Elie Ayache

God's Model vs. Market Models

A thorough reading of Bergomi sheds new perspectives on rough volatility in relation to the meaning of the options market.

Francisco Bayeu y Subías, Public domain, via Wikimedia Commons

According to a recent declaration by Gatheral, God's model is finally in our reach.[1] The ultimate explanation of option prices is about to reveal itself. There is one model, the *rough volatility model*, which finally brings together all the pieces: market microstructure, the time series of the underlying price volatility and the surface of vanilla options prices. First, market microstructure is described by a tick-by-tick model in the framework of the Hawkes process. Then some assumptions are made, which seem very reasonable in the modern setting of high-frequency trading: endogeneity of the market, absence of statistical arbitrage, asymmetry in the liquidity on the bid and ask sides of the order book, presence of meta-orders due to algorithmic trading which splits trades in time. Following Omar El Euch, these assumptions lead (asymptotically) to a rough volatility model of the log price. Empirical studies of the time series of price volatility confirm, on the other hand, its rough behavior. Finally, to crown it all, prices of vanilla options on equity indices are found to be amazingly well explained by a rough volatility model for the underlying process. In the words of El Euch: "It is really true that rough volatility models are amazingly consistent with both historical and implied volatility data."[2] This is meant to resonate with a statement first made by Gatheral: "For perhaps the first time, we have a simple consistent model of historical and implied volatility."[3] Both in the historical measure and the risk-neutral measure, the rough volatility model seems to be the unique, simple, and true model. Hence the name, 'God's model.'

Notice that the deduction of the rough volatility model (or the proof of the existence of God) starts with the underlying process then concerns itself exclusively with it. The theoretical part of the proof is the mathematical derivation of the rough volatility model from the starting point of a tick-by-tick process of the underlying price. As for the two empirical confirmations, they also hinge on the underlying process, the first by looking at its historical volatility, and the second by calibrating its parameters from the implied volatility surface. This is in total opposition with the following inaugural statement by Bergomi: "The mistake – done in master's degrees in quant finance – is to start from the assumption of a stochastic process for the underlying, a thing we're not even sure exists."[4]

Since the rough volatility model, or God's model, is essentially a model of the underlying process, Bergomi's declaration amounts to saying that we're not sure God exists. As a matter of fact, the only thing that should exist, according to Bergomi, are *market models*. Those he defines in the very opening section of his book (section 1.1 – 'Characterizing a usable model') as models that are *usable* in the market, where the market is in turn pragmatically defined as a situation where the only thing that counts is controlling the profit and loss (P&L) of the hedged position (or the portfolio composed of the derivative instrument of interest and the traded instruments that were selected for hedging), regardless of what goes underneath or what's truly underlying from God's point of view. Pragmatics against metaphysics. "A pricing equation is essentially an analytical accounting device," according to Bergomi. Derivative practitioners are "content with barely floating safely and making as few assumptions as possible about future market conditions."[5] This means that they do not look below the surface or care about the underlying process or its stochastic structure, and even less so about what this structure may imply for the future and therefore impose on it. Their accounting equation has as *equal* entries the market price of the derivative instrument they are dynamically hedging and the market prices of all the hedging instruments, and the underlying asset is merely one among them.

When the underlying asset is the only hedging instrument, this doesn't really distinguish it either. Bergomi derives the Black–Scholes–Merton (BSM) equation without assuming that the underlying price follows any stochastic process, or in other words without looking beneath the surface. Remember that we're not even sure such a thing exists. This may sound surprising when BSM is usually understood as a consequence of stochastic calculus and of the Ito differential of the hedged position, in which the term of second order in dS is precisely identified with $\sigma^2 dt$ – an intermediary result one can hardly count on without the assumption of a stochastic process and the corresponding Ito calculus. Precisely, Bergomi inverts the arrow of causality. The consequence of the mistake done in master's degrees is that students are at a loss for ideas when the underlying process is mistakenly thought by them to be the cause of everything and when they realize it is no longer Brownian motion. They worry: "What if my underlier is not lognormal? What if it's not a diffusive process with constant parameters? Do I throw BSM away?"[6] Bergomi's whole endeavor is to block these kinds of questions. He shows they should use the BSM equation regardless of the BSM model, where the model is precisely the assumption of a stochastic process of this nature. When he says the underlying process doesn't exist, he is not concerned with the metaphysical debate about the existence of objective probability or of a random generator, or about the availability in reality of the conditions that define a stochastic process according to Kolmogorov. The empirical time series of prices of the underlying asset certainly exists, and recently some fine estimations of instantaneous volatility have made it sound as if the empirical times series of the asset volatility also exists. Certainly, any quantitative researcher is welcome to try to fit these time series with a rough volatility model, and in doing so, he certainly (and rightly) can claim that the underlying process exists. So Bergomi cannot mean inexistence in that sense. I interpret him as saying that the underlying process should no longer exist in the question of the student of *derivative pricing*. Derivative pricing – and especially the market models thereof – occur, as we shall see, in a register where the underlying process is no longer the question, and in this sense, no longer exists. It no longer exists because of the derivative price, not the underlying price. So, what I find questionable in Gatheral's God's argument is not so much the analysis of market microstructure of the underlying asset or of the time series of the instantaneous volatility of its market price, but the relation of all this with the prices of options written on it. The options market should always be the revolution and even the heresy, according to me, or the very thing worthy of philosophical interest in finance, and never be identified with a sacred and final scripture.

Post-BSM world

The paradigmatic derivation of BSM took place before the existence of an options market, and so the underlying process was all that existed. Bergomi, by contrast, speaks from a *post-BSM world*, and by that we mean that options markets already exist in that world and have even matured (probably encouraged by the very usage of BSM). As a matter of fact, the crucial supposition in all of Bergomi's approach and in all of the market models is that options markets are the given. Of course, one can independently argue that options markets also existed historically before the advent of the BSM formula (Bachelier, Thorpe, etc.), and there is a nice sociological account by MacKenzie to the effect that the formula has altered the behavior of the *existing* market-makers and in this way has shaped its world.[7] What we mean by saying that the BSM paradigmatic derivation (as opposed to Bergomi's derivation of the BSM equation) took place 'before' the existence of an options market is that the options market is not an assumption of that derivation, contrary to Bergomi's and the market models. The only driver in the paradigmatic BSM is the underlying stochastic process – for instance, in the rigorization later established by Harrison and Pliska, only the stock and the riskless bond are considered to be exchanged – and what BSM achieve is the derivation of an *arbitrage-free valuation* of options and not their pricing (in the sense of the establishment and, later, the total eruption of an options market).

As a matter of fact, Harrison and Pliska explicitly call 'verbiage' any reference to a parallel options market and to options prices which one would compare with the arbitrage-free values. The options market cannot even be a consequence of BSM, since all that the investors who follow the dynamic replication strategy can achieve is, in the words of Harrison and Pliska, "to manufacture call options for themselves."[8] People think (if we may digress a little) that BSM cannot produce an options market because options are redundant in BSM. But the reality is even worse. Options simply do not exist in the BSM framework, neither as an assumption, nor as a consequence. They are not even written or named, so that we may call them redundant. The formalism does not know how to make or produce a market other than the one that is formally postulated at the beginning. To make and to produce are material processes that are foreign to the nature of the formalism. Market-makers armed with the BSM model have produced the options market thanks to the technology of the trading pits and to their material existence as traders in those pits. The technology is precisely the material process that exceeds the formal model. So strictly speaking the trader is part of the technology.

The empirical time series of prices of the underlying asset certainly exists, and recently some fine estimations of instantaneous volatility have made it sound as if the empirical times series of the asset volatility also exists

The existence of implied volatility smiles (or the net result of options traders using BSM) is to be explained by a technology rather than a theory, and that's why Bergomi's book, whose agenda is essentially the modeling of stochastic implied volatility and not of stochastic instantaneous assets' volatility, is, to our mind, a book of technology rather than a book of theory.

