Elie Ayache

God's Model vs. Market Models Part III: Man's Model

We recognize in Black–Scholes–Merton the 'model of a carpenter', and we seek ways how truthfully to generalize it.

Paul Cézanne, Public domain, via Wikimedia Commons

The defining characteristic of market models, if we recall, is that the vega-hedging instruments (vanilla options or forward variance contracts) trade in the market on an equal footing with the underlying asset. They admit of prices and not of values. They are no longer evaluated in the backward procedure of a probabilistic decision tree. Their maturity becomes irrelevant. Their prices no longer depend on unobservable volatility states (Heston models, or the like) but become themselves directly the 'volatility states.' We're not even sure we can use the label 'volatility' for them, for that notion relates to the underlying process, which, according to our reading of Bergomi, no longer exists (see Parts I and II). The market prices of the hedging instruments may be represented by their Black–Scholes–Merton (BSM) implied volatilities and it so happens that the price of a log contract (or a continuous strip of adequately weighted options) is directly assimilable to a volatility number (through BSM), but this is only a pricing convention, and it is probably the only remaining place where we will hear the word 'volatility.' Simply, the market prices of the hedging instruments are the state variables of the pricing function of the exotic option, and generally of any derivative instrument that the market model will price.

The exotic option in turn no longer admits of an expiration, and of the backward valuation procedure associated with it. The market, or the player known as Reality,[1] is such that the exotic option market price will always be given by this pricing function, with the prices of the hedging instruments as arguments. Even if we covertly approach the expiration date or the knock-out barrier of the exotic option, for all we know, the pricing function underlain by the prices of the vanilla options may still want to pull the exotic price in a different direction than the one about to be imposed by the final payoff. Who are we to reason with the market, and to impose on it the constraints of non-arbitrage? There are historically known cases of the market failing to 'converge' to the expectations of the arbitrage formula, resulting in dire consequences for the arbitrage hedge fund and even for the authors of the formula. Maybe the market expects the formula describing the payoff itself no longer to apply. Maybe events will take place, at that exact moment, which put in question all the received views of arbitrage and all the financial holy scriptures, including the written payoffs. If the trading decision underlying the market model is that the pricing function is of that particular form and admits of those particular state variables, which may not even be remotely linked to the underlying asset, then we must abide by this pricing function and trust it is the market's, who will always have the last word.[2]

I admit I may be exaggerating a little, for in practice you will always find that the pricing function is materialized by a stochastic model, which will enforce non-arbitrage and certainly convergence to the payoff function expressed with the asset price as sole underlying, but precisely the formal pricing function *is not* equal to that material model. The substance and the philosophy of the market models really take place at a distance from the boundaries of the market (expiry dates or knock-out barriers, options approaching intrinsic values, etc.). Far from those boundaries, the logic underlying the market models is, as we have argued in Parts I and II, neither ruled by arbitrage nor by probability. The essence of the market models is that the vega-hedging instruments should be independent assets. It is interesting to read Bergomi's own defense against the charge of arbitrage, at this juncture. Having realized that an arbitrage opportunity would emerge if the individual constituent options of the vanilla surface were to move totally independently of one another, he immediately corrects the view by assuming only a discrete finite set of such vanilla options and not the continuum thereof. [3] With the strikes and maturities of this discrete finite set being reasonably distant from each other, the idea is that the options prices might move a little, independently of one another, yet remain within the no-arbitrage bounds. Once again, the sacrosanct principle of non-arbitrage is recovered by the practical and always finite implementation of the market model. However, this, in my opinion, shouldn't hide the fact that market models formally escape the non-arbitrage principle, or rise above it, because they reflect the market's (infinite) pricing function and because the market is the superior (infinite) player who works indeed in mysterious ways and who will always make sure, beyond any probability structure or comprehension, that we don't get infinitely rich.

From market models to exchange models

We say that the market models *reflect* the market pricing function and do not produce it. The market pricing function is supposed to exist and always to exceed the particular market model. This reflection of the market by the model is thus always incomplete and always passive. The market, as we said, has already solved all the pricing problems. The vanilla options are priced and traded from time immemorial, and no pricing theory is needed for them apart from the market. As for the exotic option price, although seemingly produced by the market-maker, it is ultimately the result of the pricing function, which is the market's. The market-maker only makes the price as made by the market. He posts the price to the market only insofar as the market posts it back to him (see Part I). Thus, the market models are neither attempts at finding the true model of the market (leave that to Gatheral and to rough volatility[4]) nor instruments to create the market or stage its genesis: they don't place themselves in the inaugural position of BSM, when no options market existed.

