Probabilistic Aspects of Financial Risk

Hans Föllmer

Abstract. Problems arising in Finance have become a significant source of new devopments in Stochastic Analysis. We discuss some recent case studies, in particular some decomposition and representation theorems which are motivated by problems of hedging derivatives and of intertemporal consumption choice.

1. Introduction

In recent years, Mathematics has become highly visible as a key technology in the area of Finance. This is reflected in new journals, new curricula, and in new job profiles. Increasingly, advanced methods from Probability and Statistics are being applied to the analysis of financial risk in its various forms. Their impact is not only felt on a computational level. To a surprising extent, concepts of Stochastic Analysis are shaping the discourse of the field, both on the academic level and in the financial industry. Conversely, Finance has become a significant source of research problems which are of intrinsic mathematical interest, in particular in the area of stochastic analysis. The purpose of this lecture is to discuss some case studies which illustrate this converse impact.

We begin by explaining the key paradigm of a *perfect hedge*. It is concerned with financial *derivatives*, or *contingent claims*, which are defined in terms of some underlying financial asset. The mathematical analysis of such derivatives was started by Bachelier [1] who introduced Brownian motion as a model for the underlying price dynamics and computed prices for certain derivatives as expected values in this model. A nice diffusion model such as (geometric) Brownian is *complete* in the sense that any derivative can be replicated by a trading strategy in the underlying asset. From a mathematical point of view, this amounts to a version of Itô's representation theorem which represents functionals of Brownian motion as Itô integrals. From a financial point of view, the existence of a replicating strategy was the crucial insight in the work of Black and Scholes [6] and Merton [29] on arbitrage-free option pricing. Even though its standard version involves "high tech" methods from probability theory, the argument behind the pricing formula of Black and Scholes is essentially probability-free on a conceptual level; this is explained in section 1.

However, the idea of a perfect hedge is limited in scope. Realistic models will be *incomplete*, i.e., derivatives typically will carry an intrinsic risk which can not be

hedged away completely. At this stage, Probability comes in on a more fundamental level, and the hedging problem has become a source of new decomposition theorems in the theory of semimartingales. In this setting, we are going to focus on controlling the *shortfall*, i.e., that part of the contingent claim which will not be covered by the result of our hedging strategy. In section 2.1 we insist on staying on the safe side by constructing a *superhedging* strategy which keeps the shortfall down to 0. In the general case, the construction of such a superhedging strategy involves a decomposition theorem of Kramkov [24] which may be viewed as a new version of the fundamental Doob-Meyer decomposition for supermartingales. In many situations the cost of superhedging a given claim will be too high from a practical point of view. Thus, we are led to the problem of constructing strategies which are *efficient* with respect the cost and a suitably defined *shortfall risk*. In section 2.2 we discuss the solution of this problem in Föllmer and Leukert [18]; in particular we will see how it is reduced to the problem of superhedging a suitably modified claim.

As a further case study, we consider in section 3 a new variant of the problem of intertemporal consumption choice. Through the analysis of Bank and Riedel [4], it leads to a new stochastic representation problem which is closely related to the continuous-time theory of the Gittins index in El Karoui and Karatzas [10]. Its solution in Bank [2] involves a new stochastic analogue to the concept of a convex hull and is of interest from a purely mathematical point of view, quite independent of the initial motivation from Finance.

2. The idea of a perfect hedge

In this section we describe the idea of a *perfect hedge*, one of the key insights in mathematical finance. It is concerned with financial derivatives defined in terms of some underlying financial asset. A perfect hedge consists in replicating the derivative by a dynamical portfolio strategy which trades in the underlying asset. In a nice diffusion model for the stochastic fluctuation of asset prices such as geometric Brownian motion, this can indeed be done. In the usual presentation of this result a paradox arises: The argument uses "high tech" probabilistic methods, but the phenomenon itself does not really involve Probability on a conceptual level.

In the standard formalism of mathematical finance, the evolution of the (properly discounted) price of some underlying asset over a given time interval $[0,T]$ is described as a semimartingale $X = (X_t)_{t\in[0,T]}$ on a probability space $(\Omega, \mathcal{F}, P)$ with filtration $(\mathcal{F}_t)_{t\in[0,T]}$. At this point, a reader not familiar with the notions of stochastic analysis may feel somewhat discouraged. But in this section we will drastically reduce the conceptual and technical difficulties by eliminating the probabilistic ingredients altogether. The basic uncertainty about the behaviour of asset prices will be only reflected in the fact that we specify a large set Ω of possible scenarios ω, and that we do not know in advance which scenario will actually take place.

A *derivative*, or a *contingent claim*, specifies a payoff $H(\omega)$ contingent on the scenario $\omega \in \Omega$ which will be realized. For example, a European call option with strike price c and exercise time T has payoff $H(\omega) = (X_T(\omega) - c)^+$. What is the *fair price* which should be payed by the buyer of such an option? In other words, what is the fair deterministic equivalent to the random variable H? This is a classical question, and the standard answer goes back to the founding fathers of probability theory, in particular to Jacob Bernoulli. It says that you should assign probabilities to the different scenarios ω and compute the expexted value $E[H] = \int H dP$ of the random variable H with respect to the resulting probability measure P. Following Daniel Bernoulli, one might also want to add a risk premium in order to take account of risk aversion. But in our present financial context, the basic insight of Black and Scholes [6] and Merton [29] leads to a quite different result. In particular there will be no reason to add a risk premium, because the argument will show that there is no intrinsic risk. As we are ging to see, the answer does not really involve the probabilities of the different scenarios; it is the structure of the space of scenarios which matters.

