

Probabilistic Aspects of Finance

HANS FÖLLMER¹ and ALEXANDER SCHIED²

¹*Institut für Mathematik, Humboldt-Universität, 10099 Berlin, Germany*
E-mail: ¹foellmer@math.hu-berlin.de

²*Institut für Mathematik, Universität Mannheim, 68131 Mannheim, Germany*
E-mail: ²schied@uni-mannheim.de

In the past decades, advanced probabilistic methods have had significant impact on the field of finance, both in academia and in the financial industry. Conversely, financial questions have stimulated new research directions in probability. In this survey paper, we review some of these developments and point to some areas that might deserve further investigation. We start by reviewing the basics of arbitrage pricing theory, with special emphasis on incomplete markets and on the different roles played by the “real-world” probability measure and its equivalent martingale measures. We then focus on the issue of model ambiguity, also called Knightian uncertainty. We present two case studies in which it is possible to deal with Knightian uncertainty in mathematical terms. The first case study concerns the hedging of derivatives, such as variance swaps, in a strictly pathwise sense. The second one deals with capital requirements and preferences specified by convex and coherent risk measures. In the final two sections we discuss mathematical issues arising from the dramatic increase of algorithmic trading in modern financial markets.

Keywords: Arbitrage pricing theory, hedging, incomplete market, superhedging, Knightian uncertainty, model uncertainty, pathwise Itô calculus, variance swap, monetary measure of risk, convex risk measure, coherent risk measure, algorithmic trading, price impact, market impact model.

1. The coin tossing view of finance and the appearance of Brownian motion

The systematic use of advanced probabilistic methods in the context of academic Finance begins in the mid-sixties. It was pioneered at M.I.T. by Paul Samuelson [92] and greatly stimulated by the rediscovery of “Théorie de la Spéculation”, the doctoral thesis [5] of Louis Bachelier, that had been defended in Paris in 1900 based on a report by Henri Poincaré. In this thesis, Brownian motion makes its appearance as a mathematical model for the price fluctuations of a liquid financial asset. Arguing that prices should remain positive, Samuelson proposed to use geometric Brownian motion, which soon became a standard reference model. In 1973, Black and Scholes [9] and Merton [82] derived their seminal formula for the price of a call-option in this setting.

Why does Brownian motion appear in the financial context? Here is a first rough argument. At each fixed time, the price of a stock could be seen as a temporary equilibrium resulting from a large number of decisions to buy or sell, made in a random and more or less independent manner: Many coins are thrown successively, and so Brownian motion should arise as a manifestation of the central limit theorem. This is the “Coin-Tossing View of Finance”, as it is called by J. Cassidy in *How Markets Fail* [17]. This rough argument can be refined by using microeconomic assumptions on the behavior of agents and on the ways they generate a random demand, and then the application of an invariance principle typically yields a description of the price fluctuation as a solution of a stochastic differential equation driven by Brownian motion or, more generally, by a Lévy process; see, for example, [48] and the references therein.

At this point, however, it is instructive to recall the following caveat of Poincaré in *Science et Méthode* [90] as quoted in [71]:

When men are in close touch with each other, they no longer decide randomly and independently of each other, they react to the others. Multiple causes come into play which trouble them and pull them from side to side. But there is one thing that these influences cannot destroy and that is their tendency to behave like Panurge’s sheep. And it is that which is preserved.

Thus we find, right at the beginning of the use of modern probabilistic methods in finance, a warning sign pointing to interaction and herding effects which may render invalid a straightforward application of the central limit theorem.

In his “Three essays on Capital Markets” [78], David Kreps uses a different kind of argument, where geometric Brownian motion appears as a *rational expectations equilibrium*. Suppose that agents compute their demand by maximizing expected utility. If their preferences are given by power utility functions, and if their subjective expectations are described by geometric Brownian motion, then the resulting price equilibrium would indeed be a geometric Brownian motion. Thus geometric Brownian motion is described as a fixed point for an aggregation problem based on the preferences and expectations of highly sophisticated agents. Here again, Poincaré’s caveat throws some doubt on the assumptions of rationality implicit in such an argument.

Bachelier himself does not invoke the central limit theorem, nor does he argue in terms of expected utility. Instead he starts out with a simple equilibrium argument: “*It seems that the market, that is to say, the set of speculators, must not believe in a given instant in either a rise or a fall, since for each quoted price there are as many buyers as sellers.*” As a result, “*the mathematical expectation of the speculator is zero.*” Stated in modern terms, Bachelier insists that the price process should be a martingale under a probability measure P^* which describes the market’s aggregate belief. Assuming continuous paths and adding a stationarity requirement for the increments, it follows that the price process is indeed a Brownian motion.

What is the current mainstream view? To begin with, there is a broad interdis-

plinary consensus across departments of Mathematics, Finance, and Economics that the discounted price fluctuation of a liquid financial asset should be viewed as a stochastic process $X = (X_t)_{0 \leq t \leq T}$ on some underlying probability space (Ω, \mathcal{F}, P) . The intuition is typically objectivistic: Such a probability measure P exists, it can be identified at least partially by statistical and econometric methods, and it should satisfy certain *a priori* constraints. These constraints correspond to some degree of market efficiency. In its strongest form, market efficiency would require that X is a martingale under P . In the mainstream view, however, a weaker and more flexible version of market efficiency is assumed, namely the absence of safe (and not just statistical) arbitrage opportunities. In other words, the price process should not admit any trading strategy that produces a positive expected gain over the risk free return without any downside risk. If this is made precise in a suitable manner, the absence of arbitrage opportunities can be characterized by the existence of an equivalent *martingale measure*, i.e., a probability measure P^* equivalent to P such that the properly discounted price process X is a (local) martingale under P^* . This characterization is often called the Fundamental Theorem of Asset Pricing. A preliminary version appears in Harrison and Kreps [58], and its definitive form is due to Delbaen and Schachermayer [27; 28; 29]; see also Kabanov [68] and Yan [108].

Thus an economic assumption, namely the absence of arbitrage opportunities, guarantees that

$$\mathcal{P}^* \neq \emptyset,$$

if we denote by \mathcal{P}^* the set of equivalent martingale measures P^* . Due to well-known results of Jacod, Yor and others in the “general theory” of stochastic processes of the 70s and 80s, this implies that the process X is a semimartingale under the original measure P , and hence a stochastic integrator in the sense of Bichteler and Dellacherie. This allows one to apply the techniques of Itô calculus. Moreover it follows, due to a line of arguments initiated by Wolfgang Doeblin [32] and completed by I. Monroe [84; 85], that X is a Brownian motion up to a random time change. In this way Brownian motion reappears in the present general setting, although not necessarily in a very explicit manner.

2. Derivatives and the paradigm of perfect hedging

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 maturity T has payoff $H(\omega) = (X_T(\omega) - c)^+$. What is the fair price which should be paid by the buyer of such a contingent claim H ? In other words, what is the fair deterministic equivalent to the uncertain outcome 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 expected value

$$E_P[H] = \int H dP$$

of the random variable H with respect to the resulting probability measure P . Following Daniel Bernoulli [6], one might want to add a risk premium in order to take account of risk aversion. More precisely, one could describe risk aversion by a strictly increasing and concave utility function u and compute the price $\pi(H)$ of H as the certainty equivalent $u^{-1}(E_P[u(H)])$. The difference $\pi(H) - E_P[u(H)]$, which is positive by Jensen's inequality, is then interpreted as a risk premium. But in our present financial context and under the following uniqueness assumption (1), the basic insight of Black and Scholes [9] and Merton [82] leads to a quite different result. In particular there will be no reason to argue in favor of a risk premium because the following argument shows that there is no intrinsic risk in that case.