We say that the market models reflect the market pricing function and do not produce it

We witnessed, on the contrary, an attempt by Bergomi to re-embed BSM in the logic of the market models, where the option price, produced by the BSM formula, was supposed to be given by the market already, and BSM supposed to be just an accounting equation (pp. 5–7). The market models are really tools for engineers, not instruments for creators. They reflect the engineer's modesty of having to deal with an existing reality (not to construct it) and with the best way to piece it together.[5] Bergomi finds himself in the middle of the existing vanilla ocean, and he insists that the Bergomi model for the pricing of exotics is only a temporary vessel, which will always fall short of the ultimate pricing function and of the ultimate exotic price. The exotic price is thus part of the same ocean. Nobody actively makes the market or wills the market in Bergomi. Henry-Labordère, extending the engineering spirit of Bergomi, speaks of *automated* option pricing. The market is a machine; it is automated; it is self-pricing. Henry-Labordère even dispenses with the representation of the underlying dynamics (as temporary as it may already be in Bergomi). He merely asks: "Given the current market price of the underlying asset and also given market option quotes at several given strikes and terms, provide option quotes at any strike and term [in particular, exotics] in a specified set."[6]

Although the trading decision has severed option pricing from probabilistic evaluation and the underlying stochastic structure, and has reinstated the market as a whole and integral surface (a surface such that the exotic market price is only a function of the vanilla market prices, with the label 'market' intentionally kept on both sides of the equation), it hasn't really introduced the trading force, or true trading. We kept saying that the implied volatilities of the hedging instruments were modeled directly, as resulting from the proper forces of the options market and no underlying process, however the options market was always assumed. We never really investigated its genesis. Recall the trouble Bergomi had in determining whether BSM belonged in the category of market models. This is because BSM is anterior to the trading decision, and the trading decision is the inaugural step of the market models. As for the pricing function, it is, as we said, the passing observation of an engineer who picks the parts and pieces of the problem on the surface. Such modesty is great, we all agree, however its net result is that the pricing function is gotten in a passive way, and is not tensed by genesis or characterized by force. The trading decision simply amounts to making the new price a function of another price. It is an immersion within an ocean of already existing variables and variations, and is not a positive and

>

forceful decision *to trade*. As a result, all the vanilla options constituting the implied volatility surface and all the forward variance contracts constituting the forward variance curve are supposed to trade in equal measure and on an equal footing with the underlying asset, without any structure constraining their trading ranges other than the one Bergomi will eventually pick, for the subsidiary purpose of numerical tractability. At the other extreme, you have people who view derivatives as mere abstractions: they can be multiplied without a limit and there are as many derivatives as there are payoff functions of the underlying asset price. For this reason, those people think derivatives do not really exist, so how could they be traded? The underlying asset exists, by contrast. It is traded with force, an elemental force. It is traded absolutely, and constant volatility, which comes first to our mind when we conceptualize its trading, is typical of the independence of its market.[7]

Shouldn't we similarly look for the 'real' variable that is trading behind the derivatives? Shouldn't the structure both constraining and correlating the pixel movements of the volatility surface, or of the forward variance curve, really be the indication that some deeper variables, hidden beneath the volatility surface or beneath the variance curve and more fundamental than the individual prices of their individual constituents, are the things really trading and really subject to the trading force, therefore to constant volatility? Surely, God is dead and an underlying process shouldn't exist, but perhaps an underlying structure is required nevertheless in order to furnish the human trader with the handles on which to apply the trading force, as opposed to an indifferent ocean of prices, in which the relation between inputs and outputs is just an automated pricing function. An underlying stochastic structure may not only be the signal of a probabilistic evaluation, therefore of divorce with the market; to the contrary, it may reveal the parameters that are truly trading and therefore promote the market and embrace it all the better! We may wish to retain what structure there is in BSM (i.e., both the underlying Brownian motion *and* its constant volatility), in order to declare that *volatility is being traded as a result*. It is because constant volatility makes options redundant in BSM, and denies them a market, that we know *exactly* what's being achieved, in BSM, when options are traded. (We don't know exactly what's being achieved, in a stochastic volatility model, when options are traded.)