Following the 'form' of the formalism, anything that trades is supposed to admit of the corresponding price process from the start. You don't make it; it can only be given. Thus, the stock and the bond, in the BSM formalism. For options to admit a market in such a framework, they need first to be identified as exchangeable assets and to have a stochastic process written for their prices from the start, independently of the underlying asset and (of course) eventually correlated with it, a matter that is of course beyond BSM, and would even be ironic in BSM. This is exactly what Bergomi does, when he moves beyond BSM and systematically starts by giving the stochastic processes for the price of the underlying asset as well as of the hedging instruments (i.e., the vanilla options or later the forward variance contracts). This is precisely the characteristic of the market models. But let us end the digression and go back to the paradigmatic BSM.

The paradigmatic derivation of BSM starts with the underlying process and ignores the existence of an options market. It doesn't know what it means to trade options. What it is trying to achieve is to compute the initial premium such that, by trading subsequently the riskless bond and the stock in a dynamic and self-financing way, a certain contingent payoff is perfectly synthesized at a certain maturity. This is a result that has only to do with the underlying asset and its market. You can give me money that you draw from your bank in exchange for the stock I hold, and this is how we exchange the stock in the stock market at the spot price. But you can also give me money and let me hold it and manage it alone for a while, by engaging, on my side, in a trading strategy in which I dynamically deposit and draw money from my private bank, and dynamically buy and sell fractions of the underlying stock from and to other people in the market, in order that I hand you the stock at a later (predetermined) date and not on the spot, and at a (predetermined) price that is possibly different from the spot price. But this doesn't mean that the two differences we've just marked with the spot price (the time difference, due to the delayed delivery, and the space difference, due to the delivery at a different price or strike price) have been marked and written as a contract and that we have exchanged that contract. For we have certainly not discussed, at this primitive stage of the formalism, any legal issue, for instance the possibility of default on my part, and even less so considered that you might independently engage in exchanging that contract with other parties.

The paradigmatic derivation of BSM took place before the existence of an options market, and so the underlying process was all that existed

Alternatively, one can take a step outside of the trading pit and look at the stock market from outside, as a mere generator of random prices on which one wishes to evaluate lotteries, whose payoff may be quite complex for that matter. In this case, the lotteries are indeed well identified; however, we're no longer inside a market. The only constraint for the bookmaker, in that case, is that he may not be vulnerable to a Dutch book argument, or that the odds he offers be coherent. He must be invulnerable to instant arbitrage. When it is later observed that at least one lottery (the one yielding the future price of the stock as outcome) is already evaluated by the stock market and that its present value is equal to the current stock price (and I insist on saying it is evaluated and not that it is being priced, because we are now situated outside the market), and when it is later required that the bookmaker be also invulnerable to statistical arbitrage, the constraint of no arbitrage results in the risk-neutral valuation of the rest of lotteries aka derivatives, in the equivalent measure, under the underlying stochastic process. In any case, whether we are looking at dynamic replication of contingent payoffs in BSM or at risk-neutral pricing which generalizes the valuation of derivative lotteries to more general underlying processes, the initial premium or the arbitrage-free value of the derivative is supposed to exist *as soon as the underlying process exists*, and all we are trying to do is find it (or select it within a family of equivalent martingale measures when the underlying process is incomplete and there are several such measures).

Thus, BSM may assume that option value will be of the form $V(S,t)$, before finding what this form might be and what the formula might be exactly (they will derive it); however, their main implicit assumption is that this value is gotten by application of the non-arbitrage principle. This is why they are able to equate the differential value of the arbitrage (or hedge) portfolio with riskless growth under the interest rate. And how did they manage to form a riskless hedge portfolio? By application of Ito's lemma which equates the second-order term in dS with $\sigma^2 dt$, where σ is the real volatility of the underlying process. Hence the required assumption of a process, and even better, the assumption that this process is known and its volatility is known. Arbitrage theory and arbitrage-free valuation in finance in any case rely on underlying states of the world, hence on the assumption of an underlying process (that's why we said that the arbitrage-free value of the derivative exists as soon as the underlying process exists). You may not need the BSM argument of dynamic replication of the option to price it without arbitrage, and a famous paper was written once against BSM and the exactness of replication, simply by applying put–call parity to formulas derived from actuarial pricing.[9] However, actuarial pricing is based on the real probability, hence on the assumption of a real underlying process. And if, on the other hand, you tried to argue that a formal probability measure is all you need in order to produce derivative prices that are arbitrage-free, and that this measure need not relate in any way to the real historical probability or to a real and existing underlying process (i.e., you could, for instance, price options using BSM with some fixed volatility number, regardless of whether the real underlying process has stochastic volatility or jumps), you would be wrong, for arbitrage requires that your pricing operator, or formal probability measure, or risk-neutral probability, be equivalent to the real probability (i.e., agree with it on all events of measure zero); hence it cannot, on pain of statistical arbitrage, feature a constant formal volatility number when real volatility is stochastic. In a word, risk-neural probability pricing, or pricing by arbitrage, cannot dispense with the assumption of the underlying process, and with the assumption that this underlying process is known to the writer's mind.

Be that as it may, what Bergomi calls 'post-modern finance,' or 'grown-up finance,' can very well dispense with arbitrage and arbitrage-free valuation altogether![10] Thus he can speak of using the BSM pricing formula and even derive it while not assuming the BSM model (i.e., the lognormal process of the underlying stock price). What he does, instead, is weaken this assumption and content himself with controlling the P&L of the hedged position, a task that is essentially ex-post (hence the inversion of the arrow). This may sound clever and very economical, but we must be aware that what Bergomi is doing here is nothing less than dispensing with the principle of non-arbitrage altogether and throwing overboard the whole of theoretical finance.[11] BSM assume $V(S,t)$ is of a certain form and that it is derived by a non-arbitrage argument, and they manage to derive it and show what form it has by applying non-arbitrage precisely at the crucial moment, as they have announced. This, we said, necessarily implies the existence of an underlying process. Any risk-neutral valuation, even not restricted to the BSM framework, requires the assumption of an underlying process of some kind. When the underlying process is lognormal (in the real probability), there is only one risk-neutral equivalent pricing measure. If not assuming the existence of a stochastic process, a minimum assumption for Bergomi, though, seems to be that the underlying price changes are not too wild (i.e., no large jumps are assumed) or, in his words, that "the model prices risks associated with small moves,"[12] for only then could he stop the Ito expansion of his pricing function $P(S,t)$ at second-order derivatives. Note that BSM do just that, formally. However, contrary to Bergomi, they believe in the existence of an underlying process, and even better, in the full disclosure of the existence and of the nature of that process to whoever's mind is writing their paper – what we may call the paper's mind, or the paper's consciousness. Why? Because they prepare to apply the non-arbitrage principle.

BSM make $V(S,t)$ stand before they show what it is; they say it is the arbitrage-free value of the derivative, then their derivation allows them to see what form it has. Where, by contrast, does Bergomi's pricing function $P(S,t)$ stand, before he shows what form it has exactly? If not an arbitrage-free value of the derivative, what is it? Precisely, it is symbolized by P, not V. It is not a value; it is a price. There is no principle of valuation in post-modern finance except given by the market. There are no options values but only options prices affected by supply and demand. BSM did not have an options market as hypothesis or ground. They were literally building option valuation from scratch. By contrast, the implicit assumption in Bergomi, as in all market models, is that options prices are given by the market to begin with. This may sound trivial, but again this is what grown-up finance is all about. We can dispense with the assumption of a stochastic process for the underlying, or with the idea of valuing derivatives by arbitrage, *simply because we assume their prices are already given and known*. Who are we, indeed, to discuss valuation and arbitrage in the presence of the market, after the market has ruled and proposed its prices?

Who are we, indeed, to discuss valuation and arbitrage in the presence of the market, after the market has ruled and proposed its prices?

