The Irony in the Variance Swaps

Elie Ayache

rony, according to the *Oxford English Dictionary*, is (a) a figure of speech in which the intended meaning is the opposite of that expressed by the words used (b) a condition of affairs or events of a character opposite to what was, or might naturally be, expected (In French, *ironie du sort*). In the following lines, I will propose a rereading of quantitative finance where irony, as opposed to theory, emerges as a leitmotiv, perhaps even a main guide. Variance swaps will provide me with a literal rehearsal of this ironic point, as I will show that the element of irony is inscribed in the movement motivating their existence, indeed in their very contractual terms.

The irony in quantitative finance, or specifically, in derivative pricing theory, is captured by the following observation. While the typical derivative paper is expressed in words and formulas aiming at the theoretical value of the derivative, it is really intended for derivative trading, which is the domain farthest away from theory. This is irony in the first sense. And while rational option pricing, as epitomized by the work of Black, Scholes and Merton, has triggered the explosion of options markets, the ensuing liquidity of option prices has turned volatility into a traded commodity thereby contradicting the crucial theoretical assumption of constant volatility. This is irony in the second sense.

On closer look, the irony in derivative pricing seems to hinge on the double-edged phenomenon of liquidity. Breakthroughs in the pricing of derivatives and the subsequent surge of liquidity in their markets seem to be linked with the news of their successful replication. Only when Black and Scholes succeeded in dynamically replicating the option with the underlying did option trading really take off. One may wonder of course what the need for an options market may be when their valuation makes them perfectly replicable and redundant instruments. However, liquidity has another side and the irony principle is here to remind us that there is a lot more to derivatives than meets the eye of theory. The reason why new derivatives are written

At that moment Ironic Man realized the connection between his lack of superpowers and the career he had chosen for himself ...

and traded is ultimately that they are *not* replicable by the existing stock of derivatives. Replication is largely theory-dependent and if it does help trigger the explosion of derivatives markets at first, their growing liquidity is soon to take over as self-sufficient pricing source. Irony always supersedes theory.

Still, replication remains the quant's favourite exercise, so much so that we may, without much exaggeration, *define* the quant as the party who is always busy looking for the right set of conditions and the clever mathematical trick for some exotic option to be replicable by the vanillas. True, perfect replication may have sounded unrealistic when applied to the underlying alone, and it may have required, on top of meticulous and continuous dynamic rebalancing, a very stringent theoretical assumption such as a diffusion process with constant volatility. However when we move to second generation derivatives, the hope is that the scores of vanilla options, which are available *ab initio* thanks to their now unquestionable liquidity, will themselves produce the replicating portfolio and—what's even better!—*statically* so.

Variance swaps and static replication

Variance swaps fall exactly in that category of hope. Their payoff at maturity is defined by:

$$N\left(\frac{1}{T}\sum_{t=0}^{T-1}\log^2\left(\frac{S_{t+1}}{S_t}\right)-\sigma_I^2\right)$$

where N is the notional amout, T the observation period, σ_I^2 the variance strike and S_t the observed underlying price at time t. While their growing popularity

is due to their character of a pure bet on variance, the liquidity of their markets, as evident today in the huge traded volumes and the narrow bid-and-ask spreads, is mainly explained by the possibility of statically replicating them with the vanillas: the so-called *model-independent valuation of variance swaps*. So in many ways, variance swaps are a blessing. To put it in Sebastien Bossu's words: "Variance swaps are ideal instruments to bet on volatility: unlike vanilla options, [they] do not require any delta-hedging."[1] Demeterfi and co-authors voice a similar enthusiasm about volatility swaps. "Stock options are impure," they write, "they provide exposure to both the direction of the stock price and its volatility. . . The easy way to trade volatility is to use volatility swaps. . . because they provide pure exposure to volatility (and only to volatility)."[2] To crown it all, variance swaps are relatively easier to price than the vanillas themselves as they require no more than the availability of the vanilla prices!

Note here a distinction between variance swap and volatility swap. The first delivers realized variance and the second only its square root. Convexity adjustments have to be made to recover the one from the other. Unlike the variance swap, the volatility swap is not statically replicable by the vanillas. It can only be dynamically replicated and therefore requires a modelling assumption for the dynamics of the underlying and its volatility. This point is rightly emphasized by both authors. "The reason the contract is based on variance [instead of volatility]," writes Bossu, "is that only the former can be replicated with a static hedge." As for Demeterfi et al., although their paper is dedicated to volatility swaps and not to variance swaps like Bossu, they motivate their preferred treatment of the latter (which will occupy the bulk of the paper) by the following words: "Although option markets participants talk of volatility, it is *variance*, or volatility squared, that has more fundamental theoretical significance. This is so because the correct way to value a swap is to value the portfolio that replicates it, and the swap that can be replicated most reliably (by portfolios of options of varying strikes) is a *variance swap*."

