only way to allow for a worthless share in case of default is… the other way round. In other words, the default intensity should itself be made dependent on the share, and should rise dramatically with a falling share. This, we think, is the only *rational* way to account for two different behaviours, equity-like and bond-like, of the convertible bond in computational space. And once the PDE is written, together with its boundary conditions, and properly solved, then it will become evident that the upper region where the convertible is converted into equity shall in no way influence its present price. For it will simply lie outside the free boundary corresponding to early conversion. (A clever change of variables can always turn a domain bounded by a free boundary into a rectangular domain (Zhy *et al*, 1997, Zhu and Yingjun, 1999).

We hope, with this tour of default risk models, to have further strengthen the case for the PDE methodology. Our tale of the discrete and the continuous will not be complete, however, until we go back to the discrete schemes, and say at last how "they can be invested with an interest of their own." Realistic Finance, when we left it, was utterly discrete. And it was so necessarily (if only because of transaction costs). Building a realistic discrete model would therefore burden us with the task of setting the time step to its real live value, not mentioning the whole incomplete market pricing theory we would need to put in place. Now, if such a fine-grained model were implemented, we would certainly want its outputs to be stable with regard to the "real" time step. In other words, we would not like another trader, should one have a re-hedging frequency slightly different from ours, to get prices wildly at variance with ours. We couldn't even claim to have got anything like a model if such stability were not the case: the object of our science would be completely missing. One way of guaranteeing this stability, or, to put it in philosophical terms, this objectivity, is to let our discrete scheme be the consistent, stable, numerical scheme, solving some given PDE. If our discrete scheme could converge to a continuous limit, and the corresponding PDE be written, then objectivity would be granted .

The words of my game theorist now strike us with a new meaning. The infinite and the continuous, he said, are often our best way of speaking of the finite and the discrete. This maxim we now understand as saying that the infinite is often our only way of speaking of the finite, for it is the only way to grant our finite speech an object. We left the laws of nature with the Schrödinger equation and with the feeling that they could be continuous at bottom if it was. But the Schrödinger equation pertained to a very peculiar object, the wave function, whose only "physical" function was to assign probabilities to (discrete) experimental outcomes. So it all looks as if Probability were the "cement of the universe," and pretty continuous at that. The problem with Probability, however, is that it is itself a big philosophical puzzle. Our best "objective" conceptions of Probability define it as a limiting frequency of occurrence in an infinite random sequence. The finite variant runs something like this: "The probability of A is the number p between 0 and 1, such that the frequency of occurrence of A in a finite sequence has a shrinking probability to differ from p by more than a shrinking e as the sequence grows larger and larger." As Popper (1959) says : "In this formulation the word 'probability' occurs twice." And the plot thickens when you think that the cases where Probability is most urgently needed are probably those involving an event that will happen only once: the so-called single-case probability. The solution of the puzzle comes from semantics and the philosophy of meaning. It cuts through all our little worries: "Never mind if the infinite sequence exists, will exist, or will never. Our language, either scientific or ordinary, just needs the notion of objective probability. So all we need is just reference to the infinite, not its actual existence." Sometimes, all we need is reference to the infinite to make our finite and discrete language objective.

Now the trader's real world is no doubt discrete, discontinuous, and choppy like a raging sea. Stocks jump all over the screen, and volatility (my God, volatility!) is nowhere near the Black-Scholes model. But this, the trader was trained for. When to re-balance his hedge and how often to change the volatility parameter in his pricing formulas, are matters left for him to judge. The wilder the market gets and the more agitated the trader, the more he will want his tools to offer him the best grip, and to react locally in the smoothest and most continuous manner. The trader has learnt the language of delta, of gamma, and of vega, and he is out to speak it in a very loud and unruly world. But for it to be a proper language, it has to react in certain proper ways. The trader has got to have "played" exhaustively with a pricing model whose outputs are perfectly regular with respect to the inputs, for him to master the language. The intuition of delta, gamma and vega that the analytical Black-Scholes formula has helped shaping in the trader's brain, is to our mind the one valuable thing that any pricing model should preserve. This is the other reason why we think a discrete numerical scheme should behave, in the trader's hands, as perfectly as an analytical solution, and why we do not want to hear anybody say to us anymore that they didn't care whether the numerical solution of our PDE was penny accurate.