As *scenarios* we take the possible trajectories of the asset price starting from a fixed initial value x_0. Such a trajectory will be described as a positive continuous function $\omega : [0, T] \to R^+$. We use the notation $X_t(\omega) = \omega(t)$ for the corresponding price at time t; in particular we have $X_0(\omega) = x_0$. Let us introduce the space Ω of all such functions which admit a continuous quadratic variation

$$[X]_t(\omega) = \lim_n \sum_{t_i \leq t, t_i \in D_n} (X_{t_{i+1}}(\omega) - X_{t_i}(\omega))^2 \tag{1}$$

along the sequence of dyadic partitions D_n. For such functions we can use a strictly pathwise version of Itô's calculus; see Föllmer [12]. In fact, for any function $F(x, t) \in C^2$ and for any scenario $\omega \in \Omega$, Itô's formula holds in the form

$$\begin{aligned} F(X_t(\omega), t) &= F(x_0, 0) + \int_0^t F_x(X_s(\omega), s) dX_s(\omega) \\ &\quad + \int_0^t \frac{1}{2} F_{xx}(X_s(\omega), s) d[X]_s + \int_0^t F_t(X_s(\omega), s) ds \end{aligned} \tag{2}$$

The second integral is well defined as a Lebesge-Stieltjes integral, while the first exists as an Itô integral, i.e., as limit of the non-anticipating Riemann sums

$$\sum_{t_i \leq t, t \in D_n} \xi_{t_i}(\omega)(X_{t_{i+1}}(\omega) - X_{t_i}(\omega)), \tag{3}$$

along the dyadic partitions D_n, where we put $\xi_t = F_x(X_t, t)$.

From a financial point of view where ω is seen as a possible scenario for the evolution of an asset price, the Itô integral can be interpreted as the *cumulative gain* from a self-financing trading strategy defined by ξ. In fact, if $\xi_t(\omega)$ is the number of shares held at time t then the terms of the non-anticipating Riemann sums in (3) are the net gains generated by the price changes in successive periods. The

non-anticipative construction of the Itô integral correponds exactly to the natural financial requirement that the investment decision at a time does not anticipate the upcoming price change during the next time interval. With this interpretation in mind, it is quite crucial to insist on trajectories with a non-vanishing quadratic variation. Otherwise there would be arbitrage opportunities, i.e., trading strategies which would offer the chance of a strictly positive gain without ony downside risk. In fact, suppose that we would only admit scenarios ω with quadratic variation $[X](\omega) = 0$. By Itô's formula, the strategy $\xi_t(\omega) = 2(X_t(\omega) - x_0)$ would then generate the gain $(X_T(\omega) - x_0)^2$, and this is clearly an arbitrage opportunity.

From a probabilistic point of view, the set Ω provides a natural framework because the law of a typical diffusion process will be carried by this set of paths. For a specific model like geometric Brownian motion we can be more precise and describe the typical trajectories in terms of a fixed volatility, given by a strictly positive (piecewise) continuous function σ on $[0,T]$. Without introducing any probabilities at this point, we will simply restrict the discussion to the set Ω_σ of all scenarios $\omega \in \Omega$ such that the quadratic variation is absolutely continuous with $d[X]_t = \sigma^2(t)X_t^2 dt$. In this context we can now explain the idea of a perfect hedge for a derivative of the form $H(\omega) = h(X_T(\omega))$ for some continuous function $h \geq 0$; a similar discussion can be found in Bick and Willinger [5]. Let $F(x,t)$ denote the solution of the parabolic equation

$$\mathcal{L}_F^\sigma(x,t) := (\frac{1}{2}\sigma(t)^2 x^2 \frac{\partial^2}{\partial x^2} + \frac{\partial}{\partial t})F(x,t) = 0 \tag{4}$$

on the strip $R^+ \times [0,T]$ with boundary condition $F(x,T) = h(x)$. Consider the trading strategy $\xi_t = F_x(X_t, t)$. By Itô's formula, we obtain

$$H(\omega) = F(X_T(\omega), T) = F(x_0, 0) + \int_0^T \xi_s(\omega) dX_s(\omega) \tag{5}$$

since the remaining part in Itô's formula vanishes due to (4). As we have seen, the Itô integral can be viewed as the net gain resulting from the self-financing trading strategy ξ. Thus, the contingent claim $H(\omega)$ is replicated as the final value of a portfolio generated by the initial amount $F(x_0, 0)$ and the strategy ξ. This replication of the derivative works simultaneously for all scenarios in our class Ω_σ. Thus, it does not involve any risk as long as we stick to the assumption that only scenarios with volatility profile σ will come up. Under this assumption, the initial cost $F(x_0, 0)$ of such a perfect replication is clearly the right price for the derivative. Otherwise there would be an arbitrage opportunity. If, for example, the price would be higher then one could sell the derivative at that price and at the same time replicate the option at a lower cost. For any scenario one would then generate the required amount $H(\omega)$ at the final time T, and the difference between the price and the cost of replication would remain as a risk free gain.