Perhaps there is new meaning in volatility being exchanged, different from a quantitative change: a volatility that is somehow constant and not constant, a space of variation of a new kind

Observe the threshold effect: volatility is traded *because* it is recognized as the constant number (or the fixed structure) in BSM, yet its trading will entail that the number will change. According to this new view, it would be a mistake to press forward and express that change by making the BSM implied volatility stochastic, for that would result in as many stochastic processes as there are vanilla options, or the very thing we were complaining about, and an even bigger mistake to make the instantaneous volatility of the underlying asset stochastic. One should hold still at the threshold. The market-maker shows his implied volatility bid-and-ask spreads: what happens later, change or not change, is suspended. There is a new sense of variation to be found, here, for volatility. Definitely the explanation of the volatility smiles and of the volatility market does not lie in a model of stochastic instantaneous volatility (Heston, rough volatility); but perhaps it shouldn't consist, either, in modeling the implied volatilities directly, and the market models of Bergomi are perhaps rushing too quickly to the other extreme. Perhaps volatility shouldn't change, but be *ex-changed*, and instead of market models for volatility, we should be speaking of *exchange models*. Perhaps there is new meaning in volatility being exchanged, different from a quantitative change: a volatility that is somehow constant and not constant, a space of variation of a new kind.

In a confrontation with a famous econo-physicist who seemed to disagree at the time that BSM deserved the Nobel Prize for economics on grounds that volatility was constant in BSM and the model was not risky, Hélyette Geman writes: "The economy of BSM is *risky by definition* because it amounts to *exchanging* volatility. The fact that this risk should be materialized by a *single little number* makes it palpable and immediate for everybody."[8] In this, she expresses the threshold: a single little number (a constant, then), which is, for this very reason, palpable and liable to be exchanged. Geman certainly doesn't mean that BSM should be replaced by a stochastic volatility model (such as Heston); on the contrary, it should be kept, and the tension between constant volatility and exchanged volatility should be preserved![9] The rough volatility model would be great, and would deservedly be recognized as the new paradigm succeeding to BSM, if it turned out to be a tool for exchanging a single little number too, why not the key parameter H, or Hurst parameter. One of the conclusions of the rough volatility model is that the shape of the volatility surface is universal, and universally explained by no more than three parameters. This amounts to saying that the volatility surface has become the new compact object and should no longer be divided into the constituent vanilla options. This is indeed progress; however, I personally wouldn't recognize God or eternal truth in this universality, but precisely the dawn of a new market therefore of a new trading hell, with Hurst smiles now replacing the volatility smiles.[10]

Bergomi would never agree, of course, that the volatility surface is of a universal shape. How could it be when each constituent vanilla option is given total freedom of trading? To repeat, Bergomi has severed all possible links to an underlying process that might impose structure and shape or inspire a quest for the Holy Grail. As a matter of fact, he would be first to repudiate his own namesake model if the covariance structure between the hedging instruments turned out to be different one day. The inversion of the faith is such that the options market always comes first, and the underlying structure always comes later; the accounting equation comes first, and the probabilistic interpretation later. Yet, on closer scrutiny, it appears that there subsists at least one element of structure in this totally unstructured world, and that is the preliminary choice of European convex (or concave) payoffs as hedging instruments. Both the vanilla option and the power payoff (which is the basis of construction of the forward variance curve) admit of well-defined BSM implied volatilities, or a one-to-one mapping between their market price and the corresponding BSM implied volatility. This seems to suggest that their market first emerged as an exchange of the BSM volatility – the single little constant number, which provides the trading handle – exactly as voiced by Geman. Although there is no underlying process in Bergomi, and consequently no notion of an underlying asset volatility, the particular choice of the hedging instruments, which constitute the initial reality underneath which or before which we are told we must not look, seems to hint that, in the time before time, the market model was in fact an exchange model.

Volatility is exchanged in BSM because buying or selling hedged options is equivalent to a play on realized volatility. You make money or lose money depending on how your gamma gains or losses compare with the initial option premium, or how the realized volatility compares with the volatility that is implied by the premium.[11] Precisely, Bergomi introduces the vega-hedging instrument, in his proto-reasoning, in order to cancel the gamma of the total position (see Part II). He cancels the volatility play in BSM, and for this reason, it may seem that he doesn't exchange volatility, in the sense intended by Geman as the genesis of the volatility market. Bergomi skips the genesis stage, but this is because the volatility market already exists in Bergomi! Indeed, the trading of volatility that Geman intends is already all factored in by the vega-hedging instrument, which Bergomi can bring from *nowhere else but the existing options market*.[12] The decision to trade is from the start taken over by an instrument that already trades, or by the *trading decision*. The exchange model is from the start supplanted by the market model, in Bergomi, and the single little number, which we said must be kept constant in order to be ex-changed-and-not-changed, finds its correspondent in the convex (or concave) payoff of the vega-hedging instrument. The one volatility number is readily replaced by a one-to-one mapping. Bergomi shows us that gamma risk is from the start replaced by vega risk, or in other words, that the underlying process never really existed, but only the options market did. Options have always been hedged with options.