Both statements concerning the static replication of the variance swap contain the seeds of irony that I will expose later. They both intriguingly argue that variance, not volatility, has to be recognized as the most significant variable, *on no other account* than that the swap that can be statically replicated with options and can, *therefore*, be valued is the variance swap, not the volatility swap. As such they are vulnerable to irony because although they seem to want to express the fundamental reason why variance swaps have prevailed (in Demeterfi's case, the reason why variance is the most theoretically significant variable; in Bossu's, the reason why the contract is based on variance to begin with) the reason they give in reality is driven by theory and, what's even worse, by static replication. To compare, imagine Black and Scholes arguing that "Brownian motion has more fundamental theoretical significance than any other process *because* the correct way to value derivatives is to replicate them perfectly with the underlying and Brownian motion is the only dynamics where this can be done."

It is all a matter of point of view of course (irony is but a matter of point of view). We can perfectly understand Demeterfi's insistence on variance when we realize that all they really mean is that they will be basing their whole theoretical valuation on static replication with options and that this can only occur for variance. Seen from this angle, this whole advertising of variance suddenly appears for what it is: just an advance warning of the theoretical framework that will follow (diffusion) and the methodology that will be selected (static replication) rather than some general truth about variance. It makes perfect sense that Demeterfi et al. should be speaking of the "fundamental theoretical significance" of variance once it is understood that the "theory" and the "significance" are here relative to Demeterfi's paper and not general statements. What was misleading in Demeterfi's turn of phrase is just that they seemed to oppose variance, as such significantly recognized and advertised, to volatility, which is, they say, *what markets participants talk about*. "Although option markets participants talk of volatility, it is *variance*, or volatility squared, that has more fundamental theoretical significance," they write. This almost suggests that market participants should no longer talk of volatility but turn to variance instead. Pushing the thought a little further, this might even suggest that *volatility swaps*, which were originally written on the variable that is most widely and most naturally talked about, should be written on *variance* instead. (The official reason being of course that variance is the "more theoretically significant variable" but the real, unofficial reason being that it alone can be statically replicated.)

This covert recommendation is overtly endorsed by Bossu. Without hesitation he writes: "The reason the contract is based on variance is that only the former can be replicated with a static hedge." Note that Bossu is writing seven years after Demeterfi et al. This gives you the measure of the progress accomplished. What was only a prelude to a theoretical elaboration in Demeterfi's time, and almost an apology for having privileged the theoretical treatment of the variance swap rather than that of the original candidate, the volatility swap, becomes a hard fact in Bossu's, even a prescription for the terms of the contract itself. Also note that progress was not only achieved on the rhetorical side, but on the liquidity front as well. When Demeterfi et al. could only admit that "though [they are] theoretically simpler, variance swaps are less commonly traded," Bossu can write without problem today: "Variance swaps have become an increasingly popular type of 'light exotic' derivative instrument. Market participants are the major derivatives houses, hedge funds, and institutional investors. An unofficial estimate of the typical inter-broker trading volume is between \$1,000,000 and \$7,000,000 total vega notional in the European and American markets every day." (By "light exotic," Bossu probably means an exotic payoff that is less exotic than it first sounds; for instance he means that the variance swap, as exotic as its Asian-looking payoff may sound, is perfectly statically replicable by the vanillas.)

So the industry has progressed and no longer writes swaps on the variable that is most naturally talked about, but deliberately and almost "unnaturally" on the one that can be statically replicated. Witness the remarkable inversion. Methods and specific models are usually proposed to price the given, sometimes exotic, payoff structure, yet here, for the first time, a payoff structure (the variance swap) is specifically picked because it fits the given model and methodology (static replication). It almost feels as if the variance swap was invented by the quant for the quant in order to fit his framework and suit his pricing needs, not that it has emerged out of the market's own. A pure product of theory, we might say ("a pure bet on variance," "a perfect static replication"), not of nature. Beware of irony.

The irony of liquidity

I have already pointed to irony as the adverse effect of theory. There is always irony, I said, in the way the market ends up using, or rather appropriating, abusing, theory. No doubt the variance swaps were less liquid than the volatility swaps, at the time of Demeterfi's writing, because they did not correspond to the naturally talked about variable, and no doubt they are more liquid today because they can be statically hedged and the volatility swaps cannot. Hidden here are two different causes, even meanings, of liquidity, and it remains to be decided which is the right one. Let us not forget indeed that an instrument that is *truly* statically replicable by the existing ones (and by that I mean in a truly model-independent fashion) brings nothing new. It has no real reason to exist and even less so to become independently liquidly traded. For instance the *log contract* (on which, more later) is truly statically replicable by the vanilla options because, like them, it is a European payoff. Interestingly enough, it never really took off, despite Anthony Neuberger's efforts to promote it as a "pure volatility play" with better Greeks than the vanillas[3], or Demeterfi's whole point that it is the key to replicating variance. "There are no actively traded log contracts," Demeterfi et al. recognize, and this is why they have no other choice, in the remainder of their paper, but to turn to a portfolio of vanilla options as a proxy.

What I am suggesting here is that the variance swap, by contrast to the log contract, *is not truly statically replicable by the vanillas and this is why it is so liquid today!* It may have become liquid (at least more liquid than the volatility swap) because everybody *had thought* it was equivalent to the log contract and therefore statically replicable. However the reason it was invented and defined under the exact contractual terms that we know, the averaging of squared logarithms of daily returns, instead of being simply equated with the log contract, the reason that will justify its liquidity *a posteriori* although it may not have explained it *a priori*, is something different. Doubtless the variance swap would have been a poor invention if it will have never been liquid, so we may thank God that it has become so, even for the wrong reason.

