Risk Neutral Pricing

1 Introduction and History

A classical problem, coming up frequently in practical business, is the valuation of future cash flows which are somewhat risky. By the term "risky" we mean that the payment is not of a deterministic nature; rather there is some uncertainty on the amount of the future cash flows. Of course, in real life virtually everything happening in the future contains some elements of riskiness.

As a first example let us think of an investment project, say, a company plans to build a new factory. A classical way to proceed is to calculate a *net asset value*. One tries to estimate the future cash flows generated by the project in the subsequent periods. In the present example they will initially be negative; this initial investment should be compensated by positive cash flows in later periods. Having fixed these estimates of the future cash flows for all periods, one calculates a *net asset value* by discounting these cash flows to the present date. But, of course, there is uncertainty involved in the estimation of the future cash flows and people doing these calculations are, of course, aware of that. The usual way to compensate for this uncertainty is to apply an interest rate which is higher than the *riskless*[1] *rate of return* corresponding to the rate of return of government bonds.

The spread between the riskless rate of return and the interest rate used for discounting the future cash flows in the calculation of the net asset value can be quite substantial in order to compensate for the riskiness. Only if the net asset value, obtained by discounting with a rather high rate of return, remains positive, the management of the company will engage in the investment project.

Mathematically speaking, the above procedure may be described as follows: first one determines the *expected values* of the future cash flows and subsequently one discounts by using an elevated discount factor. However, there is no systematic way of mathematically approaching the question of how the degree of uncertainty in the determination of the expected values can be quantified, and in which way this should be taken into account to determine the spread between the interest rates.

We now turn to a different approach which interchanges the rôles of taking expectations and discounting in taking the riskness of the cash flows into account. This approach is used in modern mathematical finance, in particular

[1] In real life nothing is actually *riskless*: in practice, the *riskless rate of return* corresponds to government bonds (provided that the government is reliable).

in the Black Scholes formula. However, the idea goes back much further and the method was used by actuaries for centuries.

Think of a life insurance contract. To focus on the essential point we consider the simplest case: a one year death insurance. If the insured person dies within the subsequent year, the insured sum S, say $S = 1$ Euro, is paid out at the end of this year; if the insured person survives the year, nothing is paid and the contract ends at the end of the year.

To calculate the premium[2] for this contract, actuaries look up in their mortality tables[3] the probability that the insured person dies within one year. The traditional notation for this probability is q_x, where x denotes the age of the insured person.

To calculate the premium for such a one year death insurance contract, with S normalized to $S = 1$, actuaries apply the formula

$$P = \frac{1}{1+i}\, q_x. \tag{1}$$

The term q_x is just the expected value of the future cash flow and i denotes "the" interest rate: hence the premium P is the *discounted expected value* of the cash flow at the end of the year.

It is important to note that actuaries use a "conservative" value for the interest rate, e.g. $i = 3\%$. In practical terms this corresponds quite well to the "riskless rate of return". In any case, it is quite different, in practical as well as in theoretical terms, from the discount factors used to calculate the net asset value of a risky future cash flow according to the method encountered above.

But, after all, the premium of our death insurance contract also corresponds to the present value of an uncertain future cash flow! How do actuaries account for the risk involved in this cash flow, if not via an appropriate choice of the interest rate?

The answer is simple when looking at formula (1) above: apart from the interest rate i the probability q_x of dying within the next year also enters the calculation of P. The art of the actuarial profession is to choose the "good" value for q_x. Typically actuaries know very well the actual mortality probabilities in their portfolio of contracts, which often consists of several hundred thousand contracts; in other words, they have a very good understanding what the "true value" of q_x is. However, they do not apply this

[2]We do not consider costs, taxes etc, which are eventually added to this premium; we only consider the "net premium".

[3]A mortality table (horrible word!) is nothing but a list of probabilities q_x, where x runs through the relevant ages, say $x = 18, ..., 110$. The first mortality table was constructed by Edmond Halley in 1693.