It remains to characterize the exchange models of volatility, or the models that strike the right balance between the vertical structure of the rough volatility model and the horizontal and undogmatic structure of the market models (or rather, their lack of structure). How can we stop at Geman's threshold and avoid considering a continuum of European payoffs, every single one of which trades in a separate dimension? Because of this infinite proliferation, the market models in fact block any possibility that another parameter than volatility is eventually exchanged. It is true that Bergomi (as per the vocation of the market models) will always be looking for the next volatility market (i.e., the volatility market of his own vega-hedging instruments); however, it is not clear how the particular instance of the Bergomi model could, without a major re-engineering, accommodate a spot volatility surface of VIX options, or a spot forward variance curve of VIX futures, in the same way as it did the vanilla options surface or the forward variance curve. Precisely, Bergomi had chosen the particular covariance structure known as the N-factor Bergomi model as a temporary fix. The result is a time and space-homogenous break-even covariance matrix for the vega-hedging instruments. This is the generalization of the BSM break-even volatility $\hat{\sigma}$, which we had chosen to be constant, at the time, and which Geman insists should be kept constant in order to be ex-changed-and-not-changed. But the question now becomes: How could Bergomi's break-even covariance matrix be exchanged in turn? In the same way as we had exchanged volatility by trading options or variance contracts with the help of BSM, must we not expect to exchange the covariance matrix by trading the exotic option with the help of the Bergomi model? We know what Bergomi's answer would be: The N volatility processes and the covariance matrix have no special significance (p. 225). If the price manifold of the exotic option starts assuming dimensions that the present covariance structure of the underliers cannot accommodate, then this is the signal that the temporary stochastic structure should be made less stringent and the pricing function expressed differently, mathematically, but this does not mean that the pricing function has essentially changed, or that the market has changed levels and ascended from a volatility market to a market of volatility of volatility. To begin with, the word 'volatility' has no meaning in the market models, because of the absence of underlying stochastic structure. There is only a function going from prices to prices. There is no variety of the market price of the exotic option and no variety of its variation that could not ultimately be the result of the same 'original' pricing function, admitting as arguments the (once and for all) given vega-hedging instruments, only now suitably rewritten and made more complex, as far as this written expression is concerned. The covariance matrix is totally arbitrary, because Bergomi's N-factor model is arbitrary, by his own admission. I doubt it will ever achieve the same widespread acceptance as the BSM volatility. The reason why

Volatility is exchanged in BSM because buying or selling hedged options is equivalent to a play on realized volatility

>

the BSM volatility is so special and is not arbitrary is that the gamma–theta play in which it enters is the basic play, *combined with the fact* that there exist convex (or concave) payoffs whose market price maps that volatility univocally, hence takes it to the market. It is not clear, by contrast, how a covariance matrix could 'go up or down,' or equivalently, which payoff could map it univocally. Or at least, this is our whole present discussion.

It's what's between the single little number!

We are all fascinated and arrested by the BSM volatility because it is a single little number. It is the reason why the convex (or concave) European payoffs are fundamental in the market models, and why it seems as if their market prices will remain the only possible state variables of the pricing function, to be succeeded by none other. Bergomi could not have started with barrier options as hedging instruments! We, for our part, believe the key to our problem lies in this single little number, provided we turn the key *in two totally unexpected directions.*