A third degree of liquidity is here added to the first two. The first degree corresponds to the natural stage: a liquid market naturally emerges when the traded instrument corresponds to people's needs and practice. (This is the age of options markets prior to Black-Scholes, or equivalently the age of the volatility swaps.) The second corresponds to the theoretical stage: an instrument becomes liquid when it is shown to be replicable *in theory* by the existing ones. (This is the explosion of options markets after Black-Scholes, or equivalently the age of the variance swaps.) As for the third, it corresponds to the ironic stage: the liquidity of the given instrument now insinuates itself underneath the theoretical cover-up and reveals the true face of the instrument. (This is the age of the *new logic of option pricing and trading* that I have introduced elsewhere[4], or equivalently the age of the variance swap now recognized as a "heavy exotic.")

Recounting this story of irony is usually a topic for philosophical critique, however, since theory has permeated the very making of the variance swap (recall Bossu's "The reason the contract is based on variance is that only the former can be replicated with a static hedge") chances are an element of irony will have been planted too. So before I go any further, let me say a few words about derivative pricing theory and models (I have called them a "new logic") that *can* be adapted to the ironic stage. Of course you should expect this *ironist theory* to be the extreme opposite of the theory familiar to the quants, as previously defined: "Instead of looking for assumptions and tricks in order to prove that some exotic is replicable, find the reasons why it is truly exotic and may not be replicated! Instead of looking for liquidity that is the consequence of successful replication, find a benefit in a liquidity that is independent of replication!" To put it in Richard Rorty's words, ironists are people "never quite able to take themselves seriously because always aware that the terms in which they describe themselves are subject to change, always aware of the contingency and fragility of their final vocabularies, and thus of their selves."[5]

An ironist theory of derivative pricing

The ironist's first task is thus to never take seriously, i.e. take for final, any theoretical vocabulary or model she may have adopted, starting with Black-Scholes. Or should I say, ending with Black-Scholes? Although many advanced derivative pricing models have long superseded Black-Scholes, you have no idea indeed how dominant the Black-Scholes tradition still is. Think that the expression "volatility smile" is a consequence of the Black-Scholes vocabulary and wouldn't even be conceivable if Black-Scholes didn't exist. Think that Brownian volatility is just a theoretical construct needed to price options in the Black-Scholes framework, yet we keep hearing that options are a play on volatility and that options markets "imply" volatility. Think that variance is needed to price vanilla options, yet we keep hearing that vanilla options are needed to price variance swaps. Think that there is nothing special about the vanillas anyway, save that Black-Scholes first priced them, and that for all we know, variance swaps and surfaces of variance swaps may well have to become the primitive notions.

In my ironist theory of derivative pricing, only the relative liquidity of the derivative instruments counts and it is dictated by no prior model, or even less so, replication. It is perfectly conceivable that variance swaps (which are, we are told, a pure bet on variance) should be the most liquidly traded instruments in a market largely driven by volatility players, and should be used to replicate the less liquid vanillas, not the other way round. Equally, it is perfectly conceivable that the American digitals (or one-touch options), which are conceptually simpler than the vanillas, should be more liquid than the latter and help define (or calibrate) the underlying stochastic process instead of the traditional vanilla smile. There just is no predefined hierarchy between instruments of different nature. How the ironist model works is that it is first calibrated to the set of more liquid instruments (without any distinction between "exotic" and "vanilla," as this distinction itself is a heritage of Black-Scholes) and then the less liquid ones are priced by *dynamic* replication (what else?) using the former.

To allow this much flexibility, the model has of course to be quite general; in any case, more general than diffusion (and by that I mean "more general than *general* stochastic diffusion").

You must have guessed by now that my whole ironic protest against theory is in fact an assault against diffusion. Diffusion has long held us captive, to

the point where we no longer take notice of the walls of our confinement cell. Only in a diffusion framework is the variance swap equivalent to holding a log contract which is independently statically replicable by the vanillas. Only in a diffusion framework is the one-touch option, or generally the barrier option, statically replicable by a continuum of vanilla butterflies.

In my ironist derivative pricing model (we may also call it "realistic model") there is no such thing as static hedging because *truly* statically replicable instruments (and by that I mean "in a model-independent fashion") do not realistically exist or substantially trade. Of course call spreads and butterflies do exist and substantially trade, however I do not consider them as light exotics that are statically replicable by their simpler components. A call spread is a call spread and a butterfly is a butterfly. They are simply the combination of their components. By contrast, what *realistic* instruments there are in my universe (i.e. the true candidates for the calibration of the underlying process) are instruments with payoff structures not reducible to one another: for instance a vanilla option, a barrier option, a variance swap (now taken as a "heavy exotic" with the characteristically Asian feature), but *not the log contract*.

