Optimisation and Utility Functions

Walter Schachermayer

2010 Mathematics Subject Classification: 91B16, 91B24
Keywords and Phrases: Portfolio optimisation, utility functions

The story begins in St. Petersburg in 1738. There Daniel Bernoulli proposed a solution to the "St. Petersburg Paradox" by introducing the notion of a *utility function*.

The problem is formulated in somewhat flowery terms as a game. It was proposed by Nicholas Bernoulli, a cousin of Daniel, in a letter from 1713 to Pierre Raymond de Montmort. Suppose I offer you a random sum of money where the amount is determined from subsequent tosses of a fair coin in the following way. The payoff equals 2^n ducats if the first *heads* appears on the n'th toss. Of course, this event has probability 2^{-n}, so that the expected value of the payoff equals

$$\tfrac{1}{2} \times 2 + \tfrac{1}{4} \times 4 + \ldots + \tfrac{1}{2^n} 2^n + \ldots = \infty. \tag{1}$$

Here is the question: how much would you be willing to pay to me as a *fixed price* for obtaining this kind of lottery ticket?

It is instructive to discuss this question with students in a class and to ask for bids. One rarely gets a bid higher than, say, 10 ducats.

This is remarkably far away from the *expected payoff* of the game which is infinity. Clever students quickly ask a crucial question: are we allowed to play this game *repeatedly*? This would change the situation dramatically! The law of large numbers, which was already well understood in Daniel Bernoulli's times, at least in its weak form, tells you that in the long run the average win per game would indeed increase to infinity. Hence in this case, clever students would be willing to pay quite an elevated fixed price for the game.

But the flavor of the problem is that you are only offered to play the game *once*. How to determine a reasonable *value* of the game?

Daniel Bernoulli proposed *not to consider* the nominal amount of money but rather to transform the money scale onto a different scale, namely the *utility* which a person draws from the money. For a good historic account we refer to [4]. Daniel Bernoulli proposed to take $U(x) := \log(x)$ as a measure of the *utility* of having an amount of x ducats. And he gives good reasons for this choice: think of a person, an "economic agent" in todays economic lingo, who manages to increase her initial wealth $w > 0$ by 10%. Measuring utility by the logarithm then yields that the increase in utility is independent of w, namely $\log(\frac{11w}{10}) - \log(w) = \log(\frac{11}{10})$.

Bernoulli therefore passes from the expected nominal amount (1) of the game to the *expected utility* of the wealth of an agent after receiving the random amount of the game, i.e.,

$$\tfrac{1}{2}\log(w-c+2) + \tfrac{1}{4}\log(w-c+4) + \ldots + \tfrac{1}{2^n}\log(w-c+2^n) + \ldots, \qquad (2)$$

where w denotes the initial wealth of the agent and c the price she has to pay for the game. Of course, this sum now converges. For example, if $w - c = 0$, the sum equals $\log(4)$. This allows for the following interpretation: suppose the initial wealth of the agent equals $w = 4$. Then $c = 4$ would be a reasonable price for the game, as in this case the agent who uses *expected* log-*utility* as a valuation of the payoff, is indifferent between the following two possibilities:

(1) not playing the game in which case the wealth remains at $w = 4$, yielding a *certain* utility of $\log(4)$.
(2) Playing the game and paying $c = 4$ for this opportunity. This yields, by the above calculation, also an *expected* utility of $\log(4)$.

The above method today is known as "utility indifference pricing". We have illustrated it for initial wealth $w = 4$, as the calculations are particularly easy for this special value. But, of course, the same reasoning applies to general values of w. It is immediate to verify that this pricing rule yields a price $c(w)$ in dependence of the initial wealth w which is increasing in w. In economic terms this means that, the richer an agent is, the more she is willing to pay for the above game. This does make sense economically. In any case, the introduction of *utility functions* opened a perspective of dealing with the "St. Petersburg Paradox" in a logically consistent way.

Let us now make a big jump from 18'th century St. Petersburg to Vienna in the 1930's. The young Karl Menger started with a number of even younger mathematical geniuses the "Mathematische Colloquium". Participants were, among others, Kurt Gödel, Olga Taussky, Abraham Wald, Franz Alt. There also came international visitors, e.g., John von Neumann or Georg Nöbeling. In this colloquium a wide range of mathematical problems were tackled. Inspired by an open-minded banker, Karl Schlesinger, the Colloquium also dealt with a basic economic question: How are prices formed in a competitive economy? As a toy model think about a market place where "many" consumers can buy *apples*, *bananas*, and *citruses* from "many" merchants. We assume that the consumers are well informed, that they want to get the best deal for their money, and that there are no transaction costs.

This assumption implies already that the prices π_a, π_b, π_c of these goods have to be equal, for each merchant. Indeed, otherwise merchants offering higher prices than their competitors could not sell their fruits.