Let us first explore what single can mean. If we interpret the smile problem literally, we find that it doesn't involve several volatility numbers but a single number, only repeated. It doesn't involve many vanilla options at the same time, but the repetition and reconfiguration of a single vanilla option. When the break-even volatility $\hat{\sigma}$ of the BSM equation is put to the market and becomes implied volatility, or equivalently, when gamma risk is transmuted into vega risk and Bergomi resorts to the traded vanilla option in order to hedge the exotic option, this always concerns a single vanilla option, or a single convex (or concave) European payoff (pp. 15–18). The reasoning that stages the genesis of the options market concerns only one option at a time. Each single European payoff is put to trading in the market that is being created, and in which the corresponding implied volatility is exchanged. This completely separates each option from the next (no separation can be more complete because of the moment of genesis), but the problem now becomes that options, by this very fact, find themselves all immersed inside the one single reality that is being precisely created and was unconceivable before, namely the reality of the one single market. BSM creates the option market for each single option separately; yet, when they all become traded, we realize that this has now to occur in reality and no longer in the BSM genesis, and that, in reality, there is only one market. Options all sprung into that reality from BSM; that reality was constituted anew, therefore there was literally nothing, at the back of that reality, to explain the simultaneity of the options prices. The underlying stochastic volatility process which we later invent in order to account for the simultaneous prices of the vanilla options never existed. The reality, composed of the repetition of a single conceptual BSM and a single market, and later of the superposition of those instances, never became a causal reality in which time and space are shared by all and can be sliced into different locations so as to locate the different maturity dates and strike prices of the vanilla options, or equivalently a causal reality in which an underlying stochastic process can explain all the options prices at once. The genesis of reality (of the market) is separate for each option and should be indexed differently by each option; in other words, BSM knows only of one option, whose maturity T and strike K are just parameters, and it creates the market that is indexed by this option; however, once we land in that reality, we recognize it is unique and shared by all options and no longer indexed by each option separately. We inherit the problem that options now have to trade together in the same market: (T,K) are no longer parameters and become variables. This is the knot sitting at the heart of the smile problem. Options shouldn't trade together; yet they must. The solution of the smile problem is not to depose BSM; it is to superpose BSM.

Contrary to what we all think, market-makers do not trade options because volatility is stochastic or uncertain

BSM is an extraordinary formula because when fundamental value no longer existed and nothing existed but price and the random fluctuations of price, delivered to the wildest speculations of the trading crowd and to a stock exchange which, in the words of Bachelier, only speculatively "reacted to itself,"[13] this formula was able to derive the impeccable value of an option, and to derive it from nothing else but the total criticism of value, which is the volatility of the price! Indeed, what had become fundamental, in the market, after fundamental value was abolished and there subsisted only price, was the volatility of price. Volatility had become the 'fundamental value' of the market, or the fundamental concept summarizing it. And now BSM deduced the option value from that fundamental concept. Not only was their derivation impeccable; it was unassailable. For how could you criticize option valuation when it stemmed from no other concept than the very criticism of value? From this, it becomes unimaginable that the vanilla option should trade on the same level, or in the same reality, as its underlying asset. The option value was derived from the meaning of price and the meaning of the market (of the underlying asset), which is volatility. It wasn't derived on the side of price, or in competition with price. On the contrary, it rose above it. Option trading is thus unimaginable on the same trading floor as the underlying asset. Yet this is exactly what happened. Precisely because of the impeccable and unassailable derivation of option value, a god, a market-maker, descended from the conceptual heaven situated above the trading pit and wrote the option and named its price in the pit. Contrary to what we all think, market-makers do not trade options because volatility is stochastic or uncertain; they *make* options markets because volatility is constant and certain – constant and certain in the heavenly or conceptual sense, not in the quantitative sense in which something constant is about to vary and become nonconstant or something certain is about to become uncertain. It is not that options are redundant in BSM and trade in reality because BSM is obviously wrong and volatility is not constant. Rather, they trade in reality because BSM is right, even perfectly right.

BSM is sublime in providing a value, when everybody thought value no longer existed (not the 'when' of temporal coincidence, but the 'when' of opposition and even of paradoxical causation). This deserved an event, even a revolution. The revolution is that the option came back to trade on the trading floor, all the more forcefully. The major revolution or major change is not that, from constant, volatility is going to change. It is that the option's sublime and unassailable value is going to become an option's price. The fundamental and conceptual volatility, which stood for the summary and the meaning of the market of the underlying asset, is going to become implied volatility of the overlying asset. This is a qualitative, not a quantitative, change. We can even commit the radical step, and argue that implied volatility, rather than the asset's volatility, is the first incarnation of conceptual volatility. The volatility we had recognized as the fundamental value of the market was a quality, as such unquantified. One way of quantifying it is by the asset's volatility. This is the path chosen by econometricians and econo-physicists. But another way of first putting a number on the conceptual and unquantified volatility is through the option price, or implied volatility. BSM is sublime and never really cared about the time series of the underlying asset price. It chose a constant number

We can even commit the radical step, and argue that implied volatility, rather than the asset's volatility, is the first incarnation of conceptual volatility

for its volatility because that was the only sensible choice, a purely formal choice. However, when the time came to descend on the pit, in reality, BSM descended via the option. It is radically that we must now understand that the underlying process doesn't exist.