The conceptual difficulty of market models

Let us note, though, that in the first chapter of his book (p. 2) Bergomi speaks of the 'bank quants' who handed us $P(S,t)$, and not of the market. It is only later, in stochastic volatility models where vanilla options or forward variance contracts become hedging instruments in their own right, that their prices will be (precisely) given by the market, that the market model will always be initialized with their implied volatility surface or implied forward variance curve, and (if we may anticipate our conclusion a little) that the price of the exotic option, delivered by the market model and by the pricing function Bergomi will associate with it, will be supposed to be the price *always already* given by the market for that exotic option. However, in the first chapter, when considering the underlying asset price alone as market given, there is no option market yet, as Bergomi very rightly reminds us (p. 5). We are not very far from the inaugural BSM situation, and the purpose is to say of what form the pricing function the quants have coded might be. The only distance that Bergomi takes with BSM at this point, and the only difference in the reasoning, reside in the inversion of the arrows. There is no underlying process or a priori principle of valuation, so there is no way we could prescribe a priori the value of the option, and the only alternative is to see a posteriori what reasonable price the bank could have quoted for the option, using the quants pricing function, *and could continue to quote over time* without going out of business and leaving the market. This will be achieved a posteriori by controlling the P&L.

The P&L is composed of increments of the option price due to variations of the two state variables S and t underlying the pricing function $P(S,t)$, as well as increments due to variations

>

of the price of the stock itself we are holding as hedge, as well as the cost of financing this holding. Precisely, the amount Δ of stock we are holding cancels first-order variations of the option pricing function, so there remains second-order terms dS^2. As no underlying diffusion process is given a priori, the latter cannot be equated with deterministic $\sigma^2 dt$, and therefore Bergomi is left with the stochastic term. The P&L is stochastic; it is not guaranteed to be (almost surely) equal to zero as in BSM. The price $P(S,t)$ was not constructed out of perfect replication. However, the crucial observation, what Bergomi describes as a "further reasonable requirement," is that the squares of incremental returns (daily in Bergomi's example) average out over time to their realized variance (p. 4). The hope, at the end of the holding period, is to *recognize* the lognormal historical volatility of the price of the stock S. No a priori assumption requires S to follow a given process. All we know is that it is random, and we can statistically compute its historical volatility. The need to resort to this temporary fix, as Bergomi expressly reminds us, is caused by the absence of an options market. He writes: "In the absence of a volatility market for S, $\hat{\sigma}$ should be chosen as our best estimate of future realized volatility" (p. 5). It may sound circular to allude to an options market when we are in the process of deriving the BSM pricing function supposed to yield the option price, but this attitude is revealing of the market models whose main assumption, we said, is that an options market already exists. In this regard, BSM is in a kind of twilight zone. Bergomi writes the differential of the BSM pricing function handed over by the mysterious quants, while supposing that the two successive option prices whose difference is expressed by the differential are market prices already. His purpose is to show us that BSM is an accounting equation, not a theoretical model relying on an underlying process, however he cannot escape the minimum assumption of a historical volatility 'somehow' relating to an underlying process (we will discuss later the extent of this 'somehow'). It is indeed difficult for a market model, whose assumption is that an options market exists, to be in a position where its assumption must be its consequence. That's why it remains undecided, in Bergomi's account, whether BSM is a market model. In one instance he claims that "the Black–Scholes model is typical of the market models considered in this book" (p. 5) and, in a later instance, that "the local volatility model is the simplest market model" (p. 72), a statement which seems to entail that BSM, which is simpler than the local volatility model, is not really a market model.

To clear the ambiguity, we prefer to stick with the definition of market models as *models that do not rely on an underlying stochastic structure but on the existing options markets instead* (hence the unease and unsettlement in the case of BSM), keeping crucially in mind that what is typical of those models is the ex-post accounting equation (ex-post, precisely because of the absence of structure and the corresponding a priori stance) that will be formed with the derivative (exotic) instrument we are trying to price and all the hedging instruments. Precisely, the market prices of the latter are required in order to exempt us from relying on estimates of volatility (such as $\hat{\sigma}$, in the BSM case) or any other statistical parameter which directly or indirectly relates to an underlying process. The astute reader will certainly protest that this only pushes the problem one level up. For the multidimensional P&L decomposition which we will, then, be able to write for the position composed of the exotic option and the hedging instruments will in turn require an assumption of a break-even covariance matrix of some kind, and the latter will in turn directly or indirectly relate to an underlying process ruling the prices of the hedging instruments. As a matter of fact, Bergomi reiterates the complaint about the absence of a volatility market for the instruments that we are presently considering as vega-hedging instruments (the vanilla options or equivalently the forward variance contracts), or in other words, a volatility market for the volatility instruments of first level. He writes: "If there existed a market of options on ξ^T with maturities ranging from t to T, the volatility risk of ξ^T could be hedged away and the volatility of ξ^T would be derived from market implied volatilities." Failing this feat, Bergomi confesses he has "no choice but to carry a position on the realized volatility of ξ^T and thus will need to make assumptions" (p. 220). This is equivalent to the situation, in the inaugural BSM derivation, where, in the absence of an options market, the assumption was made that realized volatility would be measured somehow and supposed to be the constant $\hat{\sigma}$.

Allow me to digress a little, once again. Are we saying that those assumptions Bergomi needs to make are assumptions of a stochastic structure of some kind (this is, crucially, where he introduces the famous exponential kernels which will characterize his model and get associated with his name, and later cause the criticism of the rough volatility supporters)? Not exactly. The existence of a break-even volatility $\hat{\sigma}$ in BSM did not mean the existence of a priori structure, but only the quite defensible belief that the square of returns of the underlying asset will average out at the end of the day to a constant number $\hat{\sigma}$. The situation with multiple hedging instruments is no different. The break-even volatility is replaced by a break-even covariance matrix C_{ij}. In the first paragraph in which he considers multiple hedging instruments (p. 5), it may

Bergomi writes the differential of the BSM pricing function handed over by the mysterious quants, while supposing that the two successive option prices whose difference is expressed by the differential are market prices already

even seem that the break-even covariance matrix he considers is constant. In a later paper, Bergomi says: "Ideally we would like to be able to set these break-even levels at will."[13] The reason why C_{ij} may seem constant is that, in this preliminary discussion of multiple hedging instruments, the underlying assets in question are still implicitly supposed to be basic assets and not derivatives. As a matter of fact, Bergomi still designates them by S_i, seeming only to generalize BSM and the case of a single underlying asset. However, when the hedging instruments become vega-hedging instruments proper, after a crucial proto-reasoning that I will discuss later (i.e., vanilla options or forward variance contracts which can only be derivative on the basic underlying asset whose volatility risk they are hedging us against), the notation differs and the break-even variances and covariances are no longer constant and become explicit functions of the state variables S and ξ, $v(t,u,u',S,\xi)$ (p. 218).

Even though the defining characteristic of the market models is that the vega-hedging instruments, as derivative as they may be in reality, must be considered as independent assets,[14] when it comes to expressing the covariance matrix of their implied volatilities (i.e., of their market prices), one cannot just assume that they are independent assets in the sense of the S_i above and therefore admit of a constant covariance matrix. The way an implied volatility surface, or a forward variance curve, moves (i.e., the covariance matrix of the constituent implied volatilities of the European options or of the forward variance contracts) has to depend on the 'volatility state' and may not be constant (i.e., Bergomi cannot talk for them of squares of returns that average out to constant numbers at the end of the day, and can now only consider local averages), where 'volatility state' refers to the volatility state variable(s) that we would be considering, if we were working in the traditional framework where a stochastic volatility process is assumed for the underlying, possibly with multiple dimensions, but which *we are not considering*, precisely because we are working with a market model in which implied volatilities are modeled directly and no assumption of an underlying process is made, thus leaving us no choice but to make the covariance matrix a function of the implied volatilities ξ themselves. And why does the motion of the implied volatility surface or of the forward variance curve depend on the covert 'volatility state'? Because the corresponding instruments are covertly derivative, hence their prices depend on a covert and unsayable stochastic structure.

