Actuarial Value vs. Financial Price

Financial pricing, in its usage of the term 'real probability', has not managed to unshackle itself completely from actuarial valuation. Even though the market of contingent claims is the long-awaited technology that would finally allow us to *account* for future contingent events without using probability . . .

In Salih Neftci's *An Introduction to the Mathematics of Financial Derivatives*,[1] one can read:

> *The Girsanov theorem provides the general framework for transforming one probability measure into another 'equivalent' measure....The probabilities so transformed are called 'equivalent' because....they assign positive probabilities to the same domains.* (p. 322)

In the following chapter on the applications of equivalent measures, Neftci writes:

> *In this chapter, we show how the method of equivalent martingale measures can be applied. We use option pricing to do this. We know that there are two ways of calculating the arbitrage-free price of a European call option:*
> *1. The original Black–Scholes approach, where a riskless portfolio is formed and a partial differential equation is obtained.*
> *2. The martingale methods, where one finds a 'synthetic' probability under which the stock price process becomes a martingale.* (p. 345)

The martingale method is presumably the one that does not rely on dynamic hedging and invokes only the combination of no-arbitrage (this is what the term 'martingale' refers to) and the equivalence between the 'synthetic' probability (or changed, risk-neutral measure) and the original real one. Neftci provides the details of the derivation of option prices following the two routes. Crucially, the pricing PDE in route (2) is not obtained as in route (1), via the no-arbitrage argument imposing that the dynamically hedged portfolio should only earn the interest rate, but first, in transforming the price process of the underlying into a martingale through Girsanov's theorem then, in expressing the price process of the derivative by using Ito's Lemma, and finally, in insisting that the price process of the derivative should also be a martingale in the changed measure. To achieve the latter, the drift term of the SDE ruling the derivative is set equal to zero and this yields the same PDE as in Black–Scholes' original derivation. Neftci concludes:

> *It was shown that the martingale approach implies the same PDEs utilized by the PDE methodology [by this, Neftci means the traditional BSM approach through dynamic hedging]. The difference is that, in the martingale approach, the PDE is a consequence of risk-neutral asset pricing, whereas in the [original BSM] method, one begins with the PDE to obtain risk-free prices.* (p. 366)

Real probability measure vs. changed measure

This is all very well and I certainly believe that measure theory, Girsanov's theorem, continuous-time stochastic processes, etc. are well established and that they all exist, because they are only mathematics. On the other hand, I also believe statistics exist and I believe insurance companies do break even on average, or at least that the problem they face – that of trying to break even on average when faced with their statistical populations – is well posed. But to go back to the question of the title, about the difference between actuarial valuation and financial pricing, I ask: Why insist on calling the risk-neutral measure *equivalent* to the 'original real' one? Why even suppose that the risk-neutral measure, or the pricing operator one uses to generate arbitrage-free derivatives prices, is the result of changing the 'original real' measure in which the underlying is supposed to exhibit its real, historical statistical distribution? What if there was no such thing as the 'real probability measure?' Not that it should be unobservable or inscrutable; no, my problem is that we should find no precise meaning, but only muddled conceptions of what the real probability means.

For let us not fool ourselves, if one wishes to read into the words 'real probability distribution' something that goes beyond the mere formalism of Girsanov's theorem, which merely treats measures symmetrically and does not know what 'real' means, this 'real' distribution of the underlying that everybody is talking about has to be the one that some actuary is reading for me from the past statistics of the underlying, not the forward-looking one (whatever that means too) with which the trader is supposed to price derivatives.

It is perfectly fine to identify the probability of an event that has been recognized, or modeled, or idealized, as a member in a statistical series, with the frequency of its occurrence in the series. In that case, the word 'probability' would just be a rewording and would have no ontological implication. However, the trouble begins when you try to make sense of the probability of that single case or single event, independently of any reference to the whole series.

So again, I ask: Armed with your favorite concept of frequency-based probability, when you are squeezed in that corner facing the next draw – the next single occurrence of the event and nothing but – what exactly do you mean when you say its probability of occurring is p? Is p really its probability or does it, once again, refer to the whole series?

To repeat, this is not a problem of knowledge. Indeed, you might be prevented from knowing the frequency of the event in the series; or it might be that the series itself and the corresponding repetition of the draw are only thought-experiments whose outcomes you may speculate upon but never know. The problem I am posing is a *problem of reference*, i.e., an ontological or even logico-semantic problem. Does this probability p *truly* refer to the single event (and by 'truly' I mean: according to your own system of thought and metaphysics; according to your own usage of precise language), or does it always have, as Richard von Mises prescribes,[2] to presuppose the whole collective?