"true value" in their premium calculations: in the above example (1) they would apply a value for q_x which is substantially higher than the "true" value of q_x. Actuaries speak about mortality tables of first and second kind.

Mortality tables of second kind reflect the "true probabilities". They are only used for the internal analysis of the profitability of the insurance company. On the other hand, in the daily life of actuaries only the mortality tables of first kind come into play which display properly "modified" probabilities q_x. They are not only used for the calculation of premia, but also for all quantities of relevance involved by an insurance policy, such as surrender values, reserves etc. This constitutes a big strength of the actuarial technique: actuaries always remain in a perfectly coherent logic when doing all these calculations. This logic is that of a fair game or, mathematically speaking, of a martingale. Indeed, *if the q_x would correctly model the mortality* of the insured person and *if i were the interest rate the insurance company could precisely achieve* when investing the premia, then the premium calculation (1) would make the insurance contract a fair game.

It is important to note that this argument pertains only to a kind of virtual world, as it is precisely the task of the actuaries to choose the mortalities q_x in a prudent way, which do *not* coincide with the "true" probabilities. In the case of insurance contracts where the insurance company has to pay in the case of death, actuaries choose the probabilities q_x higher than the "true ones". This happens in the simple example considered above. On the other hand, if the insurance company has to pay when the insured person is still alive, e.g. in the case of a pension, actuaries use probabilities q_x which are lower than the "true ones", in order to be on the safe side.

We took some time and space to elaborate on these actuarial techniques as we believe that this is helpful to understand more clearly the essence of the option pricing approach of Black, Scholes, and Merton. Their wellknown model for the risky stock S and the riskfree bond are

$$\begin{aligned} dS_t &= S_t \mu dt + S_t \sigma dW_t, \\ dB_t &= B_t r dt. \end{aligned} \tag{2}$$

The task is to value a (European) derivative on the stock S at expiration time T, e.g. $C_T = (S_T - K)_+$. As presumably everybody who read until this point knows very well **eqf04/006: The Black Scholes formula**, the solution proposed by Black, Scholes, and Merton is

$$C_0 = e^{-rT} \mathbb{E}_Q[C_T]. \tag{3}$$

The above formula is a perfect analogue to the premium of a death insurance contract (1) above. The first term, taking care of the discounting, uses the "conservative" choice of a riskless interest rate r. The second term gives the expected value of the future cash flow, taken under the *risk neutral* probability measure Q. This probability measure Q is chosen in such a way that the dynamics (2) of the stock under Q become

$$dS_t = S_t r dt + S_t \sigma dW_t. \tag{4}$$

The point is that the drift term $S_t r dt$ of S under Q is in line with the growth rate of the riskfree bond

$$dB_t = B_t r dt. \tag{5}$$

The interpretation of (4) is that *if the market were correctly modelled by the probability* Q, then the market were *risk neutral.* The mathematical formulation is that $(e^{-rt}S_t)_{0\leq t\leq T}$, i.e. the stock price process discounted by the riskfree interest rate r, is a martingale under Q.

Similarly as in the actuarial context above, the mathematical model of a financial market under the risk neutral measure Q pertains to a virtual world, not to the real world. In reality, i.e. under $\mathbb{P}$, we would typically have $\mu > r$. Fixing this case, Girsanov's formula (see **eqf02/014: Equivalence of probability measures, eqf02/022: Stochastic Exponentials**) tells us precisely that the probability measure Q represents a "prudent choice of probability". It gives less weight than the original measure $\mathbb{P}$ to events which are favorable for the buyer of a stock, i.e. when S_T is large. On the other hand, Q gives more weight than $\mathbb{P}$ to unfavorable events, i.e. when S_T is small. This can be seen from Girsanov's formula

$$\tfrac{dQ}{d\mathbb{P}} = \exp\Big[-\tfrac{\mu-r}{\sigma}W_T - \tfrac{(\mu-r)^2}{2\sigma^2}T\Big]$$

and the dynamics of the stock price process S under $\mathbb{P}$ resulting from (2)

$$S_T = S_0 \exp\Big[\sigma W_T + \big(\mu - \tfrac{\sigma^2}{2}\big)T\Big].$$

Fixing a random element $\omega \in \Omega$, the Radon–Nikodym derivative $\frac{dQ}{d\mathbb{P}}(\omega)$ is small iff $W_T(\omega)$ is large, and the latter is large iff $S_T(\omega)$ is large.