Consider a financial market model such that $\mathcal{P}^* \neq \emptyset$. In many situations, and in particular for simple diffusion models such as geometric Brownian motion, the equivalent martingale measure is in fact unique, that is.,

$$|\mathcal{P}^*| = 1. \tag{1}$$

Uniqueness of the equivalent martingale measure implies that the model is *complete* in the following sense: Any contingent claim H can be represented, P -almost surely, in the form

$$H = V_0 + \int_0^T \xi_t dX_t \tag{2}$$

with some constant V_0 and some predictable process $\xi = (\xi_t)_{0 \leq t \leq T}$ such that the stochastic integral makes sense. For simple diffusion models such as geometric Brownian motion, this representation follows from Itô's theorem that functionals of Brownian motion can be represented as stochastic integrals of Brownian motion; see [91] for the general case. Since the expectation of the stochastic integral under the equivalent martingale measure P^* is zero, the constant V_0 is given by $V_0 = E^*[H]$.

In financial terms, the representation (2) amounts to a perfect replication of the contingent claim by means of a dynamic trading strategy. Indeed, Itô's non-anticipative construction of the stochastic integral allows one to interpret the stochastic integral in (2) as the cumulative net gain generated by the self-financing trading strategy consisting in holding ξ_t units of the underlying asset at each time t . The constant amount V_0 can now be viewed as the initial capital which is needed for a perfect replication, or a *perfect hedge*, of the contingent claim. But this implies that the unique arbitrage-free price of the claim is given by

$$\pi(H) = V_0 = E^*[H], \tag{3}$$

since any other price would offer the opportunity for a gain without any risk. If, for example, the actual price were higher then one could sell the claim at that price, use the smaller amount V_0 to implement the hedging strategy which generates the random amount H which has to be paid in the end, and retain the difference between the price and V_0 as a risk-free gain.

Thus, the uniqueness assumption (1) yields a simple answer to the problem of pricing and hedging financial derivatives. Note that the answer only involves the unique equivalent martingale measure P^* . The role of the probability measure P^* is to serve as a sophisticated consistency check for the pricing of assets, not for the purpose of prediction. The original probability measure P was meant to serve that purpose, but here it matters only insofar as it fixes a class of null sets. As we are going to see in Section 5.1 below, we can actually eliminate P completely if we are ready to restrict the space of possible scenarios.

3. Incompleteness as a source of new probabilistic problems

As soon as a financial market model becomes more realistic by admitting that there are more sources of uncertainty than traded financial instruments, the equivalent martingale measure is no longer unique, and this implies

$$|\mathcal{P}^*| = \infty.$$

As a result, the paradigm of a perfect hedge breaks down, and intrinsic risks appear at the level of derivatives. The model is then called *incomplete*. From a mathematical point of view, incompleteness has turned out to be a rich source of new problems in Stochastic Analysis. In particular it has motivated new versions of probabilistic decomposition theorems such as the Kunita-Watanabe decomposition and the Doob-Meyer decomposition.

Consider a derivative with non-negative payoff H and maturity date T . An admissible hedging strategy is given by an initial capital V_0 and a predictable process ξ such that the resulting portfolio process V defined by

$$V_t = V_0 + \int_0^t \xi_s dX_s \tag{4}$$

remains non-negative. At the maturity date T , any such strategy yields a decomposition

$$H = V_T + C_T$$

of the contingent claim into a part which is perfectly hedged, and hence priced by arbitrage as in the preceding section, and a remaining *hedging error* C_T . Different economic

preferences induce different choices of the strategy, and hence a different decomposition of the claim.

Suppose one wants to minimize the hedging error in a mean-square sense with respect to the given probability measure P . This will amount to a projection in the space $L^2(P)$ of the contingent claim H onto a sub-space of stochastic integrals. Under the strong form of the efficient market hypothesis, that is $P \in \mathcal{P}^*$, this projection problem is solved by using the Kunita-Watanabe decomposition in the space of square-integrable martingales; see [50]. If one drops this assumption and considers the case $P \notin \mathcal{P}^*$, the resulting decomposition problem can often be reduced to an application of the Kunita-Watanabe representation with respect to a suitable *minimal* martingale measure; cf. [47]. More generally, methods of mean-variance hedging for incomplete financial markets have been a source of new versions of the Kunita-Watanabe decomposition and of new results on closure properties of spaces of stochastic integrals with respect to a semimartingale; see, for example, the surveys [49] and [102].

From a financial point of view, however, the mean-variance approach fails to capture a basic asymmetry: The main purpose is to control the *shortfall*, defined as the positive part $C_T^+ = (H - V_T)^+$ of the hedging error. If one insists on keeping the shortfall down to 0 then one is led to a remarkable new extension of the Doob-Meyer decomposition. Consider a right-continuous version U of the process

$$U_t = \operatorname{ess\,sup}_{P^* \in \mathcal{P}^*} E^*[H | \mathcal{F}_t], \quad 0 \leq t \leq T. \quad (5)$$

Now note that U is a \mathcal{P}^* -supermartingale, that is, a supermartingale under any $P^* \in \mathcal{P}^*$. As shown in increasing generality in [37; 76; 42], any non-negative \mathcal{P}^* -supermartingale U admits a decomposition of the form

$$U_t = U_0 + \int_0^t \xi_s dX_s - A_t \quad (6)$$

with some increasing optional (but in general not predictable) process A . But the stochastic integral is a \mathcal{P}^* -local martingale, and so (6) can be viewed as a new version of the classical Doob-Meyer decomposition that holds simultaneously for all $P^* \in \mathcal{P}^*$. In the special case (5), this optional decomposition can be interpreted as a *superhedging* procedure: Starting with the initial capital $V_0 = U_0$, applying the trading strategy ξ and sequentially withdrawing the cumulative amount A_t from the generated portfolio value V_t defined in (4), one ends up with the final value $U_T = H$. Dually, U_t can be characterized as the minimal capital that is needed at time t in order to cover the contingent claim H by a dynamic trading strategy run from time t up to T .

The superhedging approach may tie down a large capital amount in order to stay on the safe side, and therefore it is usually seen as too conservative. But the mathematics of superhedging remains important even if zero tolerance for a shortfall is relaxed. Suppose,

for example, that one imposes some bound for the expected loss $E_P[\ell(C_T^+)]$, defined in terms of some convex loss function ℓ . Then the resulting problem of *efficient hedging* can be split into a statistical decision problem, which is solved by a randomized test φ , and a dynamic superhedging problem for the modified claim $\tilde{H} = \varphi H$; see [43].

More generally, the efficient hedging problem can be embedded into a problem of dynamic portfolio optimization for incomplete financial markets, where the criterion is usually formulated in terms of expected utility. There is a rich literature on such dynamic optimization problems, from the point of view of both optimal stochastic control as in [69; 70; 104] and convex duality as in [74; 75]

Note that in these optimization problems for incomplete financial markets the probability measure P does come in explicitly, in contrast to the superhedging approach. But it does so at the level of preferences, namely in the form of expected utility. As soon as one admits model uncertainty and considers robust preferences as described in Section 5.2 below, new problems of robust optimization arise; see, for example, [40; 61; 95; 96] and the survey [46]. Another new direction consists in analyzing the temporal dynamics of preference structures as in [86; 87].

4. P versus P^*

As we have seen, the standard setting in mathematical finance is probabilistic, and it involves two types of probability measures. On the one hand, it assumes that there is an objective probability measure P , often called “real world” or “historical” probability measure. On the other hand, the absence of arbitrage implies the existence of an equivalent martingale measure P^* , which should be interpreted as a consistent price system that reflects the present “market’s belief”. From a mathematical point of view, the coexistence of these measures and the explicit description of their mutual densities is a rich source of technical exercises, and the Girsanov transformation allows one to move freely back and forth between P and P^* . At a conceptual level, however, there is a crucial difference between their roles.

