2.3.2 The table of probabilistic scenarios

The fair value is, by definition, a synthetic value of the risk-neutral density of the possible final payoffs; therefore, it ignores information provided by moments of order higher than one and it does not

Figure 2.1 Comparison of the densities of two non-equity products with the same fair value

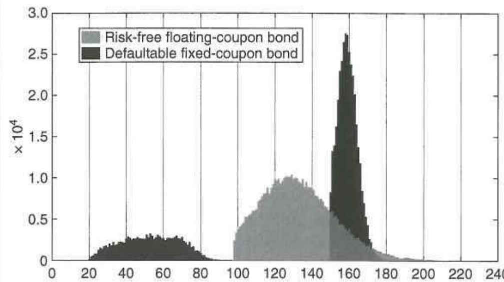

allow for the appreciation of the degree of randomness characterising the performances of a given product. In fact, the same discounted expected value may be obtained by an infinite number of density functions, even with very different shapes. These shapes completely qualify the peculiar relation between riskiness and profitability of any single product that represents the fundamental information that investors take as a reference point to assess whether an investment proposal effectively meets their needs.

The inherent conciseness of the fair value becomes particularly relevant when the risk neutral density of the product is quite irregular. Investors naturally tend to figure out the final performance by using, as its proxy, the fair value compounded at the risk-free rate until the expiry of the time horizon. Such reasoning is somewhat meaningful when the density is regular, otherwise it is biased since it implicitly assumes that an entire distribution can be exhaustively depicted by only considering the first moment, which is generally false.

Figure 2.1 clarifies this point by comparing the risk-neutral densities of two non-equity products with identical fair values (both equal to 100) and expiry dates but characterised by very heterogeneous risk–return profiles. The first security is a risk-free floating-coupon bond, while the second is a subordinated bond which pays "fair" fixed coupons, meaning that it gives an extra return over the risk-free rate in line with the one required by the market to compensate for the underlying credit risk exposure.

The difference between the two densities is striking and it is only partly signalled by the fact that the risky component (see Section 2.3.1) is 0 for the first product and 15.7 for the second one. In particular, the former density is fully contained in the region of payoffs higher than 100, meaning that the buyer of the security will achieve only positive returns. On the other hand, the other density is bimodal, as expected. Consequently, should some credit event occur over the life of the product, investors would receive only the recovery value (set equal to 20% of the notional capital) compounded at the risk-free rate until maturity, which explains the dispersion above the minimum threshold of 20 observed in the region of payoffs lower than 100. Otherwise, if no credit event occurs, then investors will benefit from positive returns that on average are higher than returns of the risk-free asset.

Comparing the two densities, it is self-evident that, despite having the same fair value, these two products carry different levels of riskiness and profitability which must be disclosed since they are suitable for investors who have very different utility functions.

The same argument applies to many non-equity return-target products featuring non-elementary payoff structures that translate into quite uneven density functions, as in the case of complex financial engineering (such as in the presence of one or more derivative components, also implicitly incorporated in the structure) but also in the case of unsophisticated products that present significant exposures to certain risk factors.

For these products the existence of a time horizon implicit in their financial engineering ensures the applicability of the contingent claim evaluation theory under the risk-neutral measure (see Section 2.2). By relying on the risk-neutral density at time T, this theory can be usefully exploited to improve the quality of the information conveyed to investors precisely in order to give a more comprehensive representation of the risk–return profile of the product.

This is partially achieved by enriching the informative set provided to investors with increased detail offered inside the financial investment table, as explained in the previous section. However, as long as the figures displayed in that table are only valid "on average", there still remains a huge information gap, which requires the development of further indicators that, starting from the risk-neutral

Figure 2.2 Densities of the same non-equity product obtained with two different models

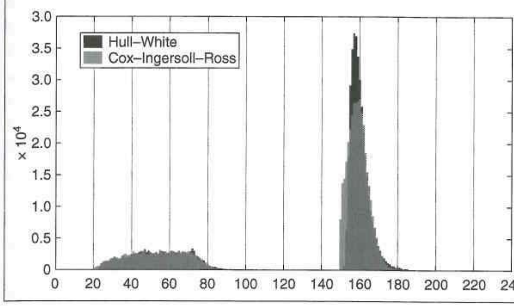

density, provide a clear and objective illustration of the levels of the possible performances and of their variability.