In many applications it is not even necessary to consider the original "true" probability measure $\mathbb{P}$. There are hundreds of papers containing the sentence: "we work under the risk neutral measure Q". This is parallel to the situation of an actuary in her daily work: she does not bother about the "true" mortality probabilities, but only about the probabilities listed in the mortality table of first kind.

The history of the valuation formula (3) goes, in fact, back much further than Black, Scholes, and Merton. Already in 1900 L. Bachelier applied this formula in his thesis [1] in order to price options. It seems worthwhile to have a closer look. Bachelier did not use a discount factor, such as e^{-rT}, in (3). The reason is that in 1900 prices underlying the option were denoted in forward prices at the Paris stock exchange (called "true prices" by Bachelier who also carefully adjusted for coupon payments; see [10] for details). When passing to forward prices the discount factor disappears as is well known. In modern terminology this fact boils down to "Black's formula".

As regards the second term in (3), Bachelier started from the very beginning with a martingale model, namely (scaled) Brownian motion (see, e.g. [10])

$$S_t = S_0 + \sigma W_t, \qquad 0 \leq t \leq T.$$

In other words, he also "works under the risk neutral probability".

In fact, in the first pages of his thesis Bachelier does speak about *two kinds of probability.* We quote from [1]:
(i) The probability which might be called "mathematical", which can be determined a priori and which is studied in games of chance.
(ii) The probability dependent on future events and, consequently impossible to predict in a mathematical manner.
This last is the probability that the speculator tries to predict.

Admitting a large portion of good will and hindsight knowledge one might interpret (i) as something like the risk neutral probability Q, while (ii) plays the role of $\mathbb{P}$.

2 Risk neutral pricing for general models

In the Black Scholes model (2) there is only one *risk neutral* measure Q under which the discounted stock price process becomes a martingale.[4]

This feature characterizes *complete financial markets* **eqf04/006: Complete markets**. In this case, we not only obtain from (3) a price C_0 for the derivative security C_T, but we get much more: in addition, the derivative can be perfectly replicated by starting at time $t = 0$ with the initial investment given by (3) and subsequent dynamical trading in the underlying stock S **eqf04/005: Replication**. This is the essence of the approach of Black,

[4] To be precise: this result only holds true if for the underlying filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \leq t \leq T}, \mathbb{P})$ we have $\mathcal{F} = \mathcal{F}_T$ and the filtration $(\mathcal{F}_t)_{0 \leq t \leq T}$ is generated by $(S_t)_{0 \leq t \leq T}$.

Scholes, and Merton; it has no parallel in the classical actuarial approach or in the work of L. Bachelier.

What happens in incomplete financial markets, i.e. when there are more than one risk neutral measure Q? It has been shown by M. Harrison and S.R. Pliska [6] that formula (3) yields *precisely* all the *consistent pricing rules* for derivatives on S, when Q runs through the set of risk neutral measures equivalent to $\mathbb{P}$. We denote the latter set by $\mathcal{M}^e(S)$. The term *consistent* means that there should be no arbitrage possibilities when all possible derivatives on S are traded at the price given by (3).

But what is the good choice of $Q \in \mathcal{M}^e(S)$? In general this question is as meaningless as the question: what is the good choice of an element in some convex subset of a vector space? In order to allow for a more intelligent version of this question one needs additional information. It is here that the original probability measure $\mathbb{P}$ comes into play again: a popular approach is to choose the element $Q \in \mathcal{M}^e(S)$ which is "closest" to $\mathbb{P}$.