For each of the consumers the market prices π_a, π_b, π_c are *given* and, depending on their *preferences* and budgets, they make their buying decisions as functions of (π_a, π_b, π_c). On the other hand, the merchants decide on these prices. For example, if the current prices are such that the apples are immediately sold out, while few people want to buy the bananas, it seems obvious that the price π_a should go up, while π_b should go down. This seems quite convincing if we only have apples and bananas, but if there are more than two goods, it is not so obvious any more how the prices for the apples and the bananas relate to the demand for citruses.

This question was already treated some 50 years earlier by Léon Walras, who was Professor of economics in Lausanne. He modeled the situation by assuming that each agent is endowed with an initial wealth w and a *utility function* U assigning to each combination (x_a, x_b, x_c) of apples, bananas, and citruses a real number $U(x_a, x_b, x_c)$. For given prices (π_a, π_b, π_c), each of the agents optimises her "portfolio" (x_a, x_b, x_c) of apples, bananas, and citruses. In this setting, we call a system of prices (π_a, π_b, π_c) an *equilibrium* if "markets clear", i.e., if for each of the three goods the total demand equals the total supply.

The obvious question is: Is there an equilibrium? Is it unique?

Léon Walras transformed the above collection of optimisation problems, which each of the "many" agents has to solve for her specific endowment and utility function, into a set of equations by setting the relevant partial derivatives zero. And then he simply counted the resulting number of equations and unknowns and noted that they are equal. At this point he concluded – more or less tacitly – that there must be a solution which, of course, should be unique as one can read in his paper "Die Gleichungen des Tausches" from 1875.

But, of course, in the 1930's such a reasoning did not meet the standards of a "Mathematische Colloquium" any more. Abraham Wald noticed that the question of existence of an equilibrium has to be tackled as a fixed point problem and eventually reduced it to an application of Brouwer's fixed point theorem. He gave a talk on this in the Colloquium and the paper was announced to appear in the spring of 1938. However, the paper was lost in the turmoil of the "Anschluss" of Austria, when the Colloquium abruptly ended, and most participants had other worries, namely organising their emigration. It was only after the war that this topic was brought up again with great success, notably by the eminent economists Kenneth Arrow and Gerard Debreu.

Finally, we make one more big jump in time and space, this time to Boston in the late 1960's. The famous economist Paul Samuelson at MIT had become interested in the problem of option pricing. Triggered by a question of Jim Savage, Paul Samuelson had re-discovered the dissertation of Louis Bachelier, entitled "Théorie de la spéculation", which Bachelier had defended in 1900 in Paris. Henri Poincaré was a member of the jury. In his dissertation Bachelier had introduced the concept of a "Brownian motion" (this is today's terminology) as a model for the price process of financial assets. He thus anticipated the work of Albert Einstein (1905) and Marian Smoluchowski (1906) who independently applied this concept in the context of thermodynamics.

Paul Samuelson proposed a slight variant of Bachelier's model, namely putting the Brownian motion W on an exponential scale, i.e.,

$$dS_t = S_t \mu dt + S_t \sigma dW_t, \qquad 0 \leq t \leq T. \tag{3}$$

Here S_t denotes the price of a "stock" (e.g. a share of Google) at time t. The initial value S_0 is known and the above stochastic differential equation models the evolution of the stock price in time. The parameter μ corresponds to the drift of the process, while $\sigma > 0$ is the "volatility" of the stock price, which models the impact of the stochastic influence of the Brownian motion W.

This model is called the "Black-Scholes model" today, as Fisher Black and Myron Scholes managed in 1973 to obtain a pricing formula for *options* on the stock S which is solely based on the "principle of no arbitrage". This result was obtained simultaneously by Robert Merton, a student of Paul Samuelson. The "Black-Scholes formula" earned Myron Scholes and Robert Merton a Nobel prize in Economics in 1997 (Fisher Black unfortunately had passed away already in 1995).

Here we want to focus on a slightly different aspect of Robert Merton's work, namely *dynamic portfolio optimisation*, which he started to investigate in the late sixties [3]. Imagine an investor who has the choice of investing either into a stock which is modeled by (3) above, or into a bond which earns a deterministic fixed interest rate, which we may assume (without loss of generality) to be simply zero. How much of her money should she invest into the stock and how much into the bond? The *dynamic* aspect of the problem is that the investor can – and, in fact, should – rebalance her portfolio *in continuous time*, i.e., at every moment.

To tackle this problem, Merton fixed a utility function $U : \mathbb{R}_+ \to \mathbb{R}$ modeling the *risk aversion* of the investor. A typical choice is the "power utility"

$$U(x) = \tfrac{x^\gamma}{\gamma}, \qquad x > 0, \tag{4}$$

where γ is a parameter in $]-\infty, 1[\, \setminus \{0\}$. Of course, the case $\gamma = 0$ corresponds to the logarithmic utility. One thus may well-define the problem of *maximising the expected utility* of terminal wealth at a fixed time T, where we optimise over all trading strategies. A similar problem can be formulated when you allow for consumption in continuous time.