Ideally, the maximum transparency on these topics is offered precisely by the availability of the risk-neutral density of the final payoffs of the product. However, it is unlikely that retail investors have the statistical and financial background necessary to autonomously handle an entire probability distribution.

In addition, most of the time, the presence of the so-called model risk must also be considered, which implies that this distribution is not unique. In fact, for any given product, once the family of models that are suitable for dealing with the associated pricing problem have been identified, the different models of this family unfortunately lead to different risk-neutral densities.

Figure 2.2 shows the risk-neutral density of the same subordinated floating-coupon bond as obtained by using two different stochastic term structure models developed respectively by Hull and White (1990) and by Cox *et al* (1985).

Although both models belong to the class of unifactorial affine short-rate models, Figure 2.2 shows that the two densities do not coincide. In this specific case the reason is that, apart from the credit risk exposure, the cashflows of the bond are strongly dependent on the time evolution of the interest rates, which varies between the two models.

The above arguments indicate that the raw disclosure of the risk-neutral density is not a viable way to fill the above-mentioned information gap, since the richness and flexibility of information connected with such density would come at the cost of an increased complexity of comprehension for the average investor.

What is really needed is a method able to efficiently exploit the information embedded in the risk-neutral density and to smooth the differences between densities generated by different models.

The solution envisaged by the first pillar is the so-called table of probabilistic scenarios (hereafter also referred to as the probability table), which summarises the salient features of the variability characterising the potential performance of the investment by partitioning the density into a few mutually exclusive macro-events which are of concrete interest for the investor.

In general terms, this solution relies on the well-known principle of reduction in granularity, which turns out to be very helpful in diminishing the relevance of the model risk. Returning to the example of Figure 2.2, an elementary partition of the two densities displayed there in two complementary events gives very similar probabilities. For instance, by defining the two events with respect to the issue price of 100, ie, "the final value of the investment is lower than the issue price" and "the final value of the investment is higher than the issue price", the probabilities of these events are the same, namely 26.8% for the first event and 73.2% for the second one.

Clearly, by choosing a different partition, the probabilities of a given event under the two models could differ. And, in general terms, the more irregular the financial structure of the product the more likely it is that the same partition applied to densities obtained from different models could exhibit some differences. But, typically they will not be significant, especially when the interest is in capturing the key message conveyed by a specific partition.

From a technical point of view, there are infinitely many ways to aggregate the elementary events of a probability distribution in macro-events which are mutually exclusive; and, also, in principle, any individual might want to know in depth specific subsets of the risk-neutral density. However, the need for endowing all investors with the same core information and to preserve the comparability across products requires consideration of the same macro-events for all non-elementary return-target products.

In this context the most reasonable choice is to identify the macro-events according to criteria apt to immediately disclose the performance risk of a product, defined as its ability to create added value for the investor with respect to the initial outlay *per se* and also to the results achievable by taking an alternative investment decision.

In the case of non-elementary return-target products representing a direct financial investment, that is, an investment that does not intervene to substitute or modify a pre-existing product or liability held by the investor, the above alternative could be any non-equity product. But, in the absence of a specific pre-existing position, the choice of a particular product available on the market would be discretionary. Hence, it is necessary to look for an alternative investment which minimises any arbitrary assumption on investors' preferences and, at the same time, is able to represent in a clear, immediate and significant way how the specific risk factors and financial structure of the non-equity product of interest will affect the payoffs that can be obtained. From this perspective, the simplest choice is to compare the risk-neutral density of the product with the density associated with the investment of the same amount of money in a risk-free asset.

The latter is intended as an investment which, over the same time horizon of the product and given an initial outlay equal to the issue price of 100, pays a return equal to that accrued at the risk-free rates of the currency area where the product is sold. As seen in Section 2.2, in finance this process is also called risk-neutral numéraire and it is modelled through the stochastic process $\{B_t\}_{t\in[0,T]}$ which is governed by Equation 2.2, ie

$$B_T := \exp\left(\int_0^T r_s \, ds\right)$$

which reveals that the unique source of uncertainty for the risk-free asset are the movements in the interest rates curve.[6] Hence, the risk-neutral density of the final values of an initial investment of 100 in the risk-free asset reproduces exactly the impact of interest rate volatility on the returns of a financial investment, ensuring that the comparison with the non-equity product highlights the influence of the specific features that characterise such a product.