As a consequence, all hedging in my ironist (dare I call it "real"?) world is model-dependent *because it is dynamic*. (What else?) By the very meaning of the observation that a given derivative instrument is not replicable in a model-independent fashion, it cannot be replicated unless some dynamics is assumed and some stochastic control problem solved. How a given instrument is priced is then by computing the initial cost of the self-financing strategy that optimally replicates its payoff. This holds true of the variance swap and the volatility swap with their payoff taken at face value, i.e. as the dreadful Asian averaging of returns. Optimal replication through minimization of variance of P&L of the hedging portfolio is a good choice because derivative prices come out linear in the payoff. The unhedged residual (the minimized variance of P&L) is eventually computed and it is crucial in risk management or even in trading decisions.

How my model achieves its ironist goal is then simply by reiterating what I have called the "irony principle": Successful replication (in my case dynamic, optimal, and model-dependent) of an exotic instrument is certainly a great booster for its liquidity, however its growing liquidity will most likely drive its price away from the cost of its replication, i.e. from the prediction of the model. This is the moment when the ironist model becomes "aware that the terms in which it describes itself are subject to change" as it now asks that the deviating exotic be included in the calibration set and it now calls for recalibration. This is the true moment of model-independence. True exotics (in the sense that their payoff is not model-independently replicable) carry price information that has to be calibrated in the underlying process as soon as it becomes liquid and it deviates. When an exotic can no longer be priced in a model-dependent way (that is, its price deviates from the model) it can only be priced in a model-independent way (that is to say, it is now priced *by the market* and a new instance of the model is calibrated against its market price).

To go back to the variance swaps, I will therefore insist that they are not statically replicable by the vanillas because the diffusion assumption which makes this possible is too frail not to give way. It fails if only because the market price of the variance swap sooner or later deviates from the price of the log contract—and this will be theoretically explainable by the possibility

of jumps in the underlying or, better, will ironically point to the *necessity* of jumps! Chances are the variance swaps will then serve as calibrating instruments alongside the vanillas to help calibrate—guess what?—the jumps that make the difference! If I may summarize in one word the irony of the variance swaps, I will say, "What a great thing indeed that their valuation should be model-independent!"

Variance swaps and the Black–Scholes tradition
How then to make sense of Demeterfi and Bossu?

Blame it on irony if you will, but although they seem to be talking about variance swaps, they are in fact talking about the *log contract* (or so I will argue), this truly "light exotic" whose delta-neutral hedging will indeed capture realized variance when there are no jumps. Surely enough, both authors do start with the variance swap and by describing its exact payoff. Bossu even goes to some lengths in presenting a real life sample of its terms and conditions, complete with the names of the parties of the trade and the detailed formula for the averaging of the square logarithms of returns. Both authors then recognize that the variance swap is an "ideal instrument to bet on volatility" (Bossu), an instrument which provides "pure exposure to volatility" (Demeterfi et al.) and does not require delta hedging like the common option. However, when they move to the quantitative part of their paper, the part concerning the valuation of the variance swap, they first make a detour in the Black-Scholes framework in order to introduce the reader to variance replication and variance (or volatility) plays. This, Demeterfi et al. call the "intuitive approach" (or "How to create a portfolio of options whose variance sensitivity is independent of stock price"). They write: "We approach variance replication by building on the reader's assumed familiarity with the standard Black-Scholes model." Only in the second section will they adopt a more rigorous and general approach which, they say, will not depend on the full validity of the Black-Scholes model. As for Bossu, he never quite moves beyond the intuitive approach.

To be quite fair, Demeterfi and Bossu couldn't have proceeded differently. They write from within the Black-Scholes tradition after all, the tradition of volatility arbitrage and delta-neutral hedging, the tradition that has invented the notion of implied volatility as opposed to realized, not to say the notion of volatility itself. To their eyes, the variance swap can only be a *volatility play*, where "volatility" is the diffusion coefficient in the Black-Scholes model, and it is not clear at this stage whether they mean it in the implied or the realized sense of the word. I am not even sure the tradition itself knows the difference.

The variance swap as constant vega
Although Demeterfi et al. begin by saying: "In this section, we explain the replicating strategy that captures *realized* variance,'' they are soon to talk of options *vegas* rather than gammas, which makes us incline towards implied variance. Their intuitive approach to variance swaps in fact consists in first noting that the vega of a single vanilla option is dependent on the stock price with a peak in the neighbourhood of the strike price and second, in finding the combination of vanilla options with different strike prices that can produce a constant vega across the broadest range of stock prices. The

answer is a portfolio Π of options with weights inversely proportional to the square of the strike prices. Demeterfi and co-authors then go on to show that this portfolio has as theoretical value:

$$\Pi\left(S, \sigma\sqrt{\tau}\right) = \frac{S - S_*}{S_*} - \log\left(\frac{S}{S_*}\right) + \frac{\sigma^2\tau}{2}$$

where the interest rate is assumed to be zero for simplicity, S_* marks the boundary between puts and calls and $\tau = T - t$ is the time to maturity.

This is how the log contract first makes its appearance. Holding short a log contract is therefore holding exposure to both the stock price and its variance, like any other non linear European payoff, with the only difference that the variance vega is here independent of the stock price. (The curvature of the log contract is such that the option is always at-the-money.) Delta hedging leaves variance as the only determinant of value of the log contract, like with any other European option. Clearly, this is variance as it appears formally in the pricing formula above. My question: How is it any different from implied variance? How can the variance swap gain exposure to *realized* variance?