In order to make this idea precise, fix a strictly convex function $V(y)$, e.g.

$$\begin{aligned} & V(y) = y\big(\ln(y) - 1\big), \quad && y > 0, \\ \text{or} \qquad & V(y) = \tfrac{y^2}{2}, && y \in \mathbb{R}. \end{aligned}$$

Determine $\hat{Q} \in \mathcal{M}^e(S)$ as the optimizer of the optimization problem

$$\mathbb{E}\big[V(\tfrac{dQ}{d\mathbb{P}})\big] \to \min! \qquad Q \in \mathcal{M}^e(S). \tag{6}$$

To illustrate things at the hand of the above examples: For $V(y) = y\big(\ln(y) - 1\big)$, this corresponds to choosing the element $\hat{Q} \in \mathcal{M}^e(S)$ minimizing the relative entropy $H(Q|\mathbb{P}) = \mathbb{E}_Q\big[\ln(\tfrac{dQ}{d\mathbb{P}})\big]$; for $V(y) = \frac{y^2}{2}$, this corresponds to choosing $Q \in \mathcal{M}^e(S)$ minimizing the L^2-norm $\|\frac{dQ}{d\mathbb{P}}\|_{L^2(\mathbb{P})} = \mathbb{E}_{\mathbb{P}}\big[(\frac{dQ}{d\mathbb{P}})^2\big]^{\frac{1}{2}}$.

Under appropriate conditions the minimization problem (6) has a solution, which then is unique by the strict convexity assumption.

There is an interesting connection with the issue of **eqf04/011: Indifference pricing**. Let $U(x)$ be the (negative) Legendre–Fenchel transform of V, i.e.

$$U(x) = \inf_y \{-xy + V(y)\}.$$

For the two examples above we obtain

$$\begin{aligned} & U(x) = -e^{-x} \\ \text{or} \qquad & U(x) = -\tfrac{x^2}{2}, \end{aligned}$$

which may be interpreted as utility functions. It turns out that – under appropriate assumptions – the optimizer $\hat{Q}$ in (6) yields precisely the marginal utility indifference pricing rule when plugged into (3) **eqf04/011: Indifference pricing**.

In particular we may conclude that pricing by marginal utility (see e.g., [8, 3, 5]) is a consistent pricing rule in the sense of Harrison and Kreps.

Walter Schachermayer

References

[1] Bachelier L. (1964) Théorie de la Spéculation. *Ann. Sci. Ecole Norm. Sup.* **17**, 21–86 (1900) English translation in: *The Random Character of stock market prices* (P. Cootner, editor), MIT Press.

[2] Black, F, Scholes, M. (1973) The pricing of options and corporate liabilities. *Journal of Political Economy* **81**, 637–659.

[3] Davis, M. (1997) Option pricing in incomplete markets. *Mathematics of Derivative Securities (M.A.H. Dempster, S.R. Pliska, editors), Cambridge University Press*, 216–226.

[4] Delbaen, F, Schachermayer, W. (2006) The Mathematics of Arbitrage. *Springer Finance*, 371.

[5] Foldes, D. (2000) Valuation and martingale properties of shadow prices: an exposition. *J. Econ. Dyn. Control* **24**, 1641–1701.

[6] Harrison, J.M, Pliska, S.R. (1981) Martingales and Stochastic Integrals in the Theory of Continuous Trading. *Stochastic Processes and their Applications* **11**, 215–260.

[7] Merton, R.C. (1973) The theory of rational option pricing. *Bell J. Econ. Manag. Sci.* **4**, 141–183.

[8] Rubinstein, M. (1976) The valuation of uncertain income streams and the pricing of options. *Bell J. Econ.* **7**, 407–426.

[9] Samuelson, P.A. (1965) Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review* **6**, 41–50.

[10] Schachermayer, W. (2003) Introduction to the Mathematics of Financial Markets. *S.* Albeverio, W. Schachermayer, M. Talagrand: Lecture Notes in Mathematics 1816 - Lectures on Probability Theory and Statistics, Saint-Flour summer school 2000 (Pierre Bernard, editor), Springer Verlag, Heidelberg , 111–177.