So far, our discussion did not involve any probabilities on the space of possible scenarios. But we can now introduce probabilities as a "high tech" device which

allows us to compute the correct price as an expected value rather than by solving a partial differential equation. By a theorem of Paul Lévy, there is exactly one probability measure P^* on the space Ω_σ such that the coordinate process $(X_t)_{t\in[0,T]}$ becomes a *martingale* under P^*, i.e.,

$$E^*[X_t|\mathcal{F}_s] = X_s \tag{6}$$

for times $s < t$. Under some growth condition on the function h, this will imply that the Itô integral in (6) has zero expectation under P^*. Thus, the price of the derivative can be computed as an expected value

$$F(X_0, 0) = E^*[H] \tag{7}$$

with respect to the probability measure P^*. In this sense, probabilities do arise. But they do not appear as objective probabilities which govern the realization of the different scenarios. Instead, the probability measure P^* plays the role of a consistent pricing rule for derivatives of the underlying asset. Consistency means, in particular, that the price $X_s(\omega)$ of the asset at time s should be equal to the conditional expectation of any future price X_t under the measure P^* given the information up to time s, and this is exactly the martingale property required for P^*.

Let us emphasize that the preceding probabilistic argument is optional, not really essential. In principle, the argument of pricing a derivative by means of a perfect replication is probability-free. But note that the volatility profile σ has to be known in advance; this is needed in order to solve in advance the parabolic differential equation which determines both the initial cost and the self-financing strategy. If this is not the case, then we have an additional source of uncertainty. Does probability come in at this second stage? Not necessarily. Of course we could model volatility as a stochastic process, and this is an important chapter in Mathematical Finance. But we could also continue to insist on a probability free approach using the following argument by T. Lyons [28]. Suppose that we are ready to assume that the unknown volatility $\sigma(t)$ will in any case stay inside some moving interval $\Sigma(t)$. Under some regularity conditions, there is a solution $F \in C^2$ of the non-linear Pucci equation

$$\mathcal{L}F(x,t) := (\sup_{\sigma(t)\in\Sigma(t)} \frac{1}{2}\sigma^2(t)x^2\frac{\partial^2}{\partial x^2} + \frac{\partial}{\partial t})F(x,t) = 0 \tag{8}$$

with terminal value $F(x,T) = h(x)$. For any scenario ω with a volatility profile σ staying inside the given moving boundaries, our pathwise Itô formula together with equation (8) implies

$$F(X_t(\omega), t) = F(x_0, 0) + \int_0^t F_x(X_s(\omega), s)dX_s(\omega) - B_t(\omega) \tag{9}$$

where

$$B_t(\omega) = \int_0^t (\mathcal{L} - \mathcal{L}^\sigma)F(X_s(\omega), s)ds \tag{10}$$

is increasing in t. From a financial point of view, this means that we can *superhedge* the contingent claim $H(\omega) = F(X_T(\omega), T)$ as follows. We start with the initial

amount $F(x_0, 0)$ and use the trading strategy $\xi_t = F_x(X_t(\omega), t)$. At the same time we cumulatively withdraw the capital $B_t(\omega)$. The resulting portfolio value at time t is given by $F(X_t(\omega), t)$, and at the final time T we have generated exactly the required payoff $H(\omega)$ of our derivative. Allocating the initial capital $F(X_0, 0)$ to the purpose of hedging the derivative, we are thus staying on the safe side. As time goes on and uncertainty is revealed, it will typically become clear at a given time that the capitel generated so far by the initial amount and by the stragegy ξ is higher than what is actually needed to stay on the safe side from that time on, and this explains the refunding scheme described by the increasing process B. As we will see in the next section, the probabilistic analogue of this approach is provided in great generality by Kramkov's optional decomposition theorem.

3. Controlling the short fall

Now we return to the standard probabilistic setting where the (properly discounted) price process is given as an adapted stochastic process $X = (X_t)_{t\in[0,T]}$ on a probability space $(\Omega, \mathcal{F}, P)$ with filtration $(\mathcal{F}_t)_{t\in[0,T]}$. From an economic point of view, the model should be such that it excludes arbitrage opportunities. One of the fundamental theorems in Mathematical Finance states that this *absence of arbitrage* is essentially equivalent to the existence of an *equivalent martingale measure*, i.e., a measure $P^* \approx P$ such that X is a martingale under P^* in the sense of (6); for a precise formulation see Delbaen and Schachermayer [9] and the references given there. Let us therefore assume

$$\mathcal{P} \neq \emptyset \tag{11}$$

where $\mathcal{P}$ denotes the class of all martingale measures $P^* \approx P$.

This condition may be viewed as a minimal version of the "Efficient Market Hypothesis". In its strong form, the efficient market hypothesis would claim that (properly discounted) prices should be martingales under the objective measure P. This is based on a rough equilibrium argument: If the present price is below the market's average expectation then the price should move immediately up, and vice versa. In fact, such an equilibrium argument was the starting point in the thesis of Bachelier [1] where Brownian motion was introduced as a model for price fluctuation on a speculative market. Condition (11) is of course much weaker since the martingale property is required only up to an equivalent change of measure. By Girsanov's theorem, it implies that X is a semimartingale under P, and so we are free to use the general tools of stochastic analysis.