The probability measure P is usually seen as a probabilistic model that tries to capture typical patterns observed in the past; under implicit stationarity assumptions, it is then used as a forward-looking prediction scheme. While it is often admitted that any specific choice of P involves a considerable amount of model risk, it is widely believed that a true probability measure exists, and that probabilistic models are getting better in approaching that reality. Bruno de Finetti [25; 26], however, would argue that the problem is more fundamental than the issue of model risk. He would put in doubt that it makes any sense to associate an objective probability $P[A]$ to a financial event of the following type:

$$A = \{\text{the sovereign bond with ISIN } x \text{ will not default}\}.$$

On the other hand, a probability $P^*[A]$, or rather an expectation $E^*[H]$ of the discounted future cash flow H generated by the bond, is assigned each day on the financial market, either directly through the present market price of the bond or by the prices of instruments such as credit default swaps (CDS) that provide insurance against a default of the bond. Thus the probability measure P^* reflects the aggregate odds of a large number of bets made on the market. This is in accordance with de Finetti's claim that *probability does not exist*, but that one can of course take bets on a given event at certain odds. De Finetti imposes consistency rules for the odds specified for different bets, and he uses an emphatic "you" to stress the subjective nature of the resulting probability measure P^* . At the level of a single agent, these consistency rules may be viewed as an overly optimistic rationality requirement. But if we replace de Finetti's "you" by "the financial market", this requirement becomes more compelling since the market is more efficient in enforcing consistency via arbitrage than any given individual. In fact, there is a close connection, both at the conceptual and technical levels, between the fundamental theorem of asset pricing and de Finetti's reconstruction of a probability measure P^* from a consistent system of bets; see, for example, [10; 94].

Apart from such foundational aspects, the attempts of predicting financial developments in terms of an "objective" probability measure P can hardly be described as a success story, especially in view of the recent financial crisis. On the other hand, a lot is known, at any given time t , about the market's present predictions of future developments in terms of a martingale measure P_t^* . More precisely, the market's view at time t is given by the conditional probability distribution

$$P_t^*[\cdot | \mathcal{F}_t] \quad \text{on} \quad \hat{\mathcal{F}}_t \tag{7}$$

where \mathcal{F}_t is the σ -field describing the available information at time t , and $\hat{\mathcal{F}}_t$ is the σ -field generated by the pay-offs of traded contingent claims with maturities $T > t$. Present prices of call or put options with maturity T provide information about the marginal distribution of $P_t^*[\cdot | \mathcal{F}_t]$ at time $T > t$, and present prices of more exotic options provide information about the multidimensional marginals. This forward-looking "lecture du marché" is an important part of current quantitative analysis.

At any given time t , the market's present view of the future as expressed in the conditional pricing measure $P_t^*[\cdot | \mathcal{F}_t]$ is consistent across different claims, and in particular it is time-consistent across different maturities $T > t$. But this consistent picture may change from day t to day $t+1$, and it may do so in a manner which is *not* time-consistent. Time-consistency across different dates t may of course be desirable from a normative point of view, and it is usually taken for granted in the mathematical finance literature. In mathematical terms, it amounts to the requirement that the conditional distributions in (7) all belong to the same martingale measure $P^* \in \mathcal{P}^*$. In the virtual world of a complete financial market model, time-consistency would thus hold automatically, due to the fact that the equivalent martingale measure is unique. In the larger world of incomplete

financial market models, and a fortiori in reality, one should expect time-inconsistency. In our standard framework, this would be described by a flow in the space \mathcal{P}^* of martingale measures. This flow could be continuous, but it also could include jumps corresponding to abrupt regime changes.

Let us denote by \mathcal{P}_{UI}^* the class of martingale measures $P^* \in \mathcal{P}^*$ such that the price fluctuation X is a uniformly integrable martingale under P^* . Typically, both \mathcal{P}_{UI}^* and $\mathcal{P}_{NUI}^* := \mathcal{P}^* \setminus \mathcal{P}_{UI}^*$ are nonempty. The behavior of X under a measure $P^* \in \mathcal{P}_{NUI}^*$ is often interpreted as a bubble; cf. [67; 66]. A regime switch from an initial martingale measure $P_0^* \in \mathcal{P}_{UI}^*$, which does not exhibit a bubble, to another martingale measure $P_1^* \in \mathcal{P}_{NUI}^*$ would thus describe the sudden appearance of a bubble as in [66]. But the flow in the space \mathcal{P}^* could also move slowly from P_0^* to P_1^* as in [7], and this would induce the slow birth of a bubble as a submartingale.

A deeper economic understanding of the dynamics of P_t^* would involve the microstructure of financial markets, i.e., the dynamic behavior of agents with heterogeneous and interacting preferences and expectations, with special emphasis on the “herding” effects which are driving bubbles and crashes. So far, there are various toy models, such as [41] and the references therein, which try to capture some of these effects. But really compelling microstructure models which offer serious possibilities for real-world prediction are not yet in sight.

There is, however, an increasing need to complement the classical microeconomic picture of noise traders and information traders by taking into account a variety of trading algorithms which are actually used on the financial market. In a way, this may make the analysis of the resulting price dynamics more tractable, since the structure of trading algorithms is more transparent and easier to model than the behavioral characteristics of individual agents. While the social utility of such algorithms may be debatable, it is important to understand their effects as clearly as possible in mathematical terms. In particular, such an understanding is crucial for any attempts to design an intelligent regulatory framework that does not create new arbitrage opportunities and thereby new sources of instability in the financial system. In Sections 6 and 7 we are going to describe some of the simplest mathematical issues which appear in connection with the interaction of trading algorithms.

5. Knightian Uncertainty

In recent years, there has been an increasing awareness, both among practitioners and in academia, of the problems caused by an excessive reliance on a specific probabilistic model and by the resulting “control illusion”; see, for example, Section 4.9 in [60]. As a result, there is a renewed focus on the issue of *model uncertainty* or *model ambiguity*, also called *Knightian uncertainty* in honor of Frank Knight [73], who introduced the

distinction between “risk” and “uncertainty” in the context of economic decision theory. Here, “risk” refers to situations where something is known about the probability measure P (“known unknowns”), while “uncertainty” refers to situations where this is not the case (“unknown unknowns”). In its analysis of the recent subprime crisis, the Turner Review [105] distinguishes between “mathematically modellable risk” and Knightian uncertainty, and thus seems to suggest that Knightian uncertainty is beyond the scope of mathematical analysis. We do not share this conclusion. To the contrary, we see Knightian uncertainty as a rich source of new mathematical problems. This is illustrated by two recent developments, where model uncertainty is taken into account explicitly. In Section 5.1 we show how some key hedging arguments in mathematical finance can be developed without even introducing any probability measure. Another example is the specification of capital requirements and of preferences in terms of convex risk measures, described in Section 5.2. Here the analysis is not tied to the specific choice of a probability measure. Instead, one considers a whole class of probabilistic models and takes a conservative worst-case approach.

5.1. Probability-free hedging

Consider a financial market with one risky and one riskless asset. In mainstream finance, the price evolution of the risky asset is usually modeled as a stochastic process defined on some probability space. Here, however, we are going to work in a strictly pathwise setting. All we assume is that the evolution of asset prices is given by one single continuous nonnegative trajectory $(X_t)_{0 \leq t \leq T}$. As before, we will suppose for simplicity that the prices of the riskless asset, or “bond”, are given by $B_t = 1$ for all t .