Note that in the Black-Scholes world (and by that I literally mean: *from inside the Black-Scholes model*), there is no distinction between the variance number you feed in the option pricing formula, the variance number you use to compute the delta-hedging ratio, and the variance of the underlying process. Also note that there is strictly speaking no meaning to the term "implied variance" as there is no meaning to the thought that option prices might be given (by the market, by fiat, etc.) and the corresponding variance computed by inversion of the pricing formula. There is even no meaning to the term "realized variance" as we are left wondering: "Realized" as opposed to what?

To talk of the "trading of variance," or of "variance bets," one has to step outside the Black-Scholes model (or rise above the *object level*) and introduce philosophical notions which disarticulate Black-Scholes as mathematical model in order to articulate the higher order concepts one has in mind, such as "forecast of future variance" (this smuggles in epistemology), "realized variance" (this smuggles in metaphysics), "the Black-Scholes formula as option pricing *tool*" (this implies truth is lying elsewhere, for instance in trading and in the market, and that the quantitative model is just an aid for the trader). These subtleties are usually unspoken when one's sole concern is option pricing. Everybody is aware of course that options are used to trade volatility or to bet on volatility, just like the variance swaps, however the topic of the option pricing paper is usually to compute the present value of a non linear payoff under some postulated underlying process. Volatility trading is left for practice, that is to say it remains outside the theory. (Often it contradicts the theory.) If Demeterfi and Bossu were to follow this line, they would stick with the pricing of an exotic option known as the "variance swap," whose payoff is the average of squared returns, and then they would see how its value reacts to variance or what not. The reason they don't is that the variance swap was *intended*, from the beginning, as a variance play. Its very design contains the necessity of speaking outside the theory or against it, like Demeterfi and Bossu do, what I have called the "seeds of irony."

For instance, this is how Demeterfi and co-authors make sense of the notion of realized variance and its effect on the value of the variance swap. They write: "For now assume we are in a Black-Scholes world where the implied volatility σ_I is the estimate [what does "estimate" mean in a Black-Scholes world?] of future realized volatility. If you take a position in the portfolio Π, the fair value you should pay at time $t = 0$ when the stock price is S_0 is:

$$\frac{S_0 - S_*}{S_*} - \log\left(\frac{S_0}{S_*}\right) + \frac{\sigma_I^2 T}{2}$$

At expiration, if the realized volatility turns out to have been σ_R [this is the step outside Black-Scholes] the initial fair value of the position captured by delta-hedging would have been:

$$\frac{S_0 - S_*}{S_*} - \log\left(\frac{S_0}{S_*}\right) + \frac{\sigma_R^2 T}{2}$$

The net P&L of the position, hedged to expiration, will be:

$$\text{P\&L} = \frac{T}{2}\left(\sigma_R^2 - \sigma_I^2\right)$$

Analyzing Demeterfi's argument, we can reconstruct it as follows. In theory, we *are* in a Black-Scholes world and the realized variance σ_R is the true variance, the only variance there is and will ever be. Delta-hedging and the underlying process unfold under this variance. Option pricing formulas, including the pricing formula of the portfolio Π above, are also supposed to rely on this variance number. However, there is one little difference in our world. Option prices are not given by the objective author of the theoretical paper (to whom there exists no doubt about the integrity of all three Black-Scholes meanings of variance and to whom knowledge is a meaningless category to begin with), but by an agent, call it the market, supposed to *know* variance. Since the category of knowledge is introduced and the integrity of the object level is broken, this "known variance" can of course be different from the true one. Therefore the portfolio Π is initially charged a different price.

As we move towards expiration, the only process that Demeterfi and co-authors describe is in fact the process of unveiling of the true number σ_R, not the process of generation of the underlying prices under σ_R or the process of actual delta-hedging. (As a matter of fact, they would be quite embarrassed to answer the question of the variance number that the delta-hedger is using in his delta-hedging formula.) So they rewind back to the time of initiation of the trade, and using the *same pricing formula*, they are now able to say: "The initial fair value would have been. . ."

The net P&L Demeterfi and co-authors express is therefore a difference of value of the position due to a difference of *implied variances* acting through the variance vega (which is aptly independent of the stock price). The "would have been" is not the unfolding, or the realization, *of* realized variance as we would have expected. It is the realization *that* the implied variance should have been σ_R instead of σ_I.

The variance swap as constant dollar gamma

One author who bases his argument on realized variance and options gammas rather than implied variance and options vegas is Bossu. In his world too, realized variance is the only true variance and it is objectively known, at least to the author himself, from the initiation of the trade. However the options market-maker knows only implied variance and he will consistently *ignore* true realized variance until expiration. In this case there will be no rewinding back as in Demeterfi's. Bossu has no choice but to unfold the *real process* until the end and to express the realized P&L of the delta-hedged vanilla option at the expiration of the trade, not at its initiation. Although not expressly stated by Bossu, delta-hedging is accomplished using *implied variance* in the formula of the Black-Scholes delta. Ahmad and Wilmott are more exacting about this in their article about delta-hedging[6], and they essentially produce the same result as Bossu, namely:

$$\text{P\&L} = \frac{1}{2}\left(\sigma_R^2 - \sigma_I^2\right)\int_0^T S_t^2 \Gamma_I dt$$

where Γ_I is the option gamma using implied variance.