Uniqueness of the equivalent martingale measure P^* is equivalent to *completeness*, i.e., any derivative H of the underlying price process admits (under some integrability conditions) a perfect replication in terms of a suitable trading strategy. As in the previous section, it follows that the price must be equal to the cost of

replication, and that it can be computed as the expectation $E^*[H]$ under the unique equivalent martingale measure P^*.

There are, however, many reasons why the reality of financial markets should be seen as being *incomplete*: Not any contingent claim admits a perfect hedge. In other words, there are intrinsic risks which cannot be hedged away completely by a strategy which uses the available assets. From a mathematical point of view, this means that there is more than one equivalent martingale measure, and since $\mathcal{P}$ is convex we have

$$|\mathcal{P}| = \infty. \tag{12}$$

In this general context let us now look at the problem of hedging a derivative described by a $\mathcal{F}_T$-measurable random variable $H \geq 0$. An *admissible strategy* is determined by some initial amount V_0 and by a predictable integrand ξ for X such that the resulting portfolio pocess

$$V_t = V_0 + \int_0^t \xi_s ds \tag{13}$$

satisfy $V_t \geq 0$ for all $t \in [0, T]$. Whichever strategy we use, we will end up at time T with a *decomposition*

$$H = V_0 + \int_0^T \xi_s ds + C_T \tag{14}$$

of the contingent claim into a part which is perfectly hedged and hence priced by arbitrage, and into the remaining *hedging error* C_T. There are different approaches to the problem of choosing a hedging strategy, and hence a specific decomposition of the contingent claim. From an economic point of view, they involve different preferences towards risk. From a mathematical point of view, they have been a significant source of new decomposition theorems in the general theory of stochastic processes.

One approach consists in choosing a strategy which minimizes the hedging error C_T in $\mathcal{L}^2(P)$. This involves a projection of H on a space of stochastic integrals. Alternatively, we could insist on a decomposition where the process C is mean-self financing in the sense that it is a martingale under the objective measure P. In the special case where the efficient market hypothesis holds in its strong form $P \in \mathcal{P}$, both formulations of the problem are equivalent, and it is solved by the classical Kunita-Watanabe decomposition for square integrable martingales; see Föllmer and Sondermann [13]. In the case $P \notin \mathcal{P}$, the two versions of the problem are different but intimately related. Often, but not always, the problem can be reduced to the Kunita-Watanabe-decomposition with respect to a suitable "minimal" martingale measure; see Föllmer and Schweizer [14]. In the general case, and in particular through the work of Martin Schweizer, the projection problem has become a source of new versions of the Kunita-Watanabe decomposition and of new closure results for spaces of stochastic integrals with respect to a semimartingale; see, e.g., Rheinländer and Schweizer [31]. For a survey of recent developments we refer to Schweizer [32].

Note, however, that there is a basic asymmetry in the financial interpretation of the hedging error which is not captured by the projection approach. It is really the *shortfall*, i.e., the positive part $C_T^+ = (H - V_T)^+$ of the hedging error, which should be kept under control. Here we will describe two approaches which focus on the shortfall. The first approach insists on staying on the safe side and keeping the shortfall down to 0. This is the idea of superhedging which has already appeared in the previous section. In our general probabilistic setting , it has led to a remarkable new version of the Doob-Meyer decomposition. The second approach consists in constructing partial hedges which are efficient in terms of the cost and of a suitably defined shortfall risk, and here the technique of superhedging comes in as a crucial tool.

3.1. Superhedging

In incomplete models it is not possible to replicate any given contingent claim H. But we could insist on keeping the shortfall C_T^+ down to 0. Thus, we would require

$$P[C_T \leq 0] = P[V_T \geq H] = 1. \tag{15}$$

The program of constructing such a strategy with a minimal initial amount has been carried out on increasing levels of generality; see El Karoui and Quenez [11], Karatzas [21] and the references given there, Kramkov [24] and Föllmer and Kabanov [16]. More precisely, let us assume

$$U_0 := \sup_{P^* \in \mathcal{P}} E^*[H] < \infty \,, \tag{16}$$

and let us define (U_t) as a right-continuous version of the process defined by

$$U_t = \operatorname*{ess.sup}_{P^* \in \mathcal{P}} E^*[H \mid \mathcal{F}_t] \,. \tag{17}$$

The process (U_t) is a $\mathcal{P}$-*supermartingale*, i.e., a supermartingale simultaneously for any $P^* \in \mathcal{P}$. In fact it is the smallest non-negative $\mathcal{P}$-supermartingale with terminal value $\geq H$.

The Doob-Meyer decomposition shows that a non-negative supermartingale with respect to a fixed probability measure can be represented as the difference of a local martingale and a predictable increasing process. Is there a reasonable analogue for $\mathcal{P}$-*supermartingales* which is valid simultaneously for all $P^* \in \mathcal{P}$? The study of this problem was initiated by El Karoui and Quenez [11] in a special context. As shown in full generality in Kramkov [24] and in Föllmer and Kabanov [16], any non-negative $\mathcal{P}$-supermartingale admits an *optional decomposition* of the form

$$U_t = U_0 + \int_0^t \xi_s dX_s - C_t \tag{18}$$

where C is an increasing optional process and ξ is an admissible strategy. Note that the stochastic integral is a local martingale simultaneously for all $P^* \in \mathcal{P}$. On the other hand, the increasing process is only optional but not necessarily predictable.