Ex-ante vs. ex-post

The whole philosophical problem of single-case objective probability can be rephrased as the problem of shifting from an ex-post stance to an ex-ante stance. If your answer to the question above is that the probability p is deemed the 'probability of the single event' only insofar as the event would display a frequency p of occurring in case the draw was suitably repeated in a series – even an ideal series that didn't exist empirically but was only a thought-experiment – then your stance will be ex-post, despite the fact that you seem to be addressing the event before it occurs when you so answer me. To really move to the ex-ante stance, you have to tell me something about the event literally *before it occurs* (this is what 'ex-ante' means literally), let alone before the whole series, possibly involving it, unfolds.

So when I say that the market price of the contingent claim exists whereas the objective probability of the event triggering it doesn't, I speak ex-ante, as this is what 'to exist' really means. To repeat, the market can give me the ex-ante price of a contingent claim whose triggering event is genuinely single-case (as this is what the market is supposed to do), while I doubt that we could ever spell out the objective probability of this single event without smuggling in statistics implicitly, therefore the ex-post stance.

Actuarial valuation vs. financial pricing

This is why I contend that financial pricing, in its usage of the term 'real probability,' has not managed to unshackle itself completely from actuarial valuation. Even though the market of contingent claims is the long-awaited technology that would finally allow us to *account* for future contingent events without using probability, simply by assigning a tradable price to the corresponding contingent claim in its market – even and especially when the triggering event is one of a kind, i.e., an event that never was and never will be a member of a statistical series (typically the default event of a corporation) – the textbook presentation of financial derivative pricing is still obsessed, or at least burdened, with the legacy of actuarial science. How? Simply when it argues (for instance, Neftci p. 319) that 'on average, the risky asset will appreciate faster than the growth of a risk-free investment' and for this reason its drift has to be changed through changing the probability measure, if it is to become a martingale.

What does 'on average' refer to, in Neftci's statement, other than a hypothetical insurance company that is supposed to hold a population of such assets and would presumably make more money by holding them than by investing risk-free (probably what AIG had in mind when they started holding CDSs)? In that case, 'average' would mean that the particular asset we are talking about is the 'average asset' of the population, or its representative. But what if there was no such population and the asset was one of a kind? Probably the insurance company would have to invent such a population and hold assets belonging, say, to the same sector, or something like that. Crucially, probability is dependent on the ability to count and to measure frequencies. But what if the event was uncountable because it was unique? Presumably, some subjective elements would have to enter into play and the unique event, or the unique asset, would have to be modeled, that is to say forced, to be a member of that statistical series or that reference class.

Or is 'on average' supposed to mean 'in the long run'? In other words, you are supposed to continue holding this single asset (and not a population thereof) while all the different rise and fall scenarios that history has in store for it unfold, until you observe – by virtue of some ergodic theorem guaranteeing that all the possible paths that are open in space for that asset will eventually unfold in time – that you have made money overall. But what if the market was a single run and not a long run of runs? What if the market was precisely like history, something that happens once and never repeats itself or gives you a second chance?

The self-sufficiency of risk-neutral pricing

It seems to me that the only reason why we need martingales in finance is to express the prices of contingent claims generally as the discounted expectation of their payoff, thus making sure we observe non-arbitrage among their instant prices. "In the absence of arbitrage possibilities, market equilibrium suggests that we can find a synthetic probability distribution such that all properly discounted asset prices behave as martingales. Because of this, martingales have a fundamental role to play in practical asset pricing," writes Neftci (p. 124). I don't know what else than 'practical asset pricing' there could be. 'Theoretical asset pricing' perhaps? Better to use the word 'valuation' instead of 'pricing' in that case. No-arbitrage is the only real (or realistic) constraint. It closes itself off to the instant market and bears no relation to an outside or to a long run, or even to the future distribution of returns of the assets. For this reason, it seems to me that the whole elaboration in terms of changing the measure and equivalence with the real measure is just lip service to the actuarial ancestry.

Characteristically, in the one passage where Neftci speaks of the real rate of return R of the risky asset, he writes:"Now consider the problem of a financial analyst who wants to obtain the fair market value of this asset today". One way to do this, Neftci suggests, is to compute the present value as the mathematical expectation of the future returns under the real probability measure: $S_t = E[S_{t+1}/(1 + R)]$. However, notes Neftci, this requires a knowledge of the distribution of R, which requires knowing the risk premium of the asset. Neftci then observes that 'knowing the risk premium before knowing the fair market value S_t is rare' (p. 320). Note the irony.