Bossu then observes that the expression of the final P&L is very similar to the payoff of the variance swap above except for the weighting $S_t^2\Gamma_I$, known as *dollar gamma*, which makes it highly path-dependent. In a vein similar to Demeterfi et al., Bossu then looks for the combination of vanilla options whose dollar gamma would be independent of the stock price, and he finds the same portfolio Π.

Calling α_I this constant dollar gamma (where the subscript I is here to remind us that the gamma is computed using implied variance) the final P&L of the hedged portfolio reduces to:

$$\text{P\&L} = \frac{\alpha_I T}{2}\left(\sigma_R^2 - \sigma_I^2\right)$$

As far as the difference of variances is concerned, this is the same expression as Demeterfi's. However the meanings of the terms are different. Recall that all delta-hedging is supposed to be performed under realized variance in Demeterfi's case, and that our interpretation of σ_R is the *implied variance* that would have had to be used in the initial pricing of the portfolio Π, had realized variance been known. By contrast, σ_R is the realized variance in Bossu's case, only the delta-hedger doesn't have knowledge of it. As a matter of fact, he computes the delta and the gamma of the portfolio Π using implied variance.

Implied variance vs. realized variance

There is a tension, of course, in wanting to remain within the Black-Scholes model (which both authors insist serves pedagogical and intuitive purposes) and wanting realized variance and implied variance to be different. To repeat, the Black-Scholes model is blind to such a distinction, and in fact the three main components of P&L of the hedged portfolio in the Black-Scholes derivation, namely the option value, its delta and the return of the underlying presuppose the same variance number. When authors like Ahmad and Wilmott, Bossu, and indeed Demeterfi et al., start writing mathematical expressions (if only to express realized P&L) where realized variance is a different number

than implied, you may legitimately wonder what their model of realized variance is. In their pedagogical examples, clearly their model is that realized variance is given and is constant, only it is different from what somebody, out there, is using as implied variance, either to price an option, or to both price an option and hedge it. But then it is a *theory* of implied variance that we would be missing. Is implied variance supposed to be the number we get by inverting our model against the options market price? This is not a slight suggestion and it has many consequences and ramifications that I have explored in an article published in the January 2006 issue of the Wilmott magazine[7]. Or maybe implied variance is just a number without origin or consequence, serving no other purpose than being different from realized variance. If so, it will not resist the tension that I mention above. Indeed you wonder: If the author of the exercise has a model for realized variance, what is to stop the option market-maker and the delta-hedger from having it too inside the exercise?

While certainly useful as a mathematical exercise ("Compute the P&L of a Black-Scholes delta-hedger who is using a different variance number in his option pricing and hedging formulas than we, the authors of the exercise, know to be the true variance") such elaborations fall short of answering big, significant, practical questions such as: "How to delta-hedge when neither the delta-hedger nor the author of the exercise know true variance, or even the process of true variance, and the only observables are the traded prices of both the underlying and the derivative?" This is a criticism I have already addressed to Ahmad and Wilmott in my article. Together with the ramifications of the concept of implied variance, it led me to a reformulation of the whole logic of option pricing and trading, where recalibration and expansion of the state space are recognized as the key phenomena.

Variance swaps and the intuitive approaches to valuing them that are proposed by Demeterfi et al. and Bossu offer me a chance to revisit this criticism. As a matter of fact, Demeterfi's approach (hedging with realized variance) and Bossu's approach (hedging with implied) reflect the exact same alternatives as proposed by Ahmad and Wilmott. The novelty here is just that what used to be a hedging issue in the hands of Ahmad and Wilmott becomes a *pricing* issue with the variance swap. This is why I maintain that the variance swap is a higher order instrument (a produce of theory, the target of irony) and that the approach to pricing it is in itself equivalent to a philosophical criticism.

If the purpose of the exercise is to show that variance swaps are pure bets on variance which do not require delta-hedging, why not stick with their initial contractual terms, the average of squared logarithms of returns? This payoff is the literal expression of realized variance and so long as the underlying process is homogeneous and geometrical (e.g. geometrical jump-diffusion, with stochastic volatility or anything you like), the returns will be independent of the stock price and the payoff of the variance swap will require no delta-hedging. Why take the trouble of the puzzling diversion into options vegas and gammas and realized variance versus implied?

The answer is that Demeterfi et al. and Bossu think and write in keeping with the Black-Scholes tradition which, like I said, has invented the notion of implied variance as opposed to realized, and of delta-hedging as a bet on variance. This is all predicated on diffusion. The bet they have in mind is a bet on diffusion variance, not on a variance whose main component may be the jumps. So they have to think of the variance swap as an instrument addressing

the needs of a delta-hedger who no longer wants his P&L to be path-dependent, or a vega trader who no longer wants his exposure to depend on stock. As a matter of fact, when it is stripped of the muddled issue of realized versus implied variance, the intuitive elaboration of Demeterfi et al. and Bossu can be interpreted as the answer to the following question: What European derivative has a vega and a dollar gamma that are independent of the stock price? This is a purely mathematical question and it is rightly posed within the confines of Black-Scholes. The answer is the log contract. (Only in their second, theoretically rigorous, section will Demeterfi and co-authors independently show that in a diffusion framework more general than Black-Scholes, the payoff of the log contract replicates the variance that is realized along the path. Their argument will be theoretical through and through and it will never leave the object level. As expected, they will thenceforth make no mention of implied variance or of the trading of variance.)