For the specific $\mathcal{P}$-supermartingale defined in (17), the optional decomposition can be interpreted as the following *superhedging* procedure: Put up the initial capital U_0, then follow the trading strategy ξ and withdraw the cumulative amount of capital C_t from the resulting portfolio as additional information comes in. As a corollary of the optional decomposition, the value U_t can be characterized as the least amount of capital needed at time t to cover the claim H by following an admissible strategy ξ from time t up to the final time T, i.e.,

$$U_t = \text{ess.inf} V_t \tag{19}$$

where V_t runs through the class of $\mathcal{F}_t$-measurable random variables ≥ 0 such that

$$V_t + \int_t^T \xi_s dX_s \geq H \quad P - a.s. \tag{20}$$

for some admissible strategy ξ. In other words, U_t is an upper bound for any arbitrage-free price of the claim computed at time t. If additional constraints are imposed on the strategies ξ then the dual description (19) of the process defined by (17) has a corresponding analogue in terms of a suitable extension of the class $\mathcal{P}$; see Karatzas [21] and Föllmer and Kramkov [15].

The minimal initial amount which is needed for staying on the safe side is given by U_0. From a practical point of view, this is usually too much. In fact, the whole idea of superhedging is more extreme resp. conservative than in any actuarial approach. It is also extreme from a mathematical point of view: For a convex derivative, the superhedging strategy can often be identified as a perfect replication strategy with respect to an associated complete model which is no longer equivalent to the initial model but sits on some kind of Choquet boundary of the initial space of scenarios.

In any case, superhedging provides the lowest upper bound U_0 for any arbitrage-free price of the given derivative. Moreover, as we shall see next, the optional decomposition of a suitably modified claim can be a crucial ingredient in the construction of hedging strategies which may be more realistic from a practical point of view.

3.2. Efficient Hedging: Cost versus Shortfall Risk

Together with the risk of a loss, superhedging will also take away the opportunity of making a profit. In fact, the cost of superhedging a contingent claim will typically exceed any actual price which could be obtained for that claim. Now suppose that the investor is unwilling to put up the initial amount U_0 required by a superhedging strategy and is therefore ready to accept some risk. What is the optimal *partial hedge* which can be achieved with a given smaller amount of capital V_0? In order

to make this question precise we need a criterion expressing the investor's attitude towards the risk of a shortfall.

One possible criterion is the probability $P[C_T > 0]$ that some shortfall occurs, i.e., that the final amount V_T generated by the strategy falls short of the amount which is required to cover the contingent claim. We could fix an upper bound α for this shortfall probability and minimize the initial cost under this constraint. Alternatively, we could fix some initial cost $V_0 < U_0$ and look for a strategy which minimizes the probability of a short fall. This problem is solved in Föllmer and Leukert [17], translating and extending an idea of Kulldorf [26] to the present financial context. The resulting strategy of *quantile hedging* arises as the superhedging strategy for a modified claim $\tilde{H}$. Typically, the modified claim takes the form of a knock-out option HI_A, and the set A is constructed as the optimal test in a Neyman-Pearson problem where the alternative is given by the objective measure P, and the compound hypothesis is defined in terms of the contingent claim and the class $\mathcal{P}$ of equivalent martingale measures. Quantile hedging may be viewed as a a dynamic version of the Value at Risk approach. It invites the same critique since it only takes into account the probability that a shortfall occurs, but not the size of the shortfall if it does occur. This motivates a modified approach where the investor's attitude towards the shortfall is specified in terms of a loss function l.

Let us assume that l is an increasing function on R^+ with $l(0) = 0$ and $E[l(H)] < \infty$; convexity of l would correspond to risk aversion. We introduce the notion of *shortfall risk*, defined as the expectation

$$E[l((H - V_T)^+)]$$

of the shortfall $C_T^+ = (H - V_T)^+$ weighted by the loss function l under the objective measure P. If we would be ready to allocate the initial amount U_0 then we could reduce the shortfall risk to 0 by using the superhedging strategy. For an initial capital $V_0 < U_0$, our problem is to find an admissible strategy which minimizes the shortfall risk under this cost constraint; alternatively, we could fix a bound on the shortfall risk and minimize the cost. In other words, we are looking for hedges which are *efficient* with respect to cost and shortfall risk. These efficient hedges interpolate in a systematic way between the extremes of a superhedge (no risk, no chance of making a profit) and no hedge (full risk of shortfall, full chance of profit), depending on the accepted level of shortfall risk.

This problem of minimizing the shortfall risk under a cost constraint is solved in Föllmer and Leukert [18]. The solution proceeds in two steps. In a first step, we consider the statistical decision problem

$$\min_{\varphi \in \mathcal{R}_0} E[l((1 - \varphi)H)] \tag{21}$$

where $\mathcal{R}_0$ denotes the class of all $\mathcal{F}_T$-measurable random values φ with values in $[0,1]$ such that

$$\sup_{P^* \in \mathcal{P}} E^*[\varphi H] \leq V_0 \ . \tag{22}$$

There exists a solution $\widetilde{\varphi}$, and we may assume $\widetilde{\varphi} = 1$ on $\{H = 0\}$. As pointed out in Delbaen and Schachermayer [9], existence results of this type follow very easily from a (simplified) version of Komlós's theorem that any sequence of random variables with bounded $\mathcal{L}^1$-norms admits a subsequence which satisfies the strong law of large numbers. Here we even have random variables bounded by 1, and in this case the theorem was already known to Rényi; see [23]. In fact, if (φ_n) is a sequence of functionsin $\mathcal{R}_0$ whose shortfall risk approaches the infimum in (21), and if $\widetilde{\varphi}$ is the limit of an almost surely convergent sequence of averages along some subsequence, then it is easy to check that $\widetilde{\varphi}$ is indeed a solution of problem (21).