Precisely, the whole conceptual difficulty of the market models lies in their desire (more than a desire, an agenda) to model the market prices of the vega-hedging instruments directly (i.e., as given by the market), without any suggestion that those prices, or equivalently the implied volatilities of vanilla options or of forward variance contracts or variance swaps, might be outputs of a model of the instantaneous volatility of the underlying asset, because this would amount to getting values and not prices. For this reason, Bergomi considers and ultimately rejects the Heston model. He describes it as an "elementary attempt at designing a model such that implied volatilities are not frozen anymore and have their own dynamics" (p. 217), but he immediately recognizes its main caveat, which is that this can only be "done by specifying an SDE for the instantaneous variance, a non-physical object." The dynamics of implied volatilities are indeed 'extracted' (as Bergomi says) in this case (i.e., derived from the SDE). They are not given by the market, and as a consequence, Heston is not a market model. By contrast, when he finally introduces the forward variance models in chapter 7, Bergomi declares: "In this chapter we model implied volatilities directly." This means they are the primitive state variables. The difficulty, as we have noted above, is that it then becomes unsayable, or taboo, to recognize that those implied volatilities will have to have their covariance matrix constrained in a way that is covertly consistent with their derivative nature (and will no longer be able to be constant, for instance, or at least to be 'chosen' as Bergomi may have wished).

Modeling implied volatilities directly

In the traditional approach in which, as opposed to the market models approach, a stochastic structure is given first (i.e., the instantaneous volatility of the underlying asset is modeled first, as with Heston above and indeed the rough volatility model), it is only a formal choice, and it is not fundamental, whether we represent the instantaneous volatility of the underlying asset directly, and write for it a stochastic process, or whether we represent it through a stochastic process that we write for the implied volatility of a certain derivative on that asset, typically an at-the-money option or a variance swap. For this implied volatility is just another way of naming the theoretical value of this derivative, *which is and always will be computed from the underlying stochastic process*, and is not its market price. Changing variables, from instantaneous volatility of the underlying asset to implied volatility of the option written on it, may be cumbersome; however, fundamentally, the model and the doctrine underlying it remain unchanged. People who have read Bergomi and adopted or adapted his model, but who have completely overlooked the radical philosophy and fundamental agenda of the market models, think that all that Bergomi has achieved is model forward variance (as a matter of fact, the whole curve thereof) instead of writing a model for instantaneous volatility. Thus Gatheral, Guyon. Precisely, the mapping may be less cumbersome with forward variance contracts than with vanilla options; however, what Bergomi tries to model are the implied volatilities of forward variance contracts directly, as given by the market (i.e., what we would get by inverting the BSM formula against their *market prices*). This is shown to be computed as a combination of the implied volatilities of vanillas (p. 141). To write

...Bergomi declares: "In this chapter we model implied volatilities directly." This means they are the primitive state variables

>

a stochastic process for implied volatilities *in the traditional approach* does not mean we are modeling the volatility market. You may ask: Why not? Why can't we write a stochastic process for options prices the same way we do for the underlying price, and call it a model of the options market, the same way we called the latter a model of the underlying market? This is because we are compelled, in the traditional approach, immediately to infer the stochastic process of the underlying asset which those options price processes imply.[15] Indeed, this is the only way, first to ensure that the options prices are arbitrage-free, and second to compute the price of exotic options. In other words, a process of instantaneous volatility of the underlying asset is always implicitly assumed, in the traditional approach. It alone explains the variability of options values for a fixed underlying price, so we may as well start with it, as we do usually, because causality always flows in this direction. And it surely wouldn't cross the mind of anybody that the options which are trading in the market by the sheer forces of supply and demand would inversely explain or cause the process of stochastic volatility of the underlying asset, in such a way that the process of prices of the former would be modeled before the latter (i.e., modeled directly, or first). Options values do not vary, in the traditional approach, because options trade freely, but because the instantaneous volatility of the underlying asset varies. To repeat, there are no options prices or options markets in the traditional approach, but only options values.

In Bergomi, by contrast, there is no underlying stochastic process and no arbitrage principle (or they are unsayable). The implied volatilities, or options market prices we are modeling *directly*, may have nothing to do, from beginning to end, with the volatility of the underlying. Or if they do, it remains a secret deeply buried under the surface of the market, and the market will have always solved for us the problem (it screens the problem off). To repeat, Bergomi does not necessarily think that the underlying process doesn't exist in itself; he only doesn't wish it to underlie derivative pricing and to confuse the student thereof. So, it may very well be the case that the underlying process exists and the instantaneous volatility of the asset is modeled, even by Bergomi himself, on a separate page, yet the BSM implied volatilities of the instruments that he is considering as hedging instruments and modeling directly, on this page, *result from the forces of the independent options market and are unrelated to the instantaneous volatility*. This is because no non-arbitrage constraints exist. The pricing function of the exotic option admits the prices of the hedging instruments as arguments, and the only doctrine or ground for faith is that it will continue to do so and to be the pricing function with those arguments, by dint of something Bergomi calls a *trading decision* (p. 16). It doesn't draw its validity from an existing stochastic structure, which might be implicit but would exist nonetheless and under which alone we would ultimately be computing the value of the exotic option. For one thing, a market price cannot be a value, or a result of computation. For another, the existing stochastic structure, as its very name indicates, can only belong in the past. By contrast, our pricing function draws its validity from the future, as the act of faith, here, is that *it will continue to be the market price of the exotic option*. The market pricing engine will always have preceded any pricing engine. For all we know, the hedging instruments underlying the market pricing function may even have nothing to do with the underlying asset, except being correlated with it (i.e., they are not derivative on it). They underlie the pricing function because they will continue to underlie it.

Perhaps the pricing function is a very peculiar function in the sense that it should precisely never admit of values as output (in mathematics, we usually say that 'a function is valued'), but always prices

We will return later to this crucial point. Suffice it to say, for the moment, that we would indeed be very disappointed if all that Bergomi was meaning to say, with his market models and the inversion of the arrows of causality and the fact of modeling the implied volatilities directly, was only a remapping and a rewording of the traditional valuation approach where the instantaneous volatility of the underlying asset is modeled as the first cause. But to end the digression and to go back to our question: "Does Bergomi have no choice but to rely on an underlying structure, when he realizes that there is (yet) no market for the volatility of volatility, and therefore has to make some assumptions about the realized volatility of the forward variance curve?" The answer is that he will make that *temporary* assumption but will, in essence, be constantly looking for the missing volatility market. He will constantly be looking forward. The vocation of the market models, if we may state it again, is always to rely on the market of options, and of options written on the options (options on VIX or variance, which Bergomi recognizes are still missing for the moment), and never to recognize a fundamental stochastic structure, although the practical solution, at any given stage, may have to assume a certain structure. Perhaps the ultimate pricing function $P(S,\xi,t)$, which is none other than the market's, is not representable or even writable and all we can do is represent it provisionally by projecting it on the valuation plane (think of the representation of a fractal curve which is always provisional). Perhaps the pricing function is a very peculiar function in the sense that it should precisely never admit of values as output (in mathematics, we usually say that 'a function is valued'), but always prices. On that, more later.

This fundamental philosophical position of market models (never recognizing ultimate structure although always assuming temporary structure) is again afforded by the reliance on the accounting equation and its ex-post character, as opposed to the a priori character of the structure.

Crucially, the paragraph in Bergomi's book that we are discussing is the one in which he offers his famous temporary structure, which will come to be known and celebrated as the *Bergomi model* (with the famous exponential kernels that will bring about all the criticism from Gatheral). It is also true that this particular temporary structure (the *N*-factor Bergomi model) allows him eventually to infer the underlying stock price process, and eventually to write a partial differential equation (PDE) with the stock price S and the N Ornstein–Ühlenbeck variables X^i as state variables; a PDE that will enable him to evaluate everything (i.e., the exotic option as well as the full vanilla surface and the full forward variance curve, or the very instruments which we were not supposed to evaluate in a model but always admit as market givens). However, as we will explain below, these two structures (the stock price process and the PDE with the N volatility processes) are only interpretations or computational artifacts. They enable us to *interpret* the market price of the exotic option as a value output by a PDE that depends on S and X^i. However, what the market price of the exotic option is and will always be, in reality, is a function (the *pricing function*) of the market prices of the hedging instruments, and what the market prices of the latter are and will always be, in reality, are the prices given by the market, and by no model.