Now we discuss the possibility of dynamic trading in such a market. To this end, consider a trading strategy $(\xi_t, \eta_t)_{0 \leq t \leq T}$, where ξ_t describes the number of shares in the risky asset and η_t the number of shares in the bond held at time t . The value of the portfolio (ξ_t, η_t) is given by

$$V_t = \xi_t X_t + \eta_t B_t = \xi_t X_t + \eta_t. \quad (8)$$

To discuss investment or hedging strategies in this framework, it is important to define self-financing trading strategies. Passing to the continuous-time limit from a discrete-time framework suggests that the strategy $(\xi_t, \eta_t)_{0 \leq t \leq T}$ should be called self-financing if the value process from (8) satisfies the relation

$$V_t = V_0 + \int_0^t \xi_s dX_s, \quad 0 \leq t \leq T, \quad (9)$$

where the integral is the limit of nonanticipative Riemann sums:

$$\int_0^t \xi_s dX_s := \lim_{n \uparrow \infty} \sum_{t_i^n \leq t} \xi_{t_{i-1}^n} (X_{t_i^n} - X_{t_{i-1}^n}). \quad (10)$$

Here we can take for instance $t_i^n = i2^{-n}$. According to the results in [38], this is possible when the trajectory X admits a continuous quadratic variation

$$[X]_t = \lim_{n \uparrow \infty} \sum_{\substack{t_i^n \leq t \\ t_i^n \leq t}} (X_{t_i^n} - X_{t_{i-1}^n})^2, \quad 0 \leq t \leq T,$$

and if ξ is of the form $\xi_t = g(X_t, A_t^1, \dots, A_t^k)$ for a continuous function g , which is differentiable in its first argument, and for continuous trajectories $(A_t^i)_{0 \leq t \leq T}$ of bounded variation. In this case, it was shown in [38] that Itô's formula holds for any C^2 -function f in the following strictly pathwise sense:

$$f(X_t) = f(X_0) + \int_0^t f'(X_s) dX_s + \frac{1}{2} \int_0^t f''(X_s) d[X]_s. \quad (11)$$

Note that the second integral in (11) can be defined as a classical Stieltjes integral, since $[X]_t$ is a nondecreasing function of t .

As pointed out in [39], it follows immediately that a non-constant trajectory X must have nontrivial quadratic variation so as to exclude arbitrage opportunities. Indeed, otherwise (11) reduces to the standard fundamental theorem of calculus, $f(X_t) = f(X_0) + \int_0^t f'(X_s) dX_s$, and by (9) the self-financing strategy $\xi_t = 2(X_t - X_0)$ and $\eta_t = (X_t - X_0)^2 - \xi_t X_t$ will generate the strictly positive wealth $V_t = (X_t - X_0)^2$ out of the initial capital $V_0 = 0$.

The probability-free trading framework sketched above can for instance be used to analyze the hedging error and the robustness of model-specific hedging strategies such as in [36] or [98]. In some special cases, it is even possible to find completely model-independent hedging strategies. We will illustrate this now for the case of a *variance swap* by transferring arguments from [88] and [33] to our probability-free setting. A variance swap is a path-dependent financial derivative with payoff

$$H = \sum_{i=1}^n (\log X_{t_{i+1}} - \log X_{t_i})^2$$

at time T , where $0 < t_1 < \dots < t_n = T$ are fixed time points. These time points are often chosen so that X_{t_i} is the closing price of the risky asset at the end of the i^{th} trading day; see, e.g., [15; 14; 53] for background on variance swaps. When n is large enough, the payoff of the variance swap can thus be approximated by the quadratic variation of $\log X$, i.e.,

$$H \approx [\log X]_T = \int_0^T \frac{1}{X_t^2} d[X]_t. \quad (12)$$

Here, the second identity follows, e.g., from Proposition 2.2.10 in [103]. On the other hand, applying Itô's formula (11) to the function $f(x) = \log x$ yields

$$\log X_T - \log X_0 = \int_0^T \frac{1}{X_t} dX_t - \frac{1}{2} \int_0^T \frac{1}{X_t^2} d[X]_t. \quad (13)$$

Putting (12) and (13) together implies that

$$H \approx \int_0^T \frac{1}{X_t^2} d[X]_t = 2 \log X_0 - 2 \log X_T + 2 \int_0^T \frac{1}{X_t} dX_t. \quad (14)$$

The Itô integral on the right-hand side of (14) can be regarded as the terminal value of the self-financing trading strategy that has zero initial investment and otherwise consists in holding $\xi_t = 2/X_t$ shares of the risky asset at each time t . To interpret the two logarithmic terms in (14), we apply the Breeden–Litzenberger formula,

$$h(X_T) = h(X_0) + h'(X_0)(X_T - X_0) + \int_0^{X_0} (K - X_T)^+ h''(K) dK + \int_{X_0}^{\infty} (X_T - K)^+ h''(K) dK \quad (15)$$

(e.g., [45], Exercise 1.3.3) to the function $h(x) = \log x$ and obtain

$$H \approx -\frac{2}{X_0}(X_T - X_0) + \int_0^{X_0} (K - X_T)^+ \frac{2}{K^2} dK + \int_{X_0}^{\infty} (X_T - K)^+ \frac{2}{K^2} dK + 2 \int_0^T \frac{1}{X_t} dX_t. \quad (16)$$

That is, H can be hedged by selling $2/X_0$ zero-price forward contracts, holding portfolios consisting of $2/K^2 dK$ “out-of-the-money” put and call options with maturity T for each strike K , and using the self-financing trading strategy with $\xi_t = 2/X_t$. The most remarkable feature of this hedging strategy is that it is model-independent. That is, (16) is valid independently of possible probabilistic dynamics of the price process X . The hedging strategy is therefore not subject to model risk that might result from a misspecification of such probabilistic dynamics.

Similar results as obtained for variance swaps are valid for so-called *Gamma* or *entropy swaps* with payoff

$$\sum_{i=1}^n X_{t_i} (\log X_{t_{i+1}} - \log X_{t_i})^2.$$

and also for *corridor variance swaps* with payoff

$$\sum_{i=1}^n \mathbf{1}_{\{A \leq X_{t_i} \leq B\}} (\log X_{t_{i+1}} - \log X_{t_i})^2,$$

for some real numbers A, B , with $A < B$. See also [24] for further extensions.

Note that the Breeden–Litzenberger formula (15) can be regarded as a simple static, and hence model-free, hedge for the option $h(X_T)$ in terms of standard “plain vanilla” put and call options. In some cases, static hedges (or superhedges) can also be constructed for path-dependent derivatives such as barrier or lookback options; see, e.g., [12; 23; 62].

If uncertainty is restricted to a suitable class of scenarios, the strictly pathwise approach can also be used to formulate the crucial hedging argument of Section 2 in a

probability-free manner. To this end, we fix a continuous volatility function $\sigma(x, t) > 0$ on $[0, \infty) \times [0, T]$ and restrict the possible scenarios to the set Ω_σ of all nonnegative continuous functions ω on $[0, T]$ such that the coordinate process $X_t(\omega) = \omega(t)$ admits an absolutely continuous quadratic variation $d[X(\omega)]_t = \sigma^2(X_t(\omega), t)X_t^2(\omega) dt$. Consider a derivative of the form $H = h(X_T)$. As explained in [8] or [39], we can now use the time-dependent extension of the pathwise Itô formula (11) to construct a perfect hedge of the form $\xi_t(\omega) = F_x(X_t(\omega), t)$, where F solves an appropriate parabolic equation with boundary condition $F(x, T) = h(x)$. Moreover, a theorem of Paul Lévy implies that there is exactly one probability measure P^* on the space Ω_σ such that the coordinate process X becomes a martingale under P^* . The price of the derivative H , defined as the initial cost of the perfect hedge, can then be computed as in (3) as the expected value $E^*[H]$ of H under the measure P^* .

In order to extend the preceding construction to more exotic options, one can use a strictly pathwise version of Malliavin calculus as recently developed in [34] and [21]. For an alternative pathwise approach in terms of *rough paths* see [80; 52].

5.2. Monetary risk measures

The *capital requirement* associated with the profits and losses, or P&L, of a given financial position is specified as the minimal capital that should be added to the position in order to make that position acceptable from the point of view of a supervising agency. This idea can be formalized as follows by the notion of a *monetary measure of risk*.

The P&L describes the uncertain net monetary outcome at the end of a given trading period, and so it will be modeled as a real-valued measurable function X on a measurable space (Ω, \mathcal{F}) of possible scenarios. We fix a linear space \mathcal{X} of such P&Ls and a nonempty subset $\mathcal{A} \subset \mathcal{X}$ associated with those positions that are deemed acceptable. We require that \mathcal{A} contains all constants and that $Y \in \mathcal{A}$ whenever $Y \geq X$ for some $X \in \mathcal{A}$. The functional ρ on \mathcal{X} defined by

$$\rho(X) := \inf\{m \in \mathbb{R} \mid X + m \in \mathcal{A}\} \tag{17}$$

is then called a *monetary risk measure*, and the value $\rho(X)$ is interpreted as the capital requirement for the financial position with P&L X .