Let us introduce the modified claim $\tilde{H} = \widetilde{\varphi} H$, and let $\tilde{U}$ denote a right-continuous version of the process

$$\tilde{U}_t = \operatorname*{ess.sup}_{P^* \in \mathcal{P}} E^*[\tilde{H} \mid \mathcal{F}_t]. \tag{23}$$

Since $\tilde{U}$ is a $\mathcal{P}$-supermartingale, we can now apply the optional decomposition theorem. As shown in Föllmer and Leukert [18], this yields the solution of our optimization problem:

Theorem. *The strategy $\tilde{\xi}$ determined by the optional decomposition of the $\mathcal{P}$-supermartingale $\tilde{U}$ generated by the modified claim $\widetilde{H} = \widetilde{\varphi} H$ minimizes the shortfall risk under the constraint that the initial cost is bounded by V_0.*

In the *complete* case where the equivalent martingale measure P^* is uniquely determined, the structure of the optimal profile $\widetilde{H} = \widetilde{\varphi} H$ can be described explicitely, and the optimal strategy consists in replicating the contingent claim $\widetilde{\varphi} H$. For a smooth convex loss function let $I = (l')^{-1}$ the inverse of l', and denote by ρ^* the Radon-Nikodym derivative of P^* with respect to P. Then the solution $\widetilde{\varphi}$ of the optimization problem (21) is given by

$$\widetilde{\varphi} = 1 - \left(\frac{I(c\rho^*)}{H} \wedge 1 \right) \quad \text{on } \{\mathrm{H} > 0\}, \tag{24}$$

where the constant c is determined by the condition $E^*[\widetilde{\varphi} H] = V_0$.

It is instructive to introduce a whole scale for the attitude towards risk and to consider the special case $l(x) = \frac{x^p}{p}$ where the risk is defined in terms of an upper partial moment of the hedging error. In the risk-averse case $p > 1$, the optimal hedge consists in replicating the modified claim

$$\varphi_p H = H - c_p (\rho^*)^{\frac{1}{p-1}} \wedge H \tag{25}$$

where the constant c_p is determined such that $E^*[\varphi_p H] = \tilde{V}_0$. In the limit $p \to \infty$ of ever increasing risk aversion, the modified claim $\varphi_p H$ converges to $(H - c)^+$, where

c is the unique constant that satisfies $E^*[c \wedge H] = E^*[H] - \tilde{V}_0$. If H is a call at strike K, then the limit for $p \to \infty$ is again a call at the higher strike $\tilde{K} = K + c$ whose arbitrage free price is given by V_0. We may also consider the case $p < 1$ where risk-averse behavior is replaced by risk-seeking behavior. As appetite for risk increases and p decreases from 1 to 0, the corresponding efficient hedges converge to the knock-out option which appears in the case of quantile hedging.

In the *incomplete* case, the problem of computing the optimal profile is much more subtle. A specific case study, where the problem consists in constructing efficient hedges for a volatility jump, can be found in Föllmer and Leukert [18]. The general case involves methods of convex duality, and in particular a deep result of Kramkov and Schachermayer [25] which describes the duality between the class of claims which can be generated by admissible strategies and a suitable set of supermartingale densities. We refer to Leukert [27] for details in the present context. Closely related results and extensions appear in Pham [31], Mnif and Pham [30], Cvitanić and Karatzas [7], and in Cvitanić, Schachermayer and Wang [8].

4. Stochastic optimization under intertemporal preferences

So far, our discussion has involved preferences on the space of portfolio values V_T generated by a trading strategy at a fixed terminal time T. Initiated by von Neumann und Morgenstern, the structure of reasonable preferences on a space of random variables has been studied in depth and is now well understood; for a survey see Karni and Schmeidler [22]. Typically, they can be represented in terms of an expected utility functional. Our optimization problem above fits into this framework: For a convex loss function l, we have maximized the functional $E[u(V_T(\omega), \omega)]$ defined in terms of the increasing concave scenario-dependent utility function $u(x, \omega) = -l(H(\omega) - x)^+$.

In the microeconomic theory of intertemporal consumption choice, the natural commodity space is no longer a space of random variables but a space of positive *optional random measures* C defined on a given time interval $[0, T]$. Such a measure describes a consumption pattern contingent on the scenario and adapted to the incoming information. Identifying measures with their distribution functions, we take as our commodity space $\mathcal{C}$ the space of all optional increasing processes $C = (C_t)_{t\in[0,T]}$ on our probability space $(\Omega, \mathcal{F}, P)$ with filtration $(\mathcal{F}_t)_{t\in[0,T]}$.