We think the BSM volatility is a single little number, but we misinterpret 'single' and 'number.' We think that the only way volatility can evolve, beyond BSM, is either by becoming multiple and no longer remaining single, or by changing numerically, as is normally expected from a number. The only possible evolutions beyond BSM seem to us either to feature as many volatility numbers as we have new tradable assets (Bergomi and market models), or to feature a stochastic volatility process in which volatility changes quantitatively (Heston, rough volatility). Yet we wish to uncover a different meaning of being single, which will allow the volatility number to remain single beyond BSM, and a different meaning of being a constant number which will allow volatility to be ex-changed-and-not-changed (i.e., to remain constant). These are the two unexpected directions in which to turn the key. What we are driving at is that the 'number' that will succeed to the BSM volatility should be multiple, of course, for multiplicity has entered into the problem whether we like it or not, yet should be considered single insofar as the algebraic operations in which a single number usually enters are concerned. *Hint*: A matrix is a multiple, yet it can be considered a single number when it lends itself to the same algebra as a single number (addition, multiplication, division, etc.). We are thus contemplating a *matrix of volatility numbers*, which is neither a covariance matrix relative to the vanilla options or forward variance contracts in their infinite number, nor a set of parameters of a stochastic volatility process such as Heston or rough volatility. What we need is a matrix corresponding to the superposition of BSM.

A covariance matrix cannot go up or down, like a single number, nor can it be mapped into a single payoff. That the BSM volatility was ex-changed did not mean, by the way, that it went up or down or changed numerically (or at least, not immediately). The threshold we had stopped at, following Geman, was exactly the following: *Because* the BSM volatility is constant, perfect replication is possible and the market-maker can write the option and present it

A covariance matrix cannot go up or down, like a single number, nor can it be mapped into a single payoff

to the market; but this does not mean that this volatility will change, because the reality of the market we will find ourselves in is revolutionary, therefore different from the initial one. Besides, the BSM volatility cannot change, for otherwise there would be no BSM and no market-maker, and no market to begin with. Recall that Geman doesn't see risk in BSM because of the stochastic volatility model that will eventually succeed to it. She insists that BSM is risky as such: "It is risky *by definition*." Therefore, the result for volatility is not a different number; it is the tension, it is the threshold. The result is an unusual algebra for the volatility number, and an unusual play between the one and the multiple. A number about to become multiple, yet that should remain single; a number about to change, yet that should remain constant. The key is the unsettlement. It is not "what's between the numbers," as the character Maximilian Cohen once exclaimed:[14] *it is what's between the single little number!* BSM is, in any case, very difficult to understand because it belongs to the genesis phase, or to the Big Bang. We cannot really reconcile it with an options market; yet it is itself the original formula that produces the options market (as witness Bergomi's difficulty in making it a market model). So a good idea may be to take a leap of faith and reveal without delay the structure of our *exchange model*, or the model that we believe is the truthful generalization of BSM, and later to return to BSM and argue, in a limit reasoning, that BSM is typical of the exchange models we have thus revealed, or a degenerate case of those exchange models, just as Bergomi has argued that BSM is typical of the market models (p. 5).

>

The regime-switching model

The exchange model is a regime-switching model: a structure where, as will become apparent, everything changes while nothing changes. The regime-switching structure is a discrete-state stochastic jump-diffusion process when seen from the outside – so it may look unimpressive to quants dealing daily with non-Markovian continuous roughened forward variance models – but its main virtues, when seen from the inside, are the *associativity* and the *genericity* of the operation of adding new regimes to the structure. The underlying asset price diffuses with constant volatility in a given regime, and can suddenly jump to a different regime through a Poisson process. The intensity of the Poisson process and the magnitude of the jump that the asset price then undergoes are contingent upon the particular transition from the initial regime to the target regime. After the jump, the asset price resumes a diffusion with a constant (and different) volatility in the target regime. Regime switches thus correlate jumps of the diffusion volatility with jumps of the underlying asset price. There is a fixed number of regimes, and switching between them continues indefinitely and is characterized by a finite transition matrix. *Associativity* is the fact that a given subset of regimes can be considered as a single 'big' regime, and switching inside the larger set can be viewed as occurring between such big regimes. There is thus already a favorable play between the single and the multiple. A regime-switching structure is indistinguishable from its own stochasticization, and it is never determined whether adding an Nth regime, say of volatility, is just expanding the current stochastic volatility model or making the volatility of volatility stochastic in turn.