The standard example of a monetary risk measure is *Value at Risk* at some level $\lambda \in (0, 1)$. For a given probabilistic model described by a probability measure P on (Ω, \mathcal{F}) , X is deemed acceptable for Value at Risk if the probability $P[X < 0]$ of a shortfall does not exceed the level λ . The resulting monetary risk measure (17) is then given, up to a minus sign, by a λ -quantile of the distribution of X under P . Value at Risk is widely used in practice. But it also has a number of deficiencies. In particular, it does not account for the size of a possible shortfall and therefore penalizes diversification while

encouraging the concentration of risk. The recognition of these deficiencies motivated the axiomatic approach to a general theory of monetary risk measures as initiated by Artzner, Delbaen, Eber, and Heath [3] in the late nineties. But there are also other drawbacks. For instance, in reaction to the recent financial crisis, *The Turner Review - A regulatory response to the global banking crisis* [105] emphasizes an excessive reliance on a single probabilistic model P and thus raises the issue of Knightian uncertainty.

We are now going to sketch some of the key ingredients in the theory of convex risk measures. As we will see, this theory does not only address the issue that diversification should not be penalized by the capital requirement. It also provides a case study on how to deal with Knightian uncertainty in a mathematical framework.

To capture the idea that diversification should be encouraged rather than penalized by a monetary risk measure, we require that the acceptance set \mathcal{A} be convex. In this case the monetary risk measure ρ defined via (17) is called a *convex risk measure*, because convexity of \mathcal{A} is equivalent to convexity of ρ . When \mathcal{A} is even a convex cone, ρ is called a *coherent risk measure*. The notion of a coherent risk measures was introduced in the seminal paper [3]; the subsequent extension from coherent to convex risk measures was introduced independently in [59], [51], and [44]. Convex duality implies that a convex risk measure typically takes the form

$$\rho(X) = \sup_{Q \in \mathcal{Q}_\rho} \{E_Q[-X] - \alpha(Q)\}, \quad (18)$$

where \mathcal{Q}_ρ is some class of probability measures and $\alpha : \mathcal{Q}_\rho \rightarrow \mathbb{R} \cup \{+\infty\}$ is a penalty function. The capital requirement is thus determined as follows: The expected loss of a position is calculated for each probability measure $Q \in \mathcal{Q}_\rho$ and penalized by the penalty $\alpha(Q)$; then one takes the worst penalized expected loss over the class \mathcal{Q}_ρ . This procedure can be interpreted as follows in the light of model uncertainty. No probability measure is fixed in advance, but probability measures $Q \in \mathcal{Q}_\rho$ do come in via convex duality and take the role of stress tests. The set \mathcal{Q}_ρ can be regarded as a class of plausible probabilistic models, in which each model $Q \in \mathcal{Q}_\rho$ is taken more or less seriously, according to the size of the penalty $\alpha(Q)$. In this way, model uncertainty is taken into account explicitly. In the special coherent case the penalty function will vanish on \mathcal{Q}_ρ , and so the representation (18) reduces to

$$\rho(X) = \sup_{Q \in \mathcal{Q}_\rho} E_Q[-X], \quad (19)$$

that is, to the worst case expected loss over the class \mathcal{Q}_ρ .

In the context of an arbitrage-free but possibly incomplete financial market model, the superhedging risk measure,

$$\rho(X) = \sup_{P^* \in \mathcal{P}^*} E^*[-X],$$

is clearly a coherent risk measure. The corresponding acceptance set \mathcal{A} consists of all X for which one can find a dynamic trading strategy with initial capital $V_0 = 0$ and final outcome V_T such that the pay-off of the combined position $X + V_T$ is nonnegative with probability one.

In the setting of mathematical finance, the history of coherent and convex risk measures begins with the seminal paper [3], as mentioned above. In a wider mathematical context, however, there is a considerable pre-history in areas such as in game theory and Choquet integration [30; 99], robust statistics [63; 64], and actuarial premium principles [31; 57].

Risk measures have also appeared implicitly in the microeconomic theory of preferences. Preferences on the space \mathcal{X} are usually represented by some utility functional U on \mathcal{X} . Under the axioms of rationality as formulated by von Neumann and Morgenstern [106] and Savage [93], U takes the form of an *expected utility*, i.e.,

$$U(X) = E_P[u(X)] \tag{20}$$

for some increasing continuous function u and some probability measure P on (Ω, \mathcal{F}) . As shown by Gilboa and Schmeidler [56] in the late eighties, a natural relaxation of the axioms of rationality implies that the linear risk measure $-E_P[\cdot]$ in (20) should be replaced by a general coherent risk measure ρ :

$$U(X) = -\rho(u(X)) = \inf_{Q \in \mathcal{Q}_\rho} E_Q[u(X)].$$

More recently, Maccheroni, Marinacci, and Rustichini [81] have relaxed the rationality axioms even further. In their axiomatic setting, ρ is now a convex risk measure, and so the numerical representation of preferences takes the form

$$U(X) = -\rho(u(X)) = \inf_{Q \in \mathcal{Q}_\rho} \{E_Q[u(X)] + \alpha(Q)\}.$$

While classical risk aversion is captured by concavity of the utility function u , the concavity of $-\rho$ corresponds to a behavioral assumption of model uncertainty aversion; see [56], [81], and also [45].

6. Price formation, market microstructure, and the emergence of algorithmic trading

When L. Bachelier and P.A. Samuelson formulated their models of asset price processes, orders were usually executed by broker signals in trading pits. But in recent years the way in which financial markets operate has changed dramatically. We are now going to

discuss some of the new challenges for mathematical finance that are resulting from this change.

In 1971, the world's first electronic stock exchange, NASDAQ, was opened. In the subsequent decades, fostered by measures of market deregulation and technological improvements, more and more trading pits were abandoned and replaced by fully electronic exchanges. Such an electronic exchange basically operates with two different kinds of orders, *limit orders* and *market orders*. A limit order is an order to buy or sell a certain amount of shares at a specific price. It is collected in an electronic *limit order book* until there is a matching sell or buy order. A market order is an order to buy or sell a certain amount of shares at the best currently available price. It thus consumes limit orders according to price priority. When the total size of all limit orders at the best price is larger than the size of the incoming matching order, limit orders are usually executed according to a first-in first-out rule. On this microscopic level, asset price dynamics are thus represented not by a one-dimensional diffusion process but by the evolution of the entire limit order book, which, from a mathematical point of view, can be regarded as a complex queuing system. As such, it can at least in principle be modeled mathematically. With a suitable model at hand, one can try to “zoom out” of the microscopic picture and characterize the limiting dynamics of the mid price (i.e., the average between the best buy and sell limit orders) on a mesoscopic diffusion scale. This can either lead to a confirmation of the standard modeling paradigms of mathematical finance or to the discovery of new types of asset price dynamics. Initial studies concerned with such questions were conducted in [4; 11; 20; 19; 22] with, e.g., [19] finding a Bachelier-type model in the diffusion limit.

The emergence of electronic trading venues facilitated the use of computers for order placement, and soon the new phenomena of *algorithmic* and *high-frequency trading* came into existence. Today, limit order books are updated in time intervals measured in milliseconds so that no human being can possibly follow the price evolution of sufficiently liquid assets. The use of computers is hence mandatory for market makers and many other traders. As a consequence, the vast majority of orders in equity markets is now placed by computer algorithms. A good description of the current state of electronic markets is given in [79].