As a corollary to the optional decomposition, the minimal cost of financing the consumption pattern C by a trading strategy on the underlying financial market can be identified as

$$\sup_{P^*\in\mathcal{P}} E^*[\int_0^T \exp(-\int_0^t r_s ds) dC_t]. \tag{26}$$

Here (r_t) is a predictable process describing the short term interest rate; in the previous sections it did not appear explicitely since all assets were already assumed to be discounted by a savings account growing at the short rate. For any preference functional U on the space of positive measures, we have a corresponding *optimization problem*

$$\max_{C \in \mathcal{C}_0} E[U(C(\omega))] \tag{27}$$

where $\mathcal{C}_0$ is the class of all optional measures $C \in \mathcal{C}$ such that the cost of financing C as defined in (26) is not greater than a given initial amount V_0.

Of course, the structure of the solution will depend both on the underlying price dynamics and on the choice of the preference functional U. A common choice consists in admitting only absolutely continuous measures $dC = c_t dt$ with a rate of consumption (c_t), and to use a functional of the form

$$U(C) = \int_0^T u(c_t, t)dt, \tag{28}$$

where u is some time-dependent utility function, typically of the form $u(x,t) = u(x)\exp(-\delta t)$ with some discounting factor δ and some classical utility function u. In this case, the corresponding optimization problem may be viewed as a space-time version of the optimization problem considered in section 2, and it can be solved in an analogous way. It turns out, however, that the choice (28) of a preference functional on absolutely continuous measures is much less canonical than the expected utility functional on a space of random variables, both from a mathematical and from a microeconomic point of view. In particular, there is no reasonable extension to the space of all measures, no reasonable transition between discrete and continuous time, and no continuity with respect to the natural weak topology on the space of measures. Hindy, Huang, Kreps [19] provide a thorough analysis why a functional of the form (28) is not satisfactory from a microeconomic point of view. Instead, they propose a utility functional of the form

$$U(C) = \int_0^T u(Y_t^C, t)dt \tag{29}$$

defined for all finite measures C on $[0,T]$, where Y_t^C is some index of past consumption up to time t, for example a weighted average

$$Y_t^C = \int_0^t \exp(-\beta(t-s))\, dC_s. \tag{30}$$

The process Y^C will be called the *level of satisfaction* generated by C.

For a utility functional U of the form (29), the structure of the optimization problem (27) becomes much richer than in the case (28). Typically, the optimal optional measure $\widetilde{C}$ is no longer absolutely continuous; it may involve consumption in gulps and at a singular rate. If the underlying price dynamics is Markovian then one can use a standard approach in terms of a Hamilton-Jacobi-Bellman equation;

see Hindy and Huang [20]. Starting with the deterministic case, Bank and Riedel [3,4] have initiated a quite different approach which is based on a new infinite-dimensional version of the Kuhn-Tucker theorem. In a general stochastic setting, and in particular in the thesis of Bank [2], this approach leads to some remarkable new problems, which are of independent interest from the point of view of general semimartingale theory.

Existence and uniqueness of a solution $\widetilde{C}$ can be shown in great generality. In the complete case, and for a functional of the form (30), Bank and Riedel [4] and Bank [2] succeed in clarifying the structure of the optimal consumption plan $\widetilde{C}$ in terms of a reference process $I = (I_t)$, called the *minimal level of satisfaction*. They show that $\widetilde{C}$ can be explicitely constructed by "tracking" I. More precisely, $\widetilde{C}$ is reconstructed from its level of satisfaction $\widetilde{Y}$, and $\widetilde{Y}$ is computed explicitely as the minimal process in the class of all feasible index processes $Y^C \geq I$ with $C \in \mathcal{C}_0$. The crucial reference process I can be characterized in terms of the solution of a remarkable new *stochastic representation problem* of the following general form. Given an optional process $Z = (Z_t)$ (in our case defined in terms of the state-price densities) and a function $f(l,t)$ decreasing in l (in our case defined in terms of u), we want to construct an optional process $L = (L_t)$ such that

$$Z_t = E[\int_t^T f(\sup_{t \leq v \leq s} L_v, s) ds | \mathcal{F}_t]. \tag{31}$$

In discussions at an Oberwolfach meeting in May 2000, N. El Karoui and P. Bank clarified the connection to the theory of the Gittins index in continuous time; see El Karoui and Karatzas [10]. Using this connection, the complete solution of the representation problem (31) is carried out in Bank [2]. In the simple special case $f(l,t) = -l$, the process L appears as the solution of the optimal stopping problem

$$L_t = \text{ess.inf} \frac{E[X_\tau - X_t | \mathcal{F}_t]}{E[\tau - t | \mathcal{F}_t]}, \tag{32}$$

where the essential infimum is taken over all stopping times τ with values in $(t,T]$. In the general case the construction of the solution is quite intricate; see Bank [2].

Thus, the initial problem of intertemporal consumption choice leads to new extensions of the theory of optimal stopping which are of intrinsic mathematical interest, quite independent of the financial motivation.

References

[1] L. Bachelier, *Théorie de la spéculation*, Ann. Sci. Ecole Norm. Sup. **17** (1900), 21-86.

[2] P. Bank, *Singular Control of optional random measures: Stochastic Optimization and Representation Problems arising in the Microeconomic Theory of Intertemporal Consumption Choice,* Dissertation, Humboldt-Universität zu Berlin (2000).

[3] P. Bank and F. Riedel, *Non-Time Additie Utility Optimization - the Case of Certainty,* J. Math. Econ., **33** (2000), 271–290.