Next Neftci offers to change the probability measure in order to get rid of the risk premium in the computations. Under this risk-neutral measure, the *present fair value* our analyst is looking for would simply be the mathematical expectation of the future market prices discounted by the interest rate. All that remains to do then is to forecast the *future market* prices S_{t+1}. This can be done, writes Neftci, probably without noticing the irony of the reversal, by "using a model that describes the dynamics of S_t and then discounting the 'average forecast' by the (known) r". In short, all an analyst has to do in order to estimate the (fair? real? fundamental?) present value of an asset is to *forecast its future* market prices. Why not simply say that its present value is equal to its present price? As a matter of fact, Neftci's notation is equivocal on this point as he speaks of forecasting S_{t+1} (supposedly the future prices) using a model for the dynamics of S_t (supposedly the present value we were after, now imperceptibly confused with the present price).

And when, in a later section, Neftci concludes that 'the synthetic probabilities [or risk-neutral measure] appear central to pricing of financial securities' and wonders where they can be got from, he suggests that the volatility parameters of the stochastic differential equation should be calibrated by the practitioner, 'based on the existence of liquid options ... that provide direct volatility quotes' (pp. 334-335). This completes the proof that the practice of derivative pricing, that is to say, of financial pricing as a whole, is totally impervious to the real probability (whatever that means) and its actuarial underpinnings. Yet you wonder: What might be the implication of this on probability and its understanding? For, surely, the textbooks of financial derivatives will never dispense with the concept of probability? It is one thing to suddenly wake up in a market where derivatives trade liquidly alongside their underlying and can be used as inputs in the pricing models; it is another thing to initiate the market on that road.

We have accustomed ourselves so well to the theoretical notion of 'random generator,' also our computerized Monte Carlo simulations seem to materialize it so well, that no one really wonders, in finance, what this means philosophically or at least semantically that we should give ourselves such a generator and write such a thing as a stochastic process

Philosophical theories of probability

The philosophical interpretation of probability has a long history and a very thick literature. It ranges from probability being only a shorthand for statistical frequency, in other words, an essentially *ex-post concept* (Richard von Mises), to probability as a propensity, i.e., an *ex-ante concept* that specifically concerns the single event and does not depend on the existence of a whole statistical series or reference class (this was advocated by Karl Popper,[3] after quantum mechanics had pressed the case of irreducible randomness that didn't relate to statistics but seemed to suggest a random generator truly at work behind each individual experiment), to subjective probability, of course, where probability is identified with the betting odds that some agent would produce (Bruno de Finetti[4]).

It is not my purpose to rehearse this debate through the very same examples that all those authors have used (usually dice, roulette wheels, mortality tables,...and quantum mechanics), but to see whether the derivatives market could not offer a fresh perspective. Note that, apart from revealing irreducible randomness in nature, quantum mechanics quickly posed deeper problems, such as what we meant by object, or property, or physical state, or even identity of the particles, etc. In itself, the algorithm for computing quantum probabilities posed no particular problem (the Born rule). As for the reason why such a random generator existed in nature, the most advanced interpretations ended up suggesting that it was not probability that we were ultimately talking about in quantum mechanics, but something else. According to Jan von Plato,[5] it is statistical physics that contributed all the tools of modern probability theory in its most advanced branches, namely stochastic processes and calculus, yet, he complains, it never was prominent in the philosophical debate of probability because it was eclipsed by quantum mechanics. (Chaos theory and chaotic determinism are a different subject, of course.) My observation is that derivative pricing is today the most advanced branch and probably for this reason: its contribution to the philosophical debate of probability is even smaller than statistical physics, not to say totally non-existent. Blame it on the excessive sophistication of the mathematics. We have accustomed ourselves so well to the theoretical notion of 'random generator,' also our computerized Monte Carlo simulations seem to materialize it so well, that no one really wonders, in finance, what this means philosophically or at least semantically that we should give ourselves such a generator and write such a thing as a stochastic process.