Here is the beautiful result by Robert Merton. Fixing the model (3) and the utility function (4), the optimal strategy consists of investing a fixed *fraction* m of one's total wealth into the stock (and the remaining funds into the bond). The value m of this fraction can be explicitly calculated from the parameters appearing in (3) and (4).

To visualize things suppose that $m = \frac{1}{2}$, so that the investor always puts half of her money into the stock and the other half into the bond. This implies that the investor sells stocks, when their prices go up, and buys them when they go down. A remarkable feature is that she should do so in continuous time which – in view of wellknown properties of Brownian trajectories – implies that the total volume of her trading is almost surely infinite, during each interval of time!

The method of Merton is dynamic programming. He defines the Hamilton–Jacobi–Bellman value-function corresponding to the above problem. In this setting he manages to explicitly solve the PDE which is satisfied by this value-function.

Of course, this so-called "primal method" is not confined to the special setting analysed by Robert Merton. It can be – and was – extended to many variants and generalisations of the above situation.

There is also a dual approach to this family of problems which was initiated in a different context by J.-M. Bismut [1]. In the Mathematical Finance community this approach is also called the "martingale method". Speaking abstractly, Merton's problem is just a convex optimisation problem over some infinite-dimensional set, namely the set of all "admissible" trading strategies. As is very wellknown, one may associate to each convex optimisation problem a "dual" problem, at least formally. The method consists in introducing (an infinite number of) Lagrange multipliers and to find a saddle point of the resulting Lagrangian function. This leads to an application of the minmax theorem. Eventually one has to optimize the Legendre transform of U over an appropriate "polar" set.

To make this general route mathematically precise, one has to identify appropriate regularity conditions, which make sure that things really work as they should, e.g., existence and uniqueness of the primal and dual optimizer as well as their differential relations. In the present case, there are two aspects of regularity conditions: on the one hand side on the model of the stock price process, e.g., (3), and on the other hand on the choice of the utility function, e.g., (4). In order to develop a better understanding of the nature of the problem, from a mathematical as well as from an economic point of view, it is desirable to identify the *natural* regularity assumptions. Ideally, they should be necessary and sufficient for a good duality theory to hold true.

In [2] this question was answered in the following way. As regards the choice of the model S for the stock price process, virtually nothing has to be assumed, except for its arbitrage freeness, which is very natural in the present context. As regards the utility function U one has to impose the condition of *"reasonable asymptotic elasticity"*,

$$\limsup_{x\to\infty} \frac{xU'(x)}{U(x)} < 1, \tag{5}$$

which is reminiscent of the $\Delta 2$ condition in the theory of Orlicz spaces. The name "asymptotic elasticity" comes from the fact that the derivative $U'(x)$, normalised by $U(x)$ and x as in (5), is called the "elasticity" of U in economics. To get a feeling for the significance of condition (5), note that for a concave, increasing function U the above limit is always less than or equal to 1. In the case of power utility (4) this limit equals $\gamma < 1$. Considering $U(x) = \frac{x}{\log(x)}$, for $x > x_0$, we find an example where the above limit equals 1, i.e., a utility function U which fails to have "reasonable asymptotic elasticity".

It turns out that condition (5) is a *necessary and sufficient* condition for the duality theory to work in a satisfactory way. If it is violated, one can find a stock price process S – in fact a rather simple and regular one – such that the duality theory totally fails. On the other hand, if it holds true, the duality theory, as well as existence and uniqueness of the primal and dual optimiser etc, works out well, even for very general stock price processes S.

There is a lot of further research on its way on related issues of portfolio optimisation. As an example, we mention the consideration of proportional transaction costs (e.g., Tobin tax) in the above problem of choosing an optimal dynamic portfolio. Of course, the most fruitful approach is the interplay between primal and dual methods.

References

[1] J.M. Bismut, Conjugate convex functions in optimal stochastic control. J. Math. Anal. Appl. 44, 384–404 (1973)

[2] D. Kramkov and W. Schachermayer, The condition on the asymptotic elasticity of utility functions and optimal investment in incomplete markets. Annals of Applied Probability 9 3, 904–950 (1999)

[3] R.C. Merton, Optimum consumption and portfolio rules in a continuous-time model. Journal of Economic Theory 3, 373–413 (1971)

[4] Gordan Zitkovic, Utility Theory: Historical Perpectives. Encyclopedia of Quantitative Finance 4, 1860–1862 (2010)

Walter Schachermayer
University of Vienna
Faculty of Mathematics
Nordbergstraße 15
1090 Vienna
Austria
walter.schachermayer@univie.ac.at