[4] P. Bank and F. Riedel, *Optimal Consumption Choice with Intertemporal Substitution,* Preprint Humboldt-Universität zu Berlin (2000).

[5] A. Bick and W. Willinger, *Dynamic Spanning without Probabilities,* Stoch. Proc. Appl. **50** (1994), 349 - 374.

[6] F. Black and M. Scholes, *The Pricing of Options and Corporate Liabilities,* J. Political Econom. **72** (1973), 637-659.

[7] J. Cvitanić and I. Karatzas, *On dynamic measures of risk,* Finance and Stochastics **3** (1999), 451 - 482.

[8] J. Cvitanić, W. Schachermayer and H. Wang, *Utility Maximization in Incomplete Markets with Random Endowment,* Preprint (2000).

[9] F. Delbaen and W. Schachermayer, *A general version of the fundamental theorem of asset pricing,* Math. Annalen **300** (1994), 463 - 520.

[10] N. El Karoui and I. Karatzas, *Dynamic Allocation Problems in Continuous Time,* Ann. Appl. Probab. **4** (1994), 255-286.

[11] N. El Karoui and M.C. Quenez, *Dynamic programming and pricing of contingent claims in an incomplete market,* SIAM J. Control and Optimization **33 (1)** (1995), 29-66.

[12] H. Föllmer,*Calcul d'Itô sans Probabilités,* In: Sém. Probabilités Strasbourg XV, Lecture Notes in Mathematics 850, Springer (1981), 143-150.

[13] H. Föllmer and D. Sondermann, *Hedging of Non-redundant Contingent Claims,* In: Contributions to Mathematical Economics, In Honor of G. Debreu (Eds. W. Hildenbrand and A. Mas-Colell), Elsevier Science Publ., North-Holland (1986), 205-223.

[14] H. Föllmer and M. Schweizer, *Hedging of Contingent Claims under Incomplete Information,* In: Applied Stochastic Analysis, eds. M.H.A. Davis and R.J. Elliott, Gordon and Breach, London (1990), 389-414.

[15] H. Föllmer and D. Kramkov, *Optional decompositions under constraints,* Prob. Theory Relat. Fields, vol. 109, (1997), 1-25.

[16] H. Föllmer and Y. Kabanov, *Optional decomposition and Lagrange multipliers,* Finance and Stochastics, vol. 2, No.1(1998), 69-81.

[17] H. Föllmer and P. Leukert, *Quantile Hedging.* Finance and Stochastics, vol. 3, No. 3 (1999), 251 - 273.

[18] H. Föllmer and P. Leukert, *Efficient hedges: cost versus shortfall risk,* Finance and Stochastics **4** (2000), 117-146.

[19] A. Hindy, C.-F. Huang and D. Kreps, *On Intertemporal Preferences in Continuous Time - the Case of Certainty,* J. Math. Econ. **21** (1992), 401 - 440.

[20] A. Hindy and C.-F. Huang, *Intertemporal Preferences for Uncertain Consumption: a Continuous-Time Approach,* Econometrica **61** (1993), 85-121.

[21] I. Karatzas, *Lectures in Mathematical Finance,* Providence: American Mathematical Society (1997)

[22] E. Karni and D. Schmeidler, *Utility Theory with Uncertainty.* In: Handbook of Mathematical Economics, vol. IV, Ch. 33, eds. W. Hildenbrand and H. Sonnenschein. Elsevier Science Publ., North-Holland (1991),1763 - 1831.

[23] J. Komlos, *A generalizaton of a problem of Steinhaus*, Acta Math. Acad. Sci. Hung. 18, 217 - 229 (1967)

[24] D.O. Kramkov, *Optional decomposition of supermartingales and hedging contingent claims in incomplete security markets,* Probability Theory and Related Fields **105** (1996), 459-479.

[25] O. Kramkov and W. Schachermayer, *The Asymptotic Elasticity of Utility Functions and Optimal Investment in Incomplete Markets,* Ann. Appl. Probab. Vol. 9, No. 3 (1999), 904-950.

[26] M. Kulldorff, *Optimal control of favorable games with a time-limit,* SIAM J. Control and Optimization **31** (1993), 52-69.

[27] P. Leukert, *Absicherungsstrategien zur Minimierung des Verlustrisikos,* Dissertation, Humboldt-Universität Berlin (1999).

[28] T. Lyons, *Uncertain volatility and the risk free synthesis of derivatives,* Appl. Math. Finance **2** (1995), 117-133.

[29] R. Merton, *Theory of Rational Option Pricing,* Bell J. Econom. Managem. Sci. **4** (1973), 141-183.

[30] M. Mnif and H. Pham, *Stochastic Optimization under Constraints,* Prepublication n^o 599, Laboratoire de Probabilité, Université de Paris VI et VII (2000).

[31] H. Pham, *Minimizing Shortfall Risk and Applications to Finance and Insurance,* Peprint Université de Marne-la-Vallée (1999).

[32] T. Rheinländer and M. Schweizer, *On L^2-Projections on a Space of Semimartingales,* Ann. Probab. **25** (1997), 1810 - 1831.

[33] M. Schweizer, *A Guided Tour through Quadratic Hedging Approaches,* Preprint TU Berlin (1999).

Department of Mathematics
Humboldt University
Unter den Linden 6
D-10099 Berlin
E-mail address: foellmer@mathematik.hu-berlin.de