So how, in the absence of stochastic structure, is any structure of the market models always temporary, and how is this related to the fundamental characteristic of the market models, which is the inversion of the arrow of causality and the replacement of ex-ante valuation or stochastic control (or the very image of depth) with an ex-post accounting equation (or floating on the surface)? And what to make of the consequence, which is that the stochastic structure and the whole of probability will turn out to be nothing more than interpretations? This will take us to a book by Shafer and Vovk, which Bergomi may not be aware of, and where it seems to us that the suggestion was made for the first time to relinquish the stochastic structure and to speak of prices in a game with Reality instead. Shafer and Vovk typically derive the BSM equation by relying on ex-post accounting rather than probabilistic a priori structure, following exactly the same lines as Bergomi.[16] We will review this overturning of the financial paradigm by Shafer and Vovk before we return to Bergomi's very closely related problem and to the two novel concepts that we read, in his book, as the ones potentially overturning derivative pricing, the *trading decision* and the *pricing function*.

Shafer and Vovk

Key in Shafer and Vovk's account is an inversion of the arrows very similar to Bergomi's. Recall that Bergomi denied the existence of a stochastic process for the underlying price S (or at least was agnostic towards such existence), and equivalently the existence of a stochastic process for the overlying price P (i.e., the derivative's price, which will demonstrably be a function $P(S, t)$ of the underlying price and time, hence the reason I say 'equivalently'), but that didn't stop him, in his derivation of the BSM equation, from assuming two different prices, both of the underlying asset and of the derivative at two different instants. These prices were given by the market and therefore considered to be ex-post (i.e., empirical observations, given without a means to see how they were given – i.e., without a generator). If we assume that the prices S and P of the underlying asset and the derivative are given at any time, yet we assume there doesn't exist any stochastic process, what is it exactly, then, that we are not assuming? What is there, in the stochastic process, over and above the empirical time series of prices? The answer is the *stochastic structure*, and with that we mean all the other possible prices. We wish the market to give us, at any time, the stock price S and the price P of a certain payoff, or a certain derivative, but we don't necessarily impose on the market to give us, at any time, the prices of all imaginable payoffs, the state-dependent as well as the path-dependent ones, for that would be equivalent to giving the (risk-neutral) probability of all the imaginable paths, and consequently all the conditional probabilities (i.e., to giving the full stochastic process).[17] Bergomi engages in a game with the market, where 'game' is supposed to mean that the market may be frivolous and hiding its intentions, and does not necessarily have to reveal its full internal mechanism or structure or how it came up with its moves, a game in which, at each step, the market announces the market prices of the hedging instruments and Bergomi announces his trading decision and the pricing function P of the exotic option, in response. The word 'game' is here to emphasize the pragmatical aspect of the process: all that counts are the results of the game, what goes on above the surface of the table and not underneath. The game has a specific (and quite economical) goal, which allows us to abstract and dispense with anything surrounding the game, what typically Shafer and Vovk call "extrinsic stochastic modeling."[18] The only result Bergomi is pursuing in the game is to control the P&L, or at least to be in a position to control it (i.e., to prevent it from growing uncontrollably positive or negative, p. 4). This is in line with Shafer and Vovk, who essentially speak of games and are interested in the fortune the player accumulates at the end of the game and not in the stochastic structure that may have driven the device the player is facing (the roulette wheel, the weather system, the financial market) to present him with the odds at which sequentially to place bets. They, too, consider the ex-post accounting equation and do not care about the probabilistic

The word 'game' is here to emphasize the pragmatical aspect of the process: all that counts are the results of the game, what goes on above the surface of the table and not underneath

>

model and its ex-ante stance. They, too, transform the theorems of probability and stochastic control into statements about break-even and the impossibility of becoming infinitely rich. They write: "Probability becomes game-theoretic as soon as we treat the expected values in a probability model as prices in a game. Many of probability's theorems turn out to be theorems about the existence of winning strategies for the player who is betting on what the world or market will do."[19]

Although Shafer and Vovk speak of probability models in general and not specifically of financial models, the word 'price' in their lexicon is meant to emphasize something practical and material (i.e., the game), as opposed to something contemplative, such as an expected value or equivalently a probability. 'Price' is here to suggest something that implies an actual unfolding, whose material and ex-post results we will only be accounting for, and not something frozen in anticipation and in the contemplation of structure. This, to our eyes, is what illuminates the opposition between price and value in Bergomi, and his talk of a *usable* (market) model as opposed to a model in charge of predicting the behavior of the market. Crucially, Shafer and Vovk write: "Defining a probability measure on a sample space means recommending a definite price for each uncertain payoff that can be defined on the sample space, a price at which one might buy or sell the payoff. Our framework requires much less than this. We may be given only a few prices."[20] This amounts to dispensing with the underlying process and with the existence of a pricing kernel, or a price for every path, or once again, with the principle of non-arbitrage. Prices are just entry tickets in the practical and economical game we happen to play; they are not supposed to reflect the whole universe at once, in the speculative (and quite exorbitant) sense of metaphysics. Shafer and Vovk continue: "Our framework differs most strikingly from the measure-theoretic framework in its ability to model open processes – processes that are open to influences we cannot model even probabilistically."[21] In the light of this important statement, we can see why the pricing process, which Bergomi has no choice but to devolve to a pricing function – because he has no choice, in the end, but to use mathematics in his technology book and his engineering task – is in fact an open process too. Probability can only be an interpretation, or a temporary diagram we need in order to represent the pricing and accomplish computations that must satisfy some consistency requirements (typically non-arbitrage constraints) if the technology is to be reliable and marketable. To that purpose, probability freezes the pricing process in a temporary structure, or in a probabilistic model, whose logic can only flow backwards, because it has no choice but to freeze the future; however, the real pricing process is supposed to remain open and to flow forwards. In this respect, it exceeds probability and certainly embeds influences that cannot be modeled probabilistically, as Shafer and Vovk point out, if only the human intervention which changes the probabilistic model when needed (recalibration).

In a brief and incisive commentary, Shafer and Vovk reveal the philosophical view underpinning this shift from probability structure to games, from ex-ante expectations to ex-post accounting and actual outcomes

The core philosophy of market models

Bergomi doesn't say it, or may even have never thought of it – probably he has no patience for philosophy or semantics – but it falls upon us, interpreters of his book, to understand his pricing function for what it is really: as something that exceeds the traditional concept of a mathematical function which takes arguments as inputs and produces values as outputs. We wish literally to understand it as a *pricing* function, as opposed to a *valuing* function. As a pricing function, it is the result of what Bergomi calls a *trading decision*; it is definitely rooted in a pragmatical doctrine, and not a function we would contemplatively decipher in the 'physical' laws of the market, and would aim ultimately to make transparent. Bergomi writes: "We *decide* to make our exotic option's price a function of other derivatives' prices" (p. 16, emphasis in the original). Because it is a decision, it remains open to adaptation. As a pricing function, it will, on the other hand, only be judged and evaluated and kept, based on its practical result, which is the ex-post accounting equation and the control of P&L. As such, the stochastic structure which lends it this temporary specific form, what Bergomi calls the "covariance structure of the model," has, as he says, "no special significance" (p. 225). This echoes perfectly what Shafer and Vovk write: "The forecasting success of a probability distribution for a sequence of events should be evaluated using only the *actual outcomes* and the sequence of forecasts (conditional probabilities) to which these outcomes give rise, without reference to other aspects of the probability distribution. […] The additional information contained in a probability measure that has these probability forecasts as conditional probabilities should not enter into the evaluation, even if the probability measure was actually constructed and the probability forecasts were derived from it."[22] By evaluation, Shafer and Vovk mean the scoring of the probability distribution or probabilistic model that was used to produce the the forecasts, or judging its success and deciding whether to keep it or not for all practical purposes. And when they add that the additional information contained in the full probability measure should not be used, this is similar to Bergomi's recommendation to disregard the specific covariance structure of the model as insignificant and never to use it in the practical evaluation of the model, or in the only practical test of its success, which is hedging. He thus repeatedly warns us against hedging the price of the exotic option, or the output of the pricing function, with the specific parameters of the model (p. 225).[23]