By now, you should be convinced that the continuous (or reference thereto) can sometimes appear as an instrument, or an accessory, to the discrete, rather than the opposite. Discrete financial models, we saw, are sometimes best interpreted as the discrete numerical schemes consistent with a continuous PDE. However, let me end this paper with a twist: *numerical schemes, destined to solve a PDE, are sometimes best convergent when interpreted as discrete financial models*. Indeed we have noticed in practice, that if we slightly "perturbed" the numerical scheme so to price exactly some known basic instruments, such as the pure bond, or the underlying stock itself, then the numerical value of instruments we needed really to compute would converge better. Sometimes the particular pricing problem imposes on you a change of variable. This trick is very useful, in that case, to get round the transformation of drift and volatility parameters. All this trick does, when you think of it, is reinterpret the discrete grid as the exhaustive collection of states of the world prevailing for the basic instruments. It is a kind of variance reduction technique. The first time I've seen it used is in Andersen and Brotherton-Ratcliffe (1998). It has recently been given a name in the literature, "tradable schemes," by J.K. Hoogland, and C.D.D Neumann (2000). When asked, my numerical analysis colleagues from the physical sciences said they had never heard of such a technique. Perhaps the reason it imposed itself in Finance is the trader's holy impatience. That your numerical scheme may give him an answer with truncation error, when he plugs into it the payoff of a pure bond or of the stock itself, is something he cannot tolerate!

REFERENCES

- Andersen, L. B. G and Brotherton-Ratcliffe, R. (1998). The equity option volatility smile: an implicit finite-difference approach, Journal of Computational Finance 1 (2), Winter 1997/98 5-37.

- Ayache, E. (2001). A Good Smile is Vital, Futures and Options World (FOW) May 2001.

- Bitbol, M. (1996) Schrödinger's Philosophy of Quantum Mechanics, Boston Studies in the Philosophy of Science, 188, Kluwer Academic Publishers.

- Black, F and Scholes, M. (1973). The Pricing of Options and Corporate Liabilities, Journal of Political Economy, 81 (May-June), 673-54.

- Carr, P., Lipton, A. and Madan, D. (2000). The Reduction Method for Valuing Derivative Securities. Risk Conference on Numerical solution of PDEs in derivatives pricing, New York, June 2000.

- Connolly, K. B., Pricing Convertible Bonds, John Wiley & Sons 1998.

- Cox, J. C., Ross, S. and Rubinstein, M. (1979). Option pricing: a simplified approach. Journal of Financial Economics 7 229-263.

- Cox, J. C., Ingersoll, J. E. and Ross. S. A. (1985). An intertemporal general equilibrium model of asset prices, Econometrica, 53 363-84.

- Henrotte, P. (1996). Construction of a State Space of Interrelated Securities with an Application to Temporary Equilibrium Theory, Economic Theory 8, 423-459.

- Henrotte, P. (2001). A Stochastic Integral for Self-Financed Portfolios, Working Paper Groupe HEC.

- Hoogland, J. K. and Neumann, C. D. D. (2000). Tradable Schemes, Centruum voor Wiskunde en Informatica (CWI), Stichting Mathematisch Centrum (SMC), Report, MAS-R0024 September 3, 2000.

- Hull, J. and White, A. (1994). Numerical Procedures for Implementing Term Structure Models II: Two-factor models, Journal of Derivatives, 2 (Winter 1994) 37-48.

- Karatzas, I. and Shreve, S. E. (1999). Methods of Mathematical Finance, Applications of Mathematics, Stochastic Modelling and Applied Probability 39, Springer 1999.

- Merton, R. C. (1973). Theory of Rational Option Pricing, Bell Journal of Economics and Management Science, 4, Spring 1973, 141-83.

- Merton, R. C. (1976). Option Pricing when Underlying Stock Returns are Discontinuous, Journal of Financial Economics 3 125-44.

- Pimbley, J. Misuse of Risk Neutrality in Credit Derivative Pricing, 10 April 2001, Risk 2001 Europe, Paris.

- Popper, K. R. (1959), The Logic of Scientific Discovery, Hutchinson & Co (Publishers) ltd.

- Wilmott, P., Dewynne, J. and Howison S. (1993). Option Pricing, Mathematical models and computation, Oxford Financial Press.

- Wilmott, P. (1998). Derivatives, The Theory and Practice of Financial Engineering, John Wiley and Sons.

- Zhy, Y.-l., Ren, H., and Xu, H. (1997). Improved effectiveness evaluating American options by the singularity separating method. Working paper, University of North Carolina at Charlotte.

- Zhu, Y.-l. and Yingjun, S. (1999). The Singularity-separating Method for Two-factor convertible bonds. Journal of Computational Finance, 3 (1), 91-110.