My contention is that, because derivatives markets are so real and so indisputable today, because prices exist *materially* and the probability of a single-case event doesn't exist or at least is still problematic, maybe we could follow through the logic of the markets and establish a link between the contingent event and the price of the corresponding contingent claim without the intermediation of probability. Note that, up to this day, it is not clear yet what Popper's propensity means. Nobody knows what it means really that a coin tossing experiment, in and by itself, should have a 0.5 propensity – Popper also calls it a 'dispositional property' – of producing heads or tails. Popper insists that propensity is real and present in the situation as a '*generating condition*' (same word as 'random generator'), i.e., it is truly ex-ante, as opposed to being just a nominalistic rephrasing of an ex-post frequency. He writes: "Like all dispositional properties, propensities exhibit a certain similarity to Aristotelian potentialities. But…they cannot…be inherent in the individual things. They are not properties inherent in the die, or in the penny, but in something a little more abstract, even though physically real: they are relational properties of the total objective situation" (p. 359).

Apart from the fact that Popper insists it really exists, this doesn't tell us, of course, what propensity really is. Popper himself closes the discussion by arguing that propensity is in the end 'what corresponds to the transition from the mathematical frequency theory of von Mises to the neoclassical or measure-theoretical treatment of probability' (p. 360). In other words, Popper evades the debate through the point I have noted above, through the sophistication of mathematics being ultimately the only true thing. He notes that measure theory 'is superior to the frequency theory, not only from a philosophical point of view but also from a purely mathematical point of view' (ibid). Why superior? Presumably because, as Popper writes, "the neoclassical theory does not attempt to give a definition of 'probability,' either on the lines of Laplace or of von Mises…Instead it takes 'probability' as anything that satisfies the rules of certain calculus…" It clearly separates the formal task of constructing a mathematical calculus of probability from the task of interpreting this calculus…' (pp. 374–375). In short, the real place of the random generator lies in mathematics not in physics and the ex-ante notion of objective probability, or propensity, remains undefined and unexplained – a lack that Popper now ironically recognizes as a philosophical superiority.

Money and break-even as primitive concepts

This leaves as the only meaningful objective probability the ex-post concept of statistical average, which depends on counting a population and cannot be single-case, and as the only meaningful ex-ante concept of probability de Finetti's subjective probability.

My observation is that both depend on money and on the existence of some financial account and that they don't stand by themselves. Subjective probability is obviously financial, because de Finetti explicitly equates it with the betting odds that a banker is supposed to quote for you. Thus, de Finetti had in mind a transaction and a price, and he fell one step short of the market of contingent claims. I think the reason why he did not fully embrace the market was that he was still keen on defining probability and that probability had to reside in the mind of a subject if it was found not to reside in nature or in some object. For surely, it could not reside in the mind of the market! As for the statistical probability, I claim it is related to money too, because in this case, the account in question is of course the insurance company's.

It is all well to define statistical probability as the limiting frequency of a certain occurrence in von Mises' collectives. However, I believe the real operational concept is that of breaking even in the long run, when somebody plays that dice, or plays those mortality tables. Characteristically, for von Mises' statistical probability to make sense, the series it is measured upon has to be 'truly random.' (Imagine that some demon is systematically drawing a series that didn't reflect the 'true' probability, say, an indefinite series of 'heads.') And how does von Mises avoid the circularity of defining 'truly random' when probability is not yet defined? By arguing that a truly random sequence is one that would be immune to gambling systems. In other words, a trading concept, or generally an accounting argument, lies at the basis of von Mises' whole edifice!

For this reason, I wish to argue that statistical regularity and the corresponding break-even in the long run are the primitive concepts and that probability, if you insist we should consider it at all, is only a derivative concept. Von Mises claims that his statistical probability is not definable without reference to the whole series or population. I elaborate this by saying that statistical probability – or actuarial probability – is not definable without prior reference to the ex-post accounting equation of the insurance company. Only because it has broken even on average, admittedly after a long history of trial and error and adjustments of the insurance premium, can the insurance company later turn back to the single case and form such a concept as the probability of death of that particular individual.[6]

What I am saying is that, appearances to the contrary, the integral comes before the integrand. First you sum up all the cases, and then you work out probability as the frequency. I am proposing that you went one step further and integrated the probability, not just against the indicating function of the event of death, but, more realistically, against the money paid in case of death. Defining probability as the limiting frequency is not enough and is not the real thing. The real thing is that only the person (either physical or moral) who has broken even on average relative to the given statistics and population can turn back and speak of the probability of the single occurrence. There would be no metaphysical coup de force in this but only a different way of slicing the account. It is not against the current of time that one should navigate in order to switch from ex-post to ex-ante, but against the current of money.