The computerization of financial markets led to some effects that can be regarded as potentially beneficial. For instance, the liquidity provided by high-frequency market makers and the competition between the growing number of electronic trading venues contributed to a significant decline of bid-ask spreads, thus reducing transaction costs for ordinary investors. There was also some hope that computer programs would act more rationally than human investors, in particular in critical situations, and thus avoid panic and herding behavior. These hopes, however, were seriously challenged by the *Flash Crash* of May 6, 2010. On that day, a sell order placed in a nervous market triggered a “hot-potato game” among the trading algorithms of high-frequency traders (HFTs),

which resulted in the steepest drop of asset prices ever, followed by a sharp recovery within 20 minutes. The following quote from [18, page 3] gives some indication that the Flash Crash was indeed generated by a feedback overflow between several trading algorithms:

... HFTs began to quickly buy and then resell contracts to each other—generating a “hot-potato” volume effect as the same positions were rapidly passed back and forth. Between 2:45:13 and 2:45:27, HFTs traded over 27,000 contracts, which accounted for about 49 percent of the total trading volume, while buying only about 200 additional contracts net.

It is an interesting challenge to understand the reasons why interacting trading algorithms can end up in such a “hot-potato game” and to reproduce this phenomenon in a mathematical model. As we will see in the next section, there are already some preliminary results that may be connected to this phenomenon.

Besides the possible creation of crash scenarios, there are also other aspects of electronic trading that are potentially problematic. For instance, certain predatory trading algorithms scan order signals for patterns resulting from the execution of large trades. Once such a large trade is detected, the predatory trading algorithm tries to make a profit by building up a position whose value will be increased by the price impact generated by the large trade; see [13; 16; 101]. To escape the adverse effects of price impact and predatory trading, many investors resort to so-called dark pools, in which orders are invisible to other market participants. But the fact that many dark pools derive the execution price of orders from the ‘lit’ market facilitates predatory trading techniques such as ‘fishing’, which are based on manipulating the price in the lit market; see [72; 77; 83].

7. Price impact and order execution

The key to understanding algorithmic trading and its potential benefits and risks is the phenomenon of *price impact*, i.e., the fact that the execution of a large order influences the price of the underlying asset. It is one of the basic mechanisms by which economic agents interact with the market and, thus, with each other. Spectacular cases in which price impact played an important role were the debacle of Metallgesellschaft in 1993, the LTCM crisis in 1998, or the unwinding of Jérôme Kerviel’s portfolio by Société Générale in 2008. But price impact can also be significant in much smaller trades, and it belongs to the daily business of many financial institutions.

The first step in understanding price impact is the execution of a single trade, a problem at which one can look on several scales. On a microscopic scale, one considers a trade that is small enough to be executed by placing a single order in a limit order book. When this order is placed as a market order, it will impact the limit order book by consuming limit orders and, if it is large enough, shift the corresponding best price and widen the bid-ask spread; see [1; 89; 107]. When it consists in placing or cancelling a limit

order, its quantitative impact is not as easy to describe but it is nevertheless existing. In either case, the impact of a trade is transient and will eventually diminish, a fact that becomes important on the next, mesoscopic level.

Many trades are too big to be executed in one single order and therefore need to be split in a series of smaller orders, sometimes called ‘child orders’, which are then spread out over a certain time interval. On this mesoscopic scale, trading algorithms are used to determine sizes and timing of each child order. These algorithms are typically based on a market impact model, i.e., a stochastic model for asset prices that takes into account the feedback effects of trading strategies. We refer to [55] for a survey on some models that are currently available. The problem of determining *optimal* trade execution strategies for a given cost criterion in a specific model has a rich structure and often leads to questions that are of intrinsic mathematical interest. It is, for instance, connected to the mathematical topics of finite-fuel control, Choquet capacity theory, and Dawson–Watanabe superprocesses. Let us briefly sketch the latter connection as established in [97]. When formulating the optimal trade execution problem as a stochastic control problem, the liquidation constraint translates into a singular terminal condition for the corresponding Hamilton–Jacobi–Bellman equation. This equation can be further reduced to a quasilinear parabolic partial differential equation with infinite terminal condition. But, according to [35], such equations are related to the Laplace functionals of Dawson–Watanabe superprocesses.

The existence or nonexistence and the structure of optimal trade execution strategies can also yield information on the viability of the underlying market impact model and perhaps even on the nature of price impact itself; see, e.g., [2; 54; 65]. For instance, it was shown in [2] that the price impact of single orders must decay as a convex function of time to exclude oscillatory trade execution strategies that are to some extent reminiscent of the “hot-potato game” mentioned earlier.

It should be pointed out that the market impact models currently available in the literature are all relatively simple. In particular, there is yet no model that combines both transience and nonlinearity of price impact in a truly convincing way.

On a macroscopic scale, the execution of the trade is seen in relation to the behavior of other agents—or algorithms—in the market. As mentioned above, the fact that an agent is executing a large trade can be betrayed to competitors for instance via the order signals created by the execution algorithm. When a competitor detects the execution of a large trade, it is generally believed that predatory trading, as described above, is the corresponding profit-maximizing strategy. This was also obtained as a mathematical result in [16] by analyzing a game-theoretic setting. By slightly extending this setting, however, it was found in [101] that predatory trading may become suboptimal in markets that are sufficiently ‘elastic’ in the sense that the price impact of orders decays very rapidly. In such markets it is instead beneficial for the competitor to cooperate with the large trader and to provide liquidity. A completely different pattern occurs, however,

when price impact is transient. Schöneborn [100] showed that in a discrete-time model with linear, exponentially decaying price impact the large trader and the competitor start a “hot-potato game” very similar to the one observed in the Flash Crash.

References

- [1] ALFONSI, A., FRUTH, A. and SCHIED, A. (2010). Optimal execution strategies in limit order books with general shape functions. *Quant. Finance* **10** 143–157.
- [2] ALFONSI, A., SCHIED, A. and SLYNKO, A. (2012). Order book resilience, price manipulation, and the positive portfolio problem. *SIAM J. Financial Math.* **3** 511–533.
- [3] ARTZNER, P., DELBAEN, F., EBER, J.-M. and HEATH, D. (1999). Coherent measures of risk. *Math. Finance* **9** 203–228.
- [4] AVELLANEDA, M. and STOIKOV, S. (2008). High-frequency trading in a limit order book. *Quant. Finance* **8** 217–224.
- [5] BACHELIER, L. (1995). *Théorie de la spéculation. Les Grands Classiques Gauthier-Villars. [Gauthier-Villars Great Classics]*. Éditions Jacques Gabay, Sceaux. Théorie mathématique du jeu. [Mathematical theory of games], Reprint of the 1900 original.
- [6] BERNOULLI, D. (1738). Specimen theoriae novae de mensura sortis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae* **5** 175–1926. Translated by L. Sommer: *Econometrica* **22** (1954), 23–36.
- [7] BIAGINI, F., FÖLLMER, H. and NEDELCO, S. (2011). Shifting martingale measures and the birth of a bubble as a submartingale. *Preprint*.
- [8] BICK, A. and WILLINGER, W. (1994). Dynamic spanning without probabilities. *Stochastic Process. Appl.* **50** 349–374.
- [9] BLACK, F. and SCHOLES, M. (1973). The pricing of options and corporate liabilities. *The Journal of Political Economy* 637–654.
- [10] BORKAR, V. S., KONDA, V. R. and MITTER, S. K. (2004). On De Finetti coherence and Kolmogorov probability. *Statist. Probab. Lett.* **66** 417–421.
- [11] BOVIER, A., ČERNÝ, J. and HRYNIV, O. (2006). The opinion game: stock price evolution from microscopic market modeling. *Int. J. Theor. Appl. Finance* **9** 91–111.
- [12] BROWN, H., HOBSON, D. and ROGERS, L. C. G. (2001). Robust hedging of barrier options. *Math. Finance* **11** 285–314.
- [13] BRUNNERMEIER, M. K. and PEDERSEN, L. H. (2005). Predatory trading. *Journal of Finance* **60** 1825–1863.
- [14] BÜHLER, H. (2006). Consistent variance curve models. *Finance Stoch.* **10** 178–203.
- [15] BÜHLER, H. (2006). Volatility markets: Consistent modeling, hedging and practical implementation PhD thesis, TU Berlin.