In a brief and incisive commentary, Shafer and Vovk reveal the philosophical view underpinning this shift from probability structure to games, from ex-ante expectations to ex-post accounting and actual outcomes. At work, indeed, is nothing less than a full revaluation of the meaning and nature of chance. "What is a stochastic mechanism?" they write, "What does it mean to suppose that a phenomenon, say the weather at a particular time and place, is generated by chance according to a particular probability measure? Scientists and statisticians who use probability theory often answer this question with a self-consciously outlandish metaphor: A demigod tosses a coin or draws from a deck of cards to decide what the weather will be, and our job is to discover the bias of the coin or the proportions of different types of cards in the deck."[24] It surely won't escape anybody that what Shafer and Vovk call the "outlandish metaphor of a demi-god tossing a coin or drawing a card" is what underlies the whole notion of the sample space Ω and of its sample elements ω, and what they call "our job" as opposed to the demigod's is that of defining the algebra of events and the probability measure (i.e., the point of view). In denouncing this 'outlandish metaphor,' Shafer and Vovk thus put in question nothing less than the whole foundation of abstract probability theory, based on measure theory. By contrast with this metaphor, which, according to Shafer and Vovk, "drives statisticians to hypothesize full probability measures for the phenomena they study and to make these measures yet more extensive and complicated whenever their details are contradicted by empirical data," the metaphor of the game allows them to "get started without a complete probability measure, such as might be defined by a biased coin or a deck of cards." In this way, they "can accommodate the idea that the phenomenon they are modeling might have only limited regularities, which permit the pricing of only some of its uncertainties." This, according to Shafer and Vovk, "encourages a minimalist philosophy" in which "the prices for each step of the process define a full probability distribution for what happens on that step," yet without this having necessarily to imply that a full probability distribution has been defined for the process.[25] Transposing this to Bergomi and to the market models, we may say that the temporary covariance structure enables the market-maker at each step to quote options prices that are arbitrage-free, but that there is no guarantee that this structure will persist over time and will continue to underlie the famous *pricing function*, or in other words to coincide with the market. This is recognizing that the recalibration of the pricing model (suddenly changing the model parameters, or shifting, say, from a two-factor Bergomi model to a three-factor model) should not itself be modeled probabilistically.[26] It is only excessive faith in probability and excessive reliance on measure theory that would make us want to internalize those changes and make the measures each time "more extensive and complicated." As Shafer and Vovk say: "We may react to empirical refutation by withdrawing some of these prices rather than adding more."[27]

Thus, we can see how far Bergomi's inaugural and emblematic statement, to the effect that an underlying stochastic process may not exist, is taking him and taking us along with him

Thus, we can see how far Bergomi's inaugural and emblematic statement, to the effect that an underlying stochastic process may not exist, is taking him and taking us along with him. To repeat, postulating this inexistence and investigating what it means exactly is the key to unlocking the meaning and philosophy of the market models. This statement of inexistence distinguishes Bergomi from all the quants of his generation. I will go even as far as claiming that it is what motivates writing his book, as opposed to just writing the model that happened to be associated with his name. For instance, Henry-Labordère introduces the expression 'market-models' in his own opus of derivative pricing; however, it is soon very clear that an underlying process is given first and that we will, for the rest of the book, inhabit the traditional space of valuation under the risk-neutral measure and no market.[28] As for readers of Bergomi, such as Guyon and Gatheral, they locate his main innovation in having succeeded to model the forward variance curve, instead of the instantaneous volatility; however, they also interpret it against the background of an existing underlying stochastic process, which reduces it, as we said above, merely to a remapping. It is true that all that a quant will care about, at the end of the day, are the quantitative capacities of the model, and since the present quantitative challenge in stochastic volatility modeling is to jointly explain the volatility smiles of SPX options and VIX options, all that the quantitative community seems to be concerned about is to improve on Bergomi, concerning this specific problem, and to propose the next quantitative model: rough Bergomi (Gatheral), in which the exponential kernels are replaced with inverse power kernels; discrete-time models (Guyon), in which the tension that prevents continuous-time models from fitting jointly the short-term SPX volatility skew and the VIX futures and options is released. Nobody really seems to pay much attention to what constitutes the cement, and I dare even say, the matter of Bergomi's book, which are his tireless comments about the particular structure having no significance, or about the need always to look for break-even conditions instead of non-arbitrage conditions, or comments about

>

every single stochastic process he may have written in his book, starting with the underlying asset and ending with the full forward variance curve (p. 219), being a *probabilistic interpretation* and not reality – I mean, not even reality as seen by the model. Surely, all these 'philosophical' warnings change nothing to the quantitative result, which is the numerical output of the Bergomi model, so why does Bergomi bother with them and litter his book with them? In order to answer this question, let's engage a little deeper in the philosophical and semantic analysis of Bergomi. Let us consider, once again, the crucial distinction he makes between the BSM model and the BSM equation. He claims to derive the second while not using the first. Since the quantitative result that any down-to-earth quant cares about is the BSM equation anyway, what could the difference be? Why make the distinction?

Semantics and modalities in Bergomi

It all starts with the hedged position, or the portfolio constituted by the option and its underlying dynamic hedge. Bergomi doesn't wish to control the P&L a priori. He doesn't force it to be equal to zero necessarily, which is an a priori condition (valid in all possible cases). He plays out the modalities differently. All he requires is that *it should not be a priori impossible that the P&L should be found to be equal to zero a posteriori*. Instead of saying: "It is necessary that the P&L should be equal to zero," the crucial weakening amounts to saying: "It is necessarily possible that the P&L could be equal to zero." There must exist at least one case in which break-even takes place. This is how Bergomi is able to derive the BSM equation without the BSM model, and the suggestion, at the end of the derivation, is to use the BSM formula with the estimate of future realized volatility $\hat{\sigma}$ (what we think ex-ante that ex-post volatility will be). But how is this different from what we've always been doing with BSM? When BSM assume a stochastic process of constant volatility σ (i.e., the BSM model) and derive their equation, the practical suggestion at the end of the exercise is also to use an estimate $\hat{\sigma}$ of future realized volatility in the formula. So, what's the difference? You may think it is only metaphysical, since the end-user of the formula will not see any difference. I believe the difference is huge, because it points to the semantics of the problem, to what the science or the technology – for Bergomi is definitely developing a technology – or the practice is ultimately saying. It points to what is really meant by pricing formula and pricing equation, or simply by pricing. Most importantly, analyzing the difference amounts to sorting out what exists, or what is known to exist and under what shape or form, or what can be written. There is, for instance, a famous paper by Ahmad and Wilmott where pricing and hedging is accomplished by the BSM formula with the only volatility that is known and observed instantly by the trader, implied volatility, and the analysis of what is being achieved, rising above the knowledge of the trader, is accomplished by decomposing the P&L under the 'real' stochastic process of the underlying price and the 'real' instantaneous volatility, supposed to be known to whoever's mind is writing the paper.[29] True, the paper is to some extent grown-up with respect to BSM, as it considers that real volatility will be different from what the user of BSM assumes; however, it is not as grown-up as Bergomi wishes it to be, because it is assuming the objective existence of a process, which may not be known to anybody except the writer of the paper. Bergomi relaxes even this.