This is exactly the feature we require in order to keep the unsettlement. Let's imagine, indeed, that we're already in the middle of the market, way past the original and genetic BSM, already in the generic market situation where derivatives have been long trading, typically vanilla options and forward variance contracts (or VIX futures). Our N-regime-switching model features more than one regime of course (otherwise it would be BSM). Diffusion volatility is constant in the regime we're in, so for anyone unaware of the regime transitions, it all looks as if we are in the BSM world of the single trading pit and the single vanilla option facing the underlying asset. But options are many, of course, and they are aware of the regimes, so in the presence of a full vanilla surface and VIX curve, the diffusion volatilities in the regimes, as well as the intensities and jump sizes of the regime-switching Poisson processes, are calibrated to both. Our model is thus instantiated as a stochastic volatility model (with jumps). Note, however, that because of the discrete character of the regimes, there is no explicit number for the volatility of volatility or for its correlation with the underlying asset price. This is an important point, perhaps the most important one. For it is not at once visible whether the volatility of volatility and its correlation with the asset price are constant or themselves stochastic.

The exchange model is a regime-switching model: a structure where, as will become apparent, everything changes while nothing changes

In any case, by solving a stochastic control problem, we know how to trade the vanilla payoffs in order to replicate perfectly any exotic payoff. Replication has to be dynamic, even for variance swaps, because of the presence of asset price jumps. VIX options are thus perfectly dynamically replicated. To speak the language of the BSM genesis, they are the 'new vanilla options' that express the 'fundamental value' of the vanilla options market (i.e., the volatility of the vanilla prices). Thanks to perfect replication with the vanillas, the market-maker is capable of writing the VIX options and offering them to trading, thus triggering an exit from this 'generalized' BSM world to the market – just as writing the first vanilla option, thanks to the original BSM and perfect replication with the underlying asset, and offering that option to the market, had ex-changed the conceptual and constant volatility and changed it into an implied volatility. The generalized BSM model is multiple-regime, of course, but it can be considered as single and simple and constant as the original BSM, once it is made a subset of the next model, following our remark above, and the objects are reassigned appropriately.

Indeed, two things will happen the next day. The vanilla options surface (and the forward variance or VIX futures curve) will be different of course, and recalibrating our N regimes will yield different parameters. On the other hand, chances are that the prices of VIX options that the market-maker had manufactured the previous day with the help of the 'generalized BSM model' will deviate from the prediction of the model, by sheer trading, thus creating a 'volatility smile.' It all looks as if our N-regime model has become stochastic, first, on account of its parameters stochastically changing the next day, and second, on account of the 'vanilla options' it was supposed to value – in this case, the VIX options – starting to exhibit 'volatility smiles.' It all looks as if the 'true' model now consists in switching between the N-regime model of day 1 and the different N-regime model of day 2. But switching between regime-switching models can only yield a new, larger, regime-switching model, by the associativity of regime-switching, so we end up with a $2N$-regime model. If modeling the underlying asset price process (or 'truth') was our objective, we wouldn't know what to do next. But since the options market is all that we care about (similarly to Bergomi), the next crucial observation is that the objects that have become tradable the next day (i.e., the set now composed of the VIX options in addition to the vanilla options – and the forward variance contracts, and the VIX futures), may not require a full $2N$-regime model to have their market prices explained, but may find a perfect fit in a 'single' N-regime model, with parameters of course different from the first two.

The associativity of regime-switching, combined with the fact that what we care about are the implied volatilities and not the asset's instantaneous volatility, makes it so that the 'single little' implied volatility number (the *N*-regime matrix) need not become multiple in order to be ex-changed. As we said, it is undecided and unsettled whether the *N*-regime model is just a model of stochastic asset volatility, which corresponds to a 'constant' implied volatility of volatility once calibrated to the vanilla and not to the VIX options, or a model of stochastic implied volatility of volatility once calibrated to the VIX options on top. The 'volatility number,' or *N*-regime matrix, need not become multiple and deal with each option separately, because it fits them all at once, and so it remains 'single.' Also, it remains 'constant,' as it is similar to its stochasticization, and does not require writing a new stochastic process. The key, therefore, is not to be found between two *N*-regime matrices. It is what's 'between' the single matrix, or the single little number. This suggests a new algebra for volatility, and we wish now to understand the 'single little volatility number' of the original BSM model in its light. Moreover, the structure is generic, because the same logic applies when a volatility index is later created for VIX options, and options are written on it in turn. These two crucial properties of the regime-switching model, the unsettlement and the genericity (which is its correlate), would be impossible, when thinking about it, if the structure wasn't discrete and the parameters populating it weren't constant.[15] This is how the puzzle is solved of whether to make the asset's volatility stochastic or the implied volatilities stochastic in order to address the smile problem. The answer is neither, or both.