- [16] CARLIN, B. I., LOBO, M. S. and VISWANATHAN, S. (2007). Episodic liquidity crises: cooperative and predatory trading. *Journal of Finance* **65** 2235-2274.
- [17] CASSIDY, J. (2009). *How markets fail: The logic of economic calamities*. Farrar, Straus & Giroux, New York, NY.
- [18] CFTC-SEC, (2010). Findings regarding the market events of May 6, 2010. Report.
- [19] CONT, R. and DE LARRARD, A. (2010). Linking volatility with order flow: heavy traffic approximations and diffusion limits of order book dynamics. *Preprint*.
- [20] CONT, R. and DE LARRARD, A. (2013). Price dynamics in a Markovian limit order market. To appear in *SIAM J. Financial Math*.
- [21] CONT, R. and FOURNIE, D. A. (2013). Functional Ito calculus and stochastic integral representation of martingales. To appear in *Ann. Probab.*
- [22] CONT, R., KUKANOV, A. and STOIKOV, S. (2010). The price impact of order book events. *Preprint*.
- [23] COX, A. M. G. and OBLÓJ, J. (2011). Robust hedging of double touch barrier options. *SIAM J. Financial Math.* **2** 141–182.
- [24] DAVIS, M., OBLÓJ, J. and RAVAL, V. (2010). Arbitrage bounds for weighted variance swap prices. *Arxiv preprint arXiv:1001.2678*.
- [25] DE FINETTI, B. (1990). *Theory of probability. Vol. 1. Wiley Classics Library*. John Wiley & Sons Ltd., Chichester. A critical introductory treatment, Translated from the Italian and with a preface by Antonio Machi and Adrian Smith, With a foreword by D. V. Lindley, Reprint of the 1974 translation.
- [26] DE FINETTI, B. (1990). *Theory of probability. Vol. 2. Wiley Classics Library*. John Wiley & Sons Ltd., Chichester. A critical introductory treatment, Translated from the Italian and with a preface by Antonio Machi and Adrian Smith, With a foreword by D. V. Lindley, Reprint of the 1975 translation.
- [27] DELBAEN, F. and SCHACHERMAYER, W. (1994). A general version of the fundamental theorem of asset pricing. *Math. Ann.* **300** 463–520.
- [28] DELBAEN, F. and SCHACHERMAYER, W. (1998). The fundamental theorem of asset pricing for unbounded stochastic processes. *Math. Ann.* **312** 215–250.
- [29] DELBAEN, F. and SCHACHERMAYER, W. (2006). *The mathematics of arbitrage. Springer Finance*. Springer-Verlag, Berlin.
- [30] DELLACHERIE, C. (1971). Quelques commentaires sur les prolongements de capacités. In *Séminaire de Probabilités, V (Univ. Strasbourg, année universitaire 1969-1970)* 77–81. Lecture Notes in Math., Vol. 191. Springer, Berlin.
- [31] DEPPEZ, O. and GERBER, H. U. (1985). On convex principles of premium calculation. *Insurance Math. Econom.* **4** 179–189.
- [32] DOEBLIN, W. (2000). *Sur l'équation de Kolmogoroff*. Éditions Elsevier, Paris C. R. Acad. Sci. Paris Sér. I Math. **331** (2000), Special Issue.
- [33] DUPIRE, B. (1993). Model art. *Risk* **6** 118–124.

- [34] DUPIRE, B. (2009). Functional Itô calculus. *Bloomberg Portfolio Research Paper*.
- [35] DYNKIN, E. B. (1992). Superdiffusions and parabolic nonlinear differential equations. *Ann. Probab.* **20** 942–962.
- [36] EL KAROUI, N., JEANBLANC-PICQUÉ, M. and SHREVE, S. E. (1998). Robustness of the Black and Scholes formula. *Math. Finance* **8** 93–126.
- [37] EL KAROUI, N. and QUENEZ, M.-C. (1995). Dynamic programming and pricing of contingent claims in an incomplete market. *SIAM J. Control Optim.* **33** 29–66.
- [38] FÖLLMER, H. (1981). Calcul d’Itô sans probabilités. In *Seminar on Probability, XV (Univ. Strasbourg, Strasbourg, 1979/1980) (French)*. *Lecture Notes in Math.* **850** 143–150. Springer, Berlin.
- [39] FÖLLMER, H. (2001). Probabilistic aspects of financial risk. In *European Congress of Mathematics, Vol. I (Barcelona, 2000)*. *Progr. Math.* **201** 21–36. Birkhäuser, Basel.
- [40] FÖLLMER, H. and GUNDEL, A. (2006). Robust projections in the class of martingale measures. *Illinois Journal of Mathematics* **50** 439–472.
- [41] FÖLLMER, H., HORST, U. and KIRMAN, A. (2005). Equilibria in financial markets with heterogeneous agents: a probabilistic perspective. *J. Math. Econom.* **41** 123–155.
- [42] FÖLLMER, H. and KABANOV, Y. M. (1998). Optional decomposition and Lagrange multipliers. *Finance Stoch.* **2** 69–81.
- [43] FÖLLMER, H. and LEUKERT, P. (2000). Efficient hedging: cost versus shortfall risk. *Finance Stoch.* **4** 117–146.
- [44] FÖLLMER, H. and SCHIED, A. (2002). Convex measures of risk and trading constraints. *Finance Stoch.* **6** 429–447.
- [45] FÖLLMER, H. and SCHIED, A. (2011). *Stochastic finance. An introduction in discrete time*, 3rd revised and extended ed. Walter de Gruyter & Co., Berlin.
- [46] FÖLLMER, H., SCHIED, A. and WEBER, S. (2009). Robust preferences and robust portfolio choice. In *Mathematical Modelling and Numerical Methods in Finance*, (P. Ciarlet, A. Bensoussan and Q. Zhang, eds.) **15** 29–88. Amsterdam: Elsevier/North-Holland.
- [47] FÖLLMER, H. and SCHWEIZER, M. (1991). Hedging of contingent claims under incomplete information. In *Applied Stochastic Analysis (London, 1989)*. *Stochastics Monogr.* **5** 389–414. Gordon and Breach, New York.
- [48] FÖLLMER, H. and SCHWEIZER, M. (1993). A microeconomic approach to diffusion models for stock prices. *Mathematical Finance* **3** 1–23.
- [49] FÖLLMER, H. and SCHWEIZER, M. (2010). The minimal martingale measure. In *Encyclopedia of Quantitative Finance* (R. Cont, ed.) 1200–1204. Wiley.
- [50] FÖLLMER, H. and SONDERMANN, D. (1986). Hedging of nonredundant contingent claims. In *Contributions to Mathematical Economics* 205–223. North-Holland, Amsterdam.