The partition technique used to make this comparison is the superimposition of the two densities. According to this technique, the risk-neutral density of the product is partitioned with respect to fixed thresholds which are exogenously identified depending on the point of zero return (ie, the final value of the investment is equal to 100) and on the risk-neutral density of the risk-free asset. This allows the information connected with the second moment of the product's probability distribution to be highlighted, and, hence, the implications of the volatility of its returns on the payoffs that investors can face at the end of the investment time horizon to be understood.

The full methodology adopted to determine the probability table is explained in Section 2.3.3. The final output is a table which displays the probabilities of four alternative macro-events and a synthetic indicator of the final value of the investment associated with each of them. For some products, depending on the specific shape of their risk-neutral density function, this information set is supplemented by further indicators in order to guarantee the complete investors' comprehension and a honest comparison with the risk-free asset.

Regardless of these technical aspects, it is clear that by properly exploiting the principle of reduction in granularity, the table of probabilistic scenarios attains the goal of reducing the amount of available data to be handled by the investor by preserving, at the same time, the core of the additional information contained in the risk-neutral probability density with respect to the fair value, so that investors achieve a much better awareness of the overall performance risk behind a non-elementary product.

In fact, although in general a higher fair value corresponds to a greater profitability of the investment, this indication is only true on average as stated above. On the other hand, the probability table has the advantage of coming from the same risk-neutral density used to determine the fair value and of preserving, at the same time, the fundamental information related to the specific shape of this distribution without being excessively exposed to the model risk.

In this way, fair values and time horizons being equal, it becomes possible to understand the effect that the investment risks have on investment performances, hence qualifying from time to time a safer product (similar in substance to the risk-free asset) or, conversely, a particularly daring product which, for example, combines a high probability of obtaining results more desirable than the risk-free asset with a non-negligible probability of receiving returns that are positive but not competitive or even negative returns.

It is worth pointing out that, for products representing a direct financial investment, the probabilistic scenarios do not allow for a

point-to-point comparison with the risk-free asset, but rather offer a comparison based on the analysis of the positioning of the product's density with respect to that of the risk-free alternative investment.

The main alternative methodology, consisting in carrying out the comparison trajectory by trajectory, would have the inconvenience of being quite difficult to disclose to retail investors, as it would produce relative performance indicators, which are unfamiliar to investors and require an excessive effort to be interpreted by them.

However, as illustrated in Section 2.3.4, the trajectory-by-trajectory comparison turns out to be very useful when the product is intended to replace a pre-existing investment (as in the case of exchange public offerings or when an investor is willing to replace a specific non-equity product they are holding), is a structured financial liability, or intervenes to modify the features of an existing liability (like a mortgage).

Both the superimposition and the trajectory-by-trajectory techniques require a subdivision of the density in a limited number of events; these classes of events are obviously coarser grained, and do not necessarily capture finer nuances related to different modelling hypotheses. Nevertheless, as noted above, this reduction of granularity is not a problem in itself, as practically any reasonable level of information within this framework can be captured, while reducing the level of complexity found in the comprehension of an abstract object such as a probability distribution.

In fact, it has to be considered that, from a practical point of view, even if, after the reduction in granularity, slight differences in the probability scenarios under different models specifications persist, usually they are not determinant variables in the investor's selection process. In other words, a significance threshold in the retail investor's perception seems to exist, and this phenomenon makes any differences of few percentage points substantially irrelevant.

A well-executed reduction in granularity eventually enables the sharing of useful information between issuers and investors, thus allowing efficient use of limited amounts of available data and helping to overcome the information asymmetry that characterises the markets of non-equity investment products.

Before presenting the details of the two methodologies, it is worth clarifying how the contents of the table of probabilistic scenarios should be interpreted. Numbers in this table support investors' understanding of the more-or-less balanced equilibrium between risks and opportunities of the product as they result from the information available at the time of issue.

This means that they come from a prospective quantitative analysis performed by using information available *ex ante*. The goal is to provide investors with a useful tool to assess and compare different investment proposals. But there is no claim about the fulfilment of those numbers at the expiry of the time horizon. Nor may it be otherwise, since (except in very unusual cases) each scenario has a certain probability of occurrence. This clarification also allows the exclusion of views that deem the representation of scenarios an additional source of liability for issuers. Although probabilistic scenarios are more detailed and, therefore, more meaningful than the mere indication of the fair value, it must always be remembered that both result from the same "raw information" given by the risk-neutral probability density of the final payoffs.