Indeed, our exchange model stands exactly between God's model and the market model

Indeed, our exchange model stands exactly between God's model and the market model. It is *man's model*, because man is single. God either gives everything or takes everything. He doesn't exchange. As for the automated or machinic vision of the market that is afforded by the market model and its pricing function, it doesn't exchange either. It exchanges only data, input and output, and its trading decision is made, as we have said, in time immemorial, or in the time before time. It is always already past the threshold point at which the single and the multiple coincide. It doesn't reproduce the trading event, and keep reproducing it, like the regime-switching model. The exchange is not a number; it is a bid-and-ask spread, or the void between the number. Only a *situated* and finite being, who is neither God nor machine, can stand in that void and endure the event of trading. God knows truth, so He cannot trade. And the machine doesn't know the *situation*, so it cannot exchange.

ENDNOTES

1. See 'God's Model vs. Market Models' (part I), *Wilmott* Sept, 34–47, and Shafer, G. and Vovk, V. (2001). *Probability and Finance: It Is Only a Game!* New York: John Wiley & Sons Ltd.
2. See 'God's Model vs. Market Models' (part II), *Wilmott* Nov, 42–48, for the two fundamental concepts of *trading decision* and *pricing function*.
3. Bergomi, L. 2016. Stochastic Volatility Modeling. Boca Raton, FL: CRC Press, p. 67. All subsequent citations from Bergomi will be identified in the text by their page number in this edition.
4. Gatheral: "I believe that rough volatility is the true model" (*Risk*, April 2021, p. 31).
5. Bergomi writes in his preface: "Modeling in finance is an engineering field: while our task as engineers is to frame problems in mathematical terms and solve them using sophisticated machinery whenever necessary, the problems themselves originate in the form of embarrassingly practical trading questions" (p. xv).
6. Henry-Labordère, P. (2011). 'Automated option pricing: Numerical methods.' Available at https://ssrn.com/abstract=1968344 or http://dx.doi.org/10.2139/ssrn.1968344.
7. When BSM postulate that the volatility of the underlying asset price is constant, they don't believe it is really constant; they just choose it to be so, for what other sensible choice do they have? This is similar to Bergomi: "In the unhedged case we were free to choose the implied volatility as our best estimate of future realized volatility and kept it constant throughout" (p. 16).
8. Geman, H. (1997). 'From Bachelier to Black–Scholes–Merton,' *Bulletin français d'actuariat*, 1(2), 55–64.
9. "An option trader knowing the ins and outs of the BSM formula can beat a trader using a state-of-the-art stochastic volatility model," writes Haug (Haug, E. [2003] 'Know your weapon,' *Wilmott*, May, 49–57).
10. A hint: Gatheral assumes *H* is constant in the rough volatility model, while recognizing, from his own time series, that it is pretty much stochastic.
11. The log contract offers the additional advantage that the hedge ratio doesn't depend on volatility.
12. Between the break-even volatility, which we can always choose as constant, and the vanilla option implied volatility, which can never be constant, there is, thus, no intermediary step. Outside the BSM formula, there is no place for options except the options market. Compared to statistics and the intellectual leeway it always offers, the market is unforgiving.
13. Bachelier, L. (2006). *Theory of Speculation: The Origins of Modern Finance*. Princeton, NJ: Princeton University Press (translated and with commentary by Mark Davis and Alison Etheridge). When he speaks, in the first page of his book, of the 'infinite number of factors' that determine the movements of the Exchange, Bachelier attempts the first analysis of market microstructure, whose conclusion is Brownian motion, later to become geometric Brownian motion and give BSM, later to become Heston, later to become rough Heston, through the continuation of the same analysis and fascination with the underlying process. However, the infinity that interests me most is the one found as soon as we turn Bachelier's first page. He writes: "There are two sorts of forward-dated transactions: forwards, options. These transactions can be combined ad infinitum, especially as one often treats several kinds of option." This is the inception of market models, according to me, where infinity concerns the options market and no longer the underlying process, and I've always been impressed that Bachelier and Bergomi have the same initials.
14. *Pi* (1998). Movie directed by Darren Aronofsky.
15. Time and space-inhomogenous models are thus out of the question in the algebra of exchange models.