- [51] FRITTELLI, M. and ROSAZZA GIANIN, E. (2002). Putting order in risk measures. *Journal of Banking & Finance* **26** 1473–1486.
- [52] FRIZ, P. K. and VICTOIR, N. B. (2010). *Multidimensional stochastic processes as rough paths. Cambridge Studies in Advanced Mathematics* **120**. Cambridge University Press, Cambridge. Theory and applications.
- [53] GATHERAL, J. (2006). *The Volatility Surface. A practitioner's Guide*. Wiley Finance.
- [54] GATHERAL, J. (2010). No-dynamic-arbitrage and market impact. *Quant. Finance* **10** 749–759.
- [55] GATHERAL, J. and SCHIED, A. (2012). Dynamical models of market impact and algorithms for order execution. To appear in *Handbook on Systemic Risk*.
- [56] GILBOA, I. and SCHMEIDLER, D. (1989). Maxmin expected utility with nonunique prior. *J. Math. Econom.* **18** 141–153.
- [57] GOOVAERTS, M. J., DE VYLDER, F. and HAEZENDONCK, J. (1984). *Insurance premiums*. North-Holland Publishing Co., Amsterdam. Theory and applications.
- [58] HARRISON, J. M. and KREPS, D. M. (1979). Martingales and arbitrage in multi-period securities markets. *J. Econom. Theory* **20** 381–408.
- [59] HEATH, D. (2000). Back to the future. *Plenary lecture, First World Congress of the Bachelier Finance Society, Paris*.
- [60] HELLWIG, M. (2009). Systemic risk in the financial sector: an analysis of the subprime-mortgage financial crisis. *De Economist* **157** 129–207.
- [61] HERNÁNDEZ-HERNÁNDEZ, D. and SCHIED, A. (2007). A control approach to robust utility maximization with logarithmic utility and time-consistent penalties. *Stochastic Process. Appl.* **117** 980–1000.
- [62] HOBSON, D. G. (1998). Robust hedging of the lookback option. *Finance Stoch.* **2** 329–347.
- [63] HUBER, P. J. (1981). *Robust statistics*. John Wiley & Sons Inc., New York. Wiley Series in Probability and Mathematical Statistics.
- [64] HUBER, P. J. and STRASSEN, V. (1973). Minimax tests and the Neyman-Pearson lemma for capacities. *Ann. Statist.* **1** 251–263.
- [65] HUBERMAN, G. and STANZL, W. (2004). Price manipulation and quasi-arbitrage. *Econometrica* **72** 1247–1275.
- [66] JARROW, R. A., PROTTER, P. and SHIMBO, K. (2007). Asset price bubbles in complete markets. In *Advances in Mathematical Finance. Appl. Numer. Harmon. Anal.* 97–121. Birkhäuser Boston, Boston, MA.
- [67] JARROW, R. A., PROTTER, P. and SHIMBO, K. (2010). Asset price bubbles in incomplete markets. *Math. Finance* **20** 145–185.
- [68] KABANOV, Y. M. (1997). On the FTAP of Kreps-Delbaen-Schachermayer. In *Statistics and control of stochastic processes (Moscow, 1995/1996)* 191–203. World Sci. Publ., River Edge, NJ. MR1647282 (2000g:91026)

- [69] KARATZAS, I., LEHOCZKY, J. P. and SHREVE, S. E. (1987). Optimal portfolio and consumption decisions for a “small investor” on a finite horizon. *SIAM J. Control Optim.* **25** 1557–1586.
- [70] KARATZAS, I. and SHREVE, S. E. (1998). *Methods of mathematical finance. Applications of Mathematics (New York)* **39**. Springer-Verlag, New York.
- [71] KIRMAN, A. (2010). The economic crisis is a crisis for economic theory. *CESifo Economic Studies* **56** 498–535.
- [72] KLÖCK, F., SCHIED, A. and SUN, Y. (2011). Price manipulation in a market impact model with dark pool. *Preprint*.
- [73] KNIGHT, F. (1921). *Risk, uncertainty, and profit*. Houghton Mifflin, Boston.
- [74] KRAMKOV, D. and SCHACHERMAYER, W. (1999). The asymptotic elasticity of utility functions and optimal investment in incomplete markets. *Ann. Appl. Probab.* **9** 904–950.
- [75] KRAMKOV, D. and SCHACHERMAYER, W. (2003). Necessary and sufficient conditions in the problem of optimal investment in incomplete markets. *Ann. Appl. Probab.* **13** 1504–1516.
- [76] KRAMKOV, D. O. (1996). Optional decomposition of supermartingales and hedging contingent claims in incomplete security markets. *Probab. Theory Related Fields* **105** 459–479.
- [77] KRATZ, P. and SCHÖNEBORN, T. (2010). Optimal liquidation in dark pools. *Preprint*.
- [78] KREPS, D. M. (1979). *Three essays on capital markets*. Institute for Mathematical Studies in the Social Sciences.
- [79] LEHALLE, C.-A. (2012). Market microstructure knowledge needed to control an intra-day trading process. *Forthcoming in Handbook on Systemic Risk*.
- [80] LYONS, T. and QIAN, Z. (2002). *System control and rough paths*. *Oxford Mathematical Monographs*. Oxford University Press, Oxford. Oxford Science Publications.
- [81] MACCHERONI, F., MARINACCI, M. and RUSTICHINI, A. (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* **74** 1447–1498.
- [82] MERTON, R. C. (1973). Theory of rational option pricing. *Bell J. Econom. and Management Sci.* **4** 141–183. MR0496534 (58 ##15058)
- [83] MITTAL, H. (2008). Are you playing in a toxic dark pool? A guide to preventing information leakage. *Journal of Trading* **3** 20-33.
- [84] MONROE, I. (1972). On embedding right continuous martingales in Brownian motion. *Ann. Math. Statist.* **43** 1293–1311.
- [85] MONROE, I. (1978). Processes that can be embedded in Brownian motion. *Ann. Probab.* **6** 42–56.
- [86] MUSIELA, M. and ZARIPHOUPOULOU, T. (2009). Portfolio choice under dynamic investment performance criteria. *Quant. Finance* **9** 161–170.

- [87] MUSIELA, M. and ZARIPHOUPOULOU, T. (2010). Stochastic partial differential equations and portfolio choice. In *Contemporary Quantitative Finance* 195–216. Springer, Berlin.
- [88] NEUBERGER, A. (1994). The log contract. *The Journal of Portfolio Management* **20** 74–80.
- [89] OBIZHAEVA, A. and WANG, J. (2013). Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets* **16** 1–32.
- [90] POINCARÉ, H. (1908). Science et méthode. *Revue scient. (5)* **10** 417–423.
- [91] REVUZ, D. and YOR, M. (1999). *Continuous martingales and Brownian motion*, third ed. *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]* **293**. Springer-Verlag, Berlin.
- [92] SAMUELSON, P. A. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial management review* **6**.
- [93] SAVAGE, L. J. (1972). *The foundations of statistics*, revised ed. Dover Publications Inc., New York.
- [94] SCHERVISH, M. J., SEIDENFELD, T. and KADANE, J. B. (2008). The fundamental theorems of prevision and asset pricing. *International journal of approximate reasoning* **49** 148–158.
- [95] SCHIED, A. (2005). Optimal investments for robust utility functionals in complete market models. *Math. Oper. Res.* **30** 750–764.
- [96] SCHIED, A. (2007). Optimal investments for risk- and ambiguity-averse preferences: a duality approach. *Finance Stoch.* **11** 107–129.
- [97] SCHIED, A. (2013). A control problem with fuel constraint and Dawson-Watanabe superprocesses. To appear in: *Ann. Appl. Probab.*
- [98] SCHIED, A. and STADJE, M. (2007). Robustness of delta hedging for path-dependent options in local volatility models. *J. Appl. Probab.* **44** 865–879.
- [99] SCHMEIDLER, D. (1986). Integral representation without additivity. *Proc. Amer. Math. Soc.* **97** 255–261.
- [100] SCHÖNEBORN, T. (2008). Trade execution in illiquid markets. Optimal stochastic control and multi-agent equilibria. Ph.D. thesis, TU Berlin.
- [101] SCHÖNEBORN, T. and SCHIED, A. (2009). Liquidation in the face of adversity: stealth vs. sunshine trading. Preprint.
- [102] SCHWEIZER, M. (2010). Mean-variance hedging. In *Encyclopedia of Quantitative Finance* (R. Cont, ed.) 1177–1181. Wiley.
- [103] SONDERMANN, D. (2006). *Introduction to stochastic calculus for finance. Lecture Notes in Economics and Mathematical Systems* **579**. Springer-Verlag, Berlin.
- [104] STOIKOV, S. F. and ZARIPHOUPOULOU, T. (2005). Dynamic asset allocation and consumption choice in incomplete markets. *Australian Economic Papers* **44** 414–454.
- [105] TURNER, A. (2009). The Turner Review: A regulatory response to the global

- banking crisis. *FSA, March*.
- [106] VON NEUMANN, J. and MORGENSTERN, O. (1980). *Theory of games and economic behavior*, third ed. Princeton University Press, Princeton, N.J.
 - [107] WEBER, P. and ROSENOW, B. (2005). Order book approach to price impact. *Quant. Finance* **5** 357-364.
 - [108] YAN, J.-A. (2002). A numeraire-free and original probability based framework for financial markets. In *Proceedings of the International Congress of Mathematicians, Vol. III (Beijing, 2002)* 861–871. Higher Ed. Press, Beijing.