I guess the reason why I insist that time should be replaced by money is that what bothers me in probability is the element of time and the time connotation of terms like 'to expect,' 'to predict,' etc. What bothers me is that we should wait for the event to happen and wonder, in the meantime, what its probability might be. I say we should wait in money, not wait in time, because what is accountable is money. In both the cases of subjective and objective probability, we accounted for the event; we didn't value it. Also, this financial underpinning of probability will allow me to argue, with all the greater force, that indeed prices of contingent claims exist and probabilities don't. For we are still missing the one configuration in which probability could be said really to exist (thereby vindicating the claim of a metaphysical realist), namely, a meaningful ex-ante concept of objective probability. We are still missing an objective concept of single-case probability – probability ▷

attaching to a singular event that is part of no population.

The market as a fundamental category

But why speak of the 'probability' of this event? Why don't we just say that we wish to account for it?

Here is my proposition: Instead of going from the derivative concept – probability – and wondering how this probability could be adapted to the genuine single case for which there is no population, no statistics, and no breaking even of an insurance company, why not branch off at the earlier step and see how the concept of break-even itself can be adapted to the single case? In other words, I wish to create the conditions of break-even for the contingent claim that is written on a singular event before even the notion of probability is constructed – for probability is a derivative concept, to repeat, and is only relative to the break-even of an insurance company over a whole population. And how do we break even when we hold a singular contingent claim? Obviously not by waiting for the long run as there is none. We break even by not waiting, by simply making sure that we could liquidate our holdings at once if we wished; in other words, by creating a market for them. Conversely, if such a liquid market existed, would we be satisfied 'valuing' our holdings other than by marking-them-to-market?

Now we see that our whole predicament comes from our unwillingness to recognize in the market a fundamental category. It is not a slight thing to have invented money, contingent claims, and the market place where they are exchanged. Exchanging is a fundamental invention. If probability is the concept that was invented precisely to suit the insurance company, in those situations where the statistical regularity (an indisputable law of nature), combined with the integrity of the account of the insurance company, created the time loop in which it looked as if the event could be addressed ex-ante and as if meaning could be given to its probability, we should look for an overall alternative to probability when there is no such statistics and no insurance company. Instead of changing the probability measure, we should change the whole concept of probability.

Why indeed doesn't the market become a pricing theory of its own, THE pricing theory? Why do we need a theory for the market? Is it because the market is complex and we need to model it? Well, I say the market is simple, not complex. Just forget the crowd that constitutes it. Simply, the market is what gives the price (and the price process) of contingent claims. And if you're not happy calling the market a theory, then just call it a technology

We are so entrapped in probability that the actuarial value of the contingent claim seems indelible from our minds. We prefer to imagine that a certain asset first admits of a value – even on pains of having artificially to create the population of which its one-time payoff would be a member – second that the market is in charge of altering or changing this value by the play of supply and demand or the fact that some players won't be content to break even on average, instead of accepting that the asset has no value but only a price through the exchange. To us, exchange can only mean price change.

Admittedly, the major conversion I am proposing, in which price absolutely replaces probability, leaves us with the later impossibility of modeling the dynamics of price. For how could we model its dynamics except through probability? But do we really need to model it, now that it is given by the market? Isn't the rule precisely constantly to recalibrate our models of the underlying dynamics to the market prices of derivatives? Better: are our derivative pricing models really models of the underlying dynamics or just risk-neutral pricing operators that allow us to capture a semblance of consistency between the instant prices of derivatives, with no idea of what will happen next apart from recalibration to the market update? Why indeed doesn't the market become a pricing theory of its own, THE pricing theory? Why do we need a theory for the market? Is it because the market is complex and we need to model it? Well, I say the market is simple, not complex. Just forget the crowd that constitutes it. Simply, the market is what gives the price (and the price process) of contingent claims. And if you're not happy calling the market a theory, then just call it a technology.

ENDNOTES

1. Neftci, S.N. 2000. *An Introduction to the Mathematics of Financial Derivatives*. Academic Press.
2. von Mises, R. 1957. *Probability, Statistics and Truth*, 2nd revised English edition. Allen and Urwin.
3. Popper, K.R. 1983. *Realism and the Aim of Science.* Routledge.
4. de Finetti, B. 1974. *Theory of Probability*, Vols. 1 and 2, English translation. Wiley.
5. von Plato, J. 1994. *Creating Modern Probability*. Cambridge University Press.
6. Ayache, E. 2011. Single-case statistics? *Wilmott* magazine. March, 12–14.