The log contract

If I were to reconstruct Demeterfi et al. and Bossu on a sound methodological basis, I would, therefore, adopt the totally symmetrical point of view of Neuberger. Neuberger is not concerned with the pricing of variance swaps *per se* or their replication. He motivates his article by first arguing that "when a modern institutional investor buys an option, he is buying an exposure to all the Greek letters." He then observes that the Greeks of the traditional vanilla option may not the best vehicle to offer the desired exposure. For instance, its delta and gamma change over time at rates highly dependent on the relative values of the spot price and the strike price, so the investor would have to worry about the timing of his trade or the passage of time on top of worrying about the stock price movement. By contrast, the delta and gamma of the log contract depend only on the stock price and are stable over time. As time and volatility are always coupled in the pricing equations, this means the delta and gamma of the log contract do not depend on volatility!

As a matter of fact, when the underlying process is homogeneous in the logarithm of the underlying (as is the case with all the diffusion or jump-diffusion models we are considering), the risk-neutral expectation of the logarithm of the return over any period of time will only depend on the time to maturity. As a consequence, the expectation of a logarithmic payoff occurring at maturity will only depend on time and the logarithm of the initial spot price. The delta-hedging ratio of the log contract is therefore always equal to $1/S$ quite independently of the stochastic nature of the diffusion coefficient, or the existence of jumps, or whether the delta-hedger knows this to be the case or not to be the case.

To my mind, this is the key observation which liberates us from the whole vexed issue of implied variance versus realized and opens the road for the log contract as the true vehicle of *realized* variance. In retrospect, it will justify Demeterfi's and Bossu's intuitions about variance trading as it suffices to recognize that all that was bothering us back then was the necessity to step outside Black-Scholes (by the very meaning of "variance trading") while artificially maintaining Black-Scholes in the mind of the delta-hedger, *who could only use implied variance in his hedging formulas*. As a matter of fact, my criticism of Ahmad and Wilmott also revolved around the methodological

unease that was caused by the fact of "knowing" realized variance for the sake of the argument and "knowing" only implied variance for the sake of hedging. This criticism doesn't stand for the log contract. But then the article of Ahmad and Wilmott wouldn't stand either, as there would be no question left concerning the volatility to use in hedging!

To show that the log contract captures realized variance, I will slightly adapt Neuberger's derivation in the following manner.

Holding an amount $\frac{1}{S_t}$ of underlying, at discrete time t, against his short in the log contract, the writer of the delta-hedged log contract will incur, at maturity, the following P&L (assuming the interest rate is zero):

$$\text{P\&L} = L_0 - \log S_T + \sum_{t=0}^{T-1} \left(\frac{S_{t+1}}{S_t} - 1 \right)$$

where L_0 is the value of the log contract at time $t = 0$.

This can be written as:

$$\text{P\&L} = L_0 - \log S_T + \sum_{t=0}^{T-1} \left(\frac{S_{t+1}}{S_t} - 1 - \log \frac{S_{t+1}}{S_t} \right) + \log S_T - \log S_0$$

Noting $y_t = \log \left(\frac{S_{t+1}}{S_t} \right)$ and using the Taylor expansion of the exponential function:

$$\text{P\&L} = L_0 - \log S_0 + \sum_{t=0}^{T-1} \left(\frac{1}{6} y_t^3 + \frac{1}{24} y_t^4 + \dots \right) + \frac{1}{2} \sum_{t=0}^{T-1} y_t^2$$

Note that the realized variance:

$$\sigma_R^2 = \frac{1}{T} \sum_{t=0}^{T-1} y_t^2$$

can here exactly be identified with the payoff of the variance swap, independently of any assumption about the data generating process that is producing S_t (is it a diffusion? a jump process?), or whether it even admits of finite moments. We are able to achieve this generality because the delta-hedging strategy of the log contract is independent of the assumption of diffusion or jumps. All we need is space homogeneity. How implied variance eventually re-enters the picture is by noting that if, at initiation, the log contract writer had priced it using Black-Scholes, then the initial price L_0 would have been equal to:

$$L_0 = \log S_0 - \frac{1}{2} \sigma_I^2 T$$

by solving the Black-Scholes PDE for the log contract.

This is no more than a re-labelling, a mere alternative pricing representation. The implied variance σ_I is just a name and has no effective meaning as it doesn't intervene in the hedging. Rearranging the terms, the P&L can now be expressed as follows:

$$\text{P\&L} = \frac{T}{2} \left(\sigma_R^2 - \sigma_I^2 \right) + \sum_{t=0}^{T-1} \left(\frac{1}{6} y_t^3 + \frac{1}{24} y_t^4 + \dots \right)$$

and compared to the expression given by Demeterfi et al. or Bossu.