We are no longer alone. There is always the market, acting as our mirror

There is total symmetry in Bergomi. He does not assume that somebody is using BSM with a wrong volatility number on one side (the trader) and that somebody else (the writer of the paper) is objectively computing the trader's P&L with the true knowledge of the volatility number. For this would amount to just another theoretical paper in finance, where a certain process with a known volatility (to the writer's mind) is given and where, instead of computing the arbitrage-free value of the option (in the writer's mind), we complicate the problem a little bit and compute the P&L of the trader who is using the wrong formula. The objective plane (the writer's mind) is given all the same; we could have used it to compute the arbitrage-free value of the option as in BSM, instead of computing the P&L of a trader who thinks he is computing the arbitrage-free value of the option and hence incurring no P&L, or who knows such a value is beyond his knowledge and he is only using a proxy (implied volatility). And if we have this knowledge on our objective plane, why didn't we lend it to the trader in the first place? Bergomi, by contrast, does not assume the existence of the objective plane (the point of view of theoretical finance). He is with the trader, alone, all the time. It is forbidden to write an objective process. Nothing is computed a priori. The only a priori prescription is that the P&L should decompose into two terms such that it may not be impossible that the P&L should be zero a posteriori.

We may call this weakening, or this inversion of the arrows, a new theory of derivative valuation, which is no longer founded on the principle of non-arbitrage, because no underlying process is assumed a priori, either to be written by the trader (whose mind would coincide, then, with BSM) or by someone watching the trader from the objective plane. This new theory of valuation without the principle of non-arbitrage is just a *pricing*, now understood in the sense that prices are provided by the market, or by a bank who consistently showed them to the market and consistently controlled (ex-post) the P&L of its positions, hence was not driven outside the market (or the game) – and hence became a market-maker and the prices it is showing became the market. As a matter of fact, Bergomi calls 'prices' the result of the pricing formula when the P&L of the hedged position is decomposed appropriately and displays appropriately the *possibility* of a break-even volatility number, or covariance matrix. He does so when he explicitly wonders whether the result produced by formulas that do not so appropriately behave, what he calls *unusable models*, is a price (p. 455). Usually, in financial theory, price means a martingale (i.e., precisely something gotten by a non-arbitrage argument). By contrast, there is a philosophical shift in Bergomi, because we are now deemed able to price something (the new meaning of valuation), not when we produce an arbitrage-free value for it, but when the pricing formula we used (which may have been originally handed over to us by a mysterious quant) and posted to the market, and was posted back to us by the market (i.e., we weren't driven outside the market by using it), allowed us to control the P&L a posteriori.

We are no longer alone. There is always the market, acting as our mirror. We produce the pricing formulas, and as a matter of fact, we try to derive them or see what form they may have, but ultimately they are produced by the market, they are there, given, and we may only act ex-post relative to them, and only weaken down our assumption to controlling the P&L through the peculiar rearrangement of the modalities above. This is what it takes to reverse the arrows and to relax the assumption of the underlying process. This is what it means that the underlying process no longer exists. *Its nonexistence means exactly the existence of the derivatives market*. We no longer impose non-arbitrage on the market, or the arrow going our way – for who are we to impose anything on the market? Instead, we hold that non-arbitrage must apply *because* there exists the market: the arrow goes in the other way and the market has already solved the problem for us. We don't question how the market started. We immediately start in the middle of the reflective loop: posting the pricing formula, being given the pricing formula in return. Our only link with what might look as an underlying assumption is that, statistically (i.e., ex-post), the returns of the prices of the underlying hedging instruments average out to their variances and there is at least the possibility that break-even statistically occurs.

Rough volatility

From this, we can realize the total blunder of those whom Bergomi calls the "expectation calculators" (p. xv) or those who criticize the simplistic assumptions of BSM as regards the underlying process.[30] Anyone aiming to model accurately the underlying process is of course welcome to do so, however it is time we understood this has nothing to do with the derivatives market. In the derivatives market, as said, there doesn't exist an underlying process, where its nonexistence is of course not meant in the physical sense but in the sense of the absence of appropriate register. We simply *are not* in this register. There is nothing wrong with God's rough volatility model being supported by a time series analysis of real volatility or by an argument probing the market microstructure and the order books. There may even be nothing more sublime than this fine analysis of the underlying. But where this really becomes indefensible and highly debatable is the moment when the authors seek decisive evidence, to support their model, from the options market which, they tell us, can be fitted exactly by a rough volatility underlying process. Now, where did this fancy idea come from? Who ever said the valuation theory of options *in the options market*, or for short, *option pricing*, should be based on risk-neutral expectation under the underlying process? To repeat, it would be a great disappointment if all that Bergomi had meant, in insisting so much on modeling the implied volatilities directly, on the inversion of the arrows and on the disconnect between the pricing function and the particular covariance structure that we temporarily lend to it – in short, all the import of the market models – was just that option pricing should be conducted under the 'equivalent martingale measure of the market,' as the saying goes, or in other words, by following the non-arbitrage principle. For that requires a real underlying stochastic process to be given in the real measure, to which the pricing measure must be equivalent, and requires no market. To call it the 'martingale measure of the market,' or to think that the 'market' selects its pricing measure, is only metaphorical and remains of course very far from the idea of a real market. This, once again, would be valuation and not pricing.

The market is a trading arena and a place of constant recalibration of the pricing models, not of their selection. Compare two philosophies: (a) one in which we are never told how the 'market' finds its martingale measure: Is it equilibrium in the long run? Informational efficiency? A long process of trial and error by options traders who deeply learn the rough volatility model – so deeply they don't even realize it – until someone (Gatheral?) finally reveals to them they've always been using this model without knowing?[31] and (b) one in which the aim is to put in our hands the very tools and technology with which to make the market ourselves. We are not told anything either, in the second philosophy, only this time it is because telling and knowledge (as in: "I *know* the underlying process is BSM") are completely irrelevant. Bergomi doesn't criticize BSM for the simplicity of its assumptions, but for its *usage* as a model; for then, as he very rightly says, if the underlying process is shown by time series analysis no longer to be Brownian motion, what should we do and what should we use instead? The rough volatility proponents commit a bigger sin than just beginning with BSM as a model, and subsequently wondering what to do in case it no longer held, because they provide the terminus. Not only do they acquiesce in the blunder of questioning the BSM model and what to do when it fails, but they furnish the answer. The final answer is rough volatility: let us look no further. There is already a clever twist in rough volatility: volatility is not just stochastic, it is rough. Rough volatility is already trying to outsmart us and deliver to us the last word. Nothing wrong with that, mind you. But where this becomes a bit foolish, is to top up the cleverness with the argument from the options market. The authors lock themselves completely in the final stop: where to go from there? According to them, the problem is over. Not only is the underlying process truly rough volatility, following the time series analysis and God's own design, but even better, the options market

It just felt too tempting to replace the exponential kernels in Bergomi's model with the inverse power function

>

confirms it. One is tempted to reverse Bergomi's worry for them: What if the options market no longer concurred?

It just felt too tempting to replace the exponential kernels in Bergomi's model with the inverse power function. Granted, this fixed the behavior of the short-term at-the-money skew, but the price to pay was rough volatility and the path-dependent nature of the model. Thank God, the analysis of market microstructure and of times series of 'instantaneous' volatility seemed to confirm the rough volatility process, but the problem we now faced was how to deal with the loss of the Markovian property, which is key in all practical applications and feasible pricing algorithms. Here again, Bergomi seemed exactly to offer the mold: the forward variance curve which he models directly. For it turned out that the rough volatility model becomes Markovian when the full forward variance curve is used as a state variable. Turning the path variable into a state variable is indeed a well-known technique to make a path-dependent pricing problem amenable to resolution by a PDE, even if its dimension is infinite as a result (further reductions will deal with that). As a matter of fact, modeling the forward variance curve directly had also led to path-dependency in Bergomi, and his choice of the exponential kernels was his way of dealing with it. But where I say there is misusage of Bergomi by the rough volatility model is that the forward variance curve, or the infinite number of forward variance contracts ξ^T composing it, were supposed to remain independent assets in Bergomi, and we were always supposed to use all of them for hedging. This is Bergomi's continuous insistence on the particular covariance structure (in his case, the N-factor Bergomi model) being only a numerical and computational expedient, which should never dictate or simplify the hedging, for that would contradict the philosophy of the market models. The rough volatility proponents recommend just the opposite. In total opposition to Bergomi, they insist their model is a two-factor model composed of the asset price and its instantaneous volatility, hence one volatility instrument (an at-the-money option, a single variance contract) is sufficient for perfect dynamic hedging in conjunction with the underlying asset. The path variable which helps them turn the process Markovian, or the full forward variance curve, was never meant as an independently trading instrument, therefore destined to hedging, any more than the average of the closing underlying prices, which we use as a path variable in the PDE pricing of Asian options, is meant to be trading independently of the stock.