The rest of the argument consists in noting that the terms under the summation symbol are negligible under diffusion. But be aware that they are not so negligible and may even become dominant in the presence of jumps of substantial size! This concludes the proof that the log contract is the replica of the variance swap in the absence of jumps. The remainder of the traditional paper is usually dedicated to the static replication of the log contract by the vanillas. Instead of shorting a log contract against your writing a variance swap you are advised to buy the weighted combination of vanilla options computed above. However, this is only a geometrical exercise (how to approach a continuously curved payoff with a bunch of sharp-cornered ones) and has no financial bearing.

The so-called model-independent valuation of the variance swap

Yet, for some reason, people are over-excited about this result. They think the options market contains information about the expected realized variance and, what's more, completely independently of any model, as it takes only the market prices of the vanillas, and no option pricing model, to compute the price of the log contract that they statically replicate. Some options exchanges, Bossu reports, have even launched volatility indices following the weighting methodology of the replicating portfolio, for instance the new Chicago Board Options Exchange SPX Volatility Index (VIX) and the Deutsche Börse VSTOXX Volatility Index.

To my mind, the only thing the log contract has accomplished is push the boundary of model-dependency a little further, with the zest of irony that you should always expect when theory gets mixed up with the market. (Ironically, the thought here is that the market knows variance independently of theory.) And I would re-describe the remarkable fact of model-independence as follows.

From the market price of a single option you used to imply volatility (or variance) and everyone agreed this was model-dependent as it depended on the Black-Scholes model. As a matter of fact, what you implied was instantaneous variance, the one required in the infinitesimal re-hedging of the Black-Scholes arbitrage portfolio. How the log contract changes this picture is that, thanks to the homogeneity of its payoff (no singularity in space localized at the strike price) and to the consequence that its delta and gamma no longer depend on time (constant dollar gamma), what you can imply from its re-hedging is not something local, instantaneous variance, but something global, variance realized over the whole path, and what you can fit in this interval is not just diffusion with constant variance (the Black-Scholes model), but more generally diffusion with stochastic variance. From the price of the log contract you can therefore imply the expected realized variance *when no jumps enter in that realization*. True this sounds model-independent, however it is identical with the good old, model-dependent, Black-Scholes implied variance. The only difference is that the inference is now generalized to the class of models falling under the umbrella of stochastic diffusion.

Delocalizing, or homogenizing, the sharp-cornered payoff of the vanilla option (thus turning it into the log contract) comes at a price and that

is that the strike gets spread out too and must now span a whole range of values. This is the essence of the static replication. While we do this, and our class of models is concomitantly extended from constant to stochastic variance, we won't be surprised if the vanillas that we require come out smiled relative to Black-Scholes. Their prices are "model-independent" only to the extent that they are tautologically given (supposedly by the market).

We can equally argue that the price of the single option in the Black-Scholes setting was model-independent too. What was model-dependent is the belief that this alone could impose instantaneous variance. Likewise, what is model-dependent here is the belief that the bunch of vanillas replicating the log contract can impose the expected realized variance. In reality they can't, because of the jumps that may stand in the way. No wonder that the market should price expected realized variance through the *real* variance swap and nothing else and that it, in turn, should be valued by an ironist model specifically with jumps!

In sum, inferring "Black-Scholes implied volatility'' from the price of a single option is no more model-dependent than inferring the "diffusion implied expected realized volatility" from the price of a bunch of options. Ironically, what the joy of quoting "model-independent" variance swaps prices has created (as well as the joy of trading them in volumes potentially as liquid as the replicating vanillas) is the perfect opportunity for the jumps in the not so diffusive equity markets to come jumping to the fore!

FOOTNOTES

1. Sebastian Bossu, "Introduction to Variance Swaps," Wilmott magazine, March 2006.
2. Kresimir Demeterfi, Emanuel Derman, Michael Kamal, Joseph Zou, "More Than You Ever Wanted To Know About Volatility Swaps," Goldman Sachs Quantitative Strategies Research Notes, March 1999.
3. Anthony Neuberger, "The Log Contract and Other Power Contracts," in The Handbook of Exotic Options, I. Nelken, editor, Irwin 1996, 200–212. Neuberger defines the log contract as "a European-style contingent claim whose value at maturity is equal to the natural logarithm of the price of the underlying asset." Another reason for the log contract's lack of liquidity may be that its payoff unrealistically blows up to infinity when the underlying stock goes to zero. While certainly a nice theoretical construct which, as we shall see, addresses variance replication, the log contract, when pushed to the wrong extreme, starts to look again like the mathematical fantasy that it has always been, quite removed from the meaningful financial contracts that trade in the real market.
4. "What's on a Man's Mind: Elie Ayache of ITO 33 unveils *OpSession*: A New Logic for Option Pricing and Trading,'' Wilmott magazine, May 2006.
5. Richard Rorty, *Contingency, Irony, and Solidarity*, Cambridge University Press 1989, 73–74.
6. Riaz Ahmad and Paul Wilmott, "Which Free Lunch Would You Like Today, Sir?: Delta Hedging, Volatility Arbitrage and Optimal Portfolios," Wilmott magazine, November 2005.
7. Elie Ayache, "What Is Implied by Implied Volatility?", Wilmott magazine, January 2006.