The rough volatility model would be okay and would, as a matter of fact, fit perfectly well within the category of market models and the compass of Bergomi's book, if it were interpreted as a *temporary covariance structure*, akin to Bergomi's model, that we are temporarily lending to a forward variance market model in which the implied volatilities of the forward variance contracts are modeled directly, and the latter are meant as the irreducible set of hedging instruments. Only, we should no longer call it 'God's model' in this case, and certainly not be so triumphant about it. This is not only a semantic debate, or my attempt at spoiling God's banquet. Hedging is important. Ultimately, it is what pricing is all about. I certainly can understand that two quants may disagree on the particular model, and split hairs about which model better calibrates SPX smiles jointl y with VIX. But I cannot allow that an issue as important as hedging remain unsettled: Do I hedge with a single contract or with the full curve? This question has nothing to do with the quality of the hedge, or how to improve it. For it is the question of whether to rely on a market model, where only options markets exist and there exists no underlying stochastic process, or on an expectation calculation, where we might happen to know what probability, volatility, and even rough volatility mean, but we don't know what a volatility market means.

I certainly can understand that two quants may disagree on the particular model, and split hairs about which model better calibrates SPX smiles jointly with VIX

To be continued.

ENDNOTES

1. Citing from the article 'Quants of the year: Jim Gatheral and Mathieu Rosenbaum' (*Risk*, February 2021): "Rough volatility resulted not from a eureka moment but from incremental steps Gatheral and Rosenbaum made over the years, with the contribution of other researchers who worked with them. […] [Gatheral] saw patterns in the tick-by-tick data of market microstructure that shaped his ideas on how volatility behaves. Then, in 2012, Gatheral was shown an application of fractional Brownian motion originally developed by Elisa Alòs and became quickly convinced the true model of volatility must follow that form. 'I knew it had to be correct. I started to call fractional Brownian motion God's model,' he says."
2. Omar El Euch. Quantitative Finance under rough volatility. Numerical Analysis [math.NA]. Sorbonne Université, 2018. English. NNT: 2018SORUS172.
3. Jim Gatheral. Rough volatility with Python. https://tpq.io/p/rough_volatility_with_python.html. Accessed 29 June 2021.
4. Lorenzo Bergomi. Communication to Dan Tudball.
5. Lorenzo Bergomi 2016. *Stochastic Volatility Modeling*. Boca Raton, FL: CRC Press, pp. 1–2. All subsequent citations from Bergomi will be identified in the text by their page number in this edition.
6. Lorenzo Bergomi. Communication to Dan Tudball.
7. MacKenzie, D. (2003). 'An equation and its worlds: bricolage, exemplars, disunity and performativity in financial economics,' *Social Studies of Science*, 33, 831–868.
8. Harrison, J. M. and Pliska, S. R. (1981). 'Martingales and stochastic integrals in the theory of continuous trading,' *Stochastic Processes and their Applications*, 11, 215–260.
9. Haug, E. G. and Taleb, N. N. (2008). 'Why we have never used the Black–Scholes–Merton option pricing formula,' *Wilmott*, Jan, 72–79.
10. Lorenzo Bergomi. Communication to Dan Tudball.
11. Indeed Bergomi never speaks of non-arbitrage; only of prices not being nonsensical. It is true the PDE he will ultimately write admits of a probabilistic interpretation, hence produces prices that are invulnerable to a Dutch book argument, but the PDE is not the ultimate theory in Bergomi; it is just a numerical expedient. The ultimate pricing engine is the market. Its prices may be arbitrage-free at any moment and a pricing kernel may exist at any moment; however, this kernel may be changing, in time, in ways that cannot be modeled probabilistically.
12. Lorenzo Bergomi. Communication to Dan Tudball.
13. Bergomi, L. (2017). 'Local-stochastic volatility: models and non-models,' *Risk*, Aug, 78–83.
14. "Market models aim at treating vanilla options on the same footing as the underlying itself: vanilla option prices observed at $t = 0$ are initial values of hedge instruments to be used as inputs in the model" (p. 25). "A market model takes as inputs vanilla option prices, in addition to the spot. The former are treated as hedge instruments, on the same footing as the underlying itself" (p. 77).
15. See Cont, R., da Fonseca, J., and Durrleman, V. (2002). 'Stochastic models of implied volatility surfaces,' *Economic Notes*, 31(2), 361–377.
16. They write: "The price the [BSM] formula gives for [the derivative] is its game-theoretic price, […], in a game between Investor and Market in which Investor is allowed to continuously adjust the amount of the stock he holds. […] To the protocol [of the game], Black and Scholes added the assumption that [the price of the underlying stock] follows a geometric Brownian motion. […] But the derivation of the Black–Scholes formula *does not actually use the full force of this assumption*" (Shafer, G. and Vovk, V. [2001]. *Probability and Finance: It Is Only a Game!* New York: John Wiley & Sons Ltd, pp. 217–218, our emphasis).
17. Note that this has nothing to do with the market being complete or incomplete. The whole notion of a market being complete or incomplete is set against the assumption of an underlying stochastic process. Suppose it is a jump-diffusion with stochastic volatility and stochastic jumps. In this case, there exist several families of arbitrage-free values for derivatives or several martingale measures. Suppose the vanilla options prices are available in the market, and the barrier options prices are not. The market is then incomplete, and two martingale measures can agree on the vanillas but disagree on the barriers. The market becomes complete when it 'selects' its martingale measure, which then becomes unique. Shafer and Vovk also speak of the derivatives prices being partially and not totally available in the market; however, no underlying stochastic process exists in their case. There is no probability to begin with, and the prices are precisely what replace it.
18. Shafer and Vovk, op. cit. p. 3.
19. Ibid. p. 2.
20. Ibid.
21. Ibid. p. 3.
22. Ibid. p. 58.
23. "Thus, with regard to deltas, the deformation modes of the variance curve generated by the processes have no special significance. Model factors simply set the structure and rank of the break-even covariance matrix of the gamma/theta P&L of a hedged position. […] It is important to stress that calculation of deltas is not connected in any way to the covariance structure of the hedging instruments in the model at hand" (p. 225).
24. Shafer and Vovk, op. cit. p. 22.
25. Ibid.
26. This is how Bergomi reacts to model recalibration: "Model parameters do not generate any P&L leakage, and only give rise to discrete P&Ls when they are changed and the option position is remarked with new parameter values" (p. 470).
27. Shafer and Vovk, op. cit. p. 22.
28. Henry-Labordère, P. (2009). *Analysis, Geometry, and Modeling in Finance: Advanced Methods in Option Pricing*. Boca Raton, FL: CRC Press.
29. Ahmad, R. and Wilmott, P. (2005). 'Which free lunch would you like today, sir?: Delta hedging, volatility arbitrage and optimal portfolios,' *Wilmott*, Nov, 64–79.
30. "One may think that a model derives its legitimacy and usefulness from the accuracy with which it captures the historical dynamics of the underlying security – hence the scorn demonstrated by econometricians and econo-physicists for the Black–Scholes model and its simplistic assumptions, upon first encounter" (p. 1).
31. MacKenzie once wondered what option pricing model traders could be enacting after the October 1987 crash and the emergence of volatility smiles (i.e., after having long enacted the BSM model). Gatheral seems finally to bring the answer: markets have been enacting the rough volatility model.