20 The Origins of Expected Utility Theory

Yvan Lengwiler*

This short contribution is not about Vinzenz Bronzin or about option pricing. Rather, the topic I would like to address is another important piece of economic theory, namely the theory of expected utility maximization. It is interesting to note just how many thinkers have contributed to it, and at the same time to realize that the earliest statements of the theory were the most powerful ones, and were followed by weaker conceptions. It just took the field of economics a surprisingly long time to grasp its full potential. I believe that the history of this great piece of theory is instructive, because it is an example of a powerful idea that was assimilated only very slowly and in a roundabout fashion.

20.1 Introduction

Expected utility theory consists of two components. The first component is that people use or should use the expected value of the utility of different possible outcomes of their choices as a guide for making decisions. I say "use or should use" because the theory can be interpreted in a positive or a normative fashion. With "expected value" we mean the weighted sum, where the weights are the probabilities of the different possible outcomes. This component, which I discuss in section 2, goes back to the Blaise Pascal's writings of mid-17th century.

The second component is the idea or insight that more of the same creates additional utility only with a decreasing rate. This assumption of decreasing marginal utility plays a very central role in economics in general, but as we will see, is actually older than the marginalist school with which we would typically associate this idea. I discuss some of the contributions of the marginalist school in section 3.

In section 4, I talk about the additional insight that is possible by combining both components. It is this combination that gives rise to the concept of risk aversion and implies the demand for diversification and insurance. When we use the term "expected utility theory", we typically mean the combination of these two components.

Section 5 is a digression into the problems connected with unbounded utility functions. These problems relate to Pascal's original writings, but may also be relevant for the way we use expected utility theory today.

* Universität Basel, Switzerland. *yvan.lengwiler@unibas.ch*. I thank Heinz Zimmermann and Ralph Hertwig for useful remarks.

20.2 Pascal and God

The first component is old, very old. In fact, it is as old as probability theory itself. In the mid-17th century, Blaise Pascal (1670) presented a peculiar argument explaining why believing in god is rational, and not believing is not rational. This argument, known as “Pascal’s wager”, is an arbitrage or hedging argument. I do not know the psychological or social circumstances that Pascal was subject to when proposing this argument, but to me it seems quite far-fetched and artificial, especially since it can easily be invalidated, even within the framework of expected utility maximization that Pascal proposes. The wager works as follows. Consider a binomial world: either god exists or god does not exist. You have to decide on which of these two cases you bet by choosing whether to be religious or not. Pascal proposes the following payoffs:

	god exists	god does not exist
living as if god exists	$-C+\infty$	$-C$
living as if god does not exist	$U-\infty$	U

U is the utility provided by an earthly life unconstrained by religion. C is the disutility from living a god-abiding life.[1] Pascal argues that both, C and U, are finite, whereas the stakes are infinite in the case that god exists, simply because afterlife is infinitely longer than earthly life. If god exists, believers will spend an eternal afterlife in heaven, collecting an infinite amount of utility; non-believers will receive infinite disutility by spending eternity in hell. Obviously, if the prior probability of god existing is strictly positive (even if arbitrarily small), choosing to be religious is the best reply. So, people should choose to be religious simply in order to hedge the risk of eternal damnation and bet on the possibility of eternal bliss.

Pascal’s wager has generated a lively debate in philosophy, maybe in part because there are so many obvious arguments against it. One obvious, and in my view devastating objection, is the many gods objection.[2] It runs as follows: maybe there is a god, but it is unclear what type of god it is. Several types are advertised on earth right now: there is the christian faction, the muslim faction, the hindu faction, all of them with various sub-types, and also several smaller enterprizes. How would a god, type-X, treat an atheist compared with a believer of a god, type-Y? Of course, one could try to worship all the proposed gods, but

[1] Actually, the sign of C is not important. Whether living a religious life provides positive or negative utility is immaterial because the absolute level of utility has no meaning. The assumption is simply that $-C<U$. Pascal argues that despite this assumption it is still rational to be religious.

[2] Diderot (1875–1877) is generally acclaimed to be the first to make this objection by noting that “An Imam could reason just as well this way”.

would portfolio diversification work in this case? Maybe god demands exclusive devotion?

More generally, if nothing is known about god, it is essentially random what the right thing to do is. Maybe god dislikes obedient believers in general but prefers critical minds, and thus treats atheists the best? Or maybe he just likes people with blue hair. So we should all color our hair or wear a wig?

Another argument, which I have not read before, but which comes naturally to an economist, is discounting. Let a stay in heaven yield a flow of g utils, and a stay in hell yields a flow of $-h$ utils. Similarly, a stay on earth without religious constraints yields a flow of u utils, and with constraints it yields a flow of $-c$ utils. The person discounts future utils with a rate of r. Let T be the remaining length of the person's earthly life (assumed, for simplicity, not to be stochastic). Then Pascal's payoff matrix presents itself as follows,

	god exists	god does not exist
living as if god exists	$-C+G$	$-C$
living as if god does not exist	$U-H$	U

where

$$C := \int_0^T c\exp(-rt)dt = \frac{c}{r}(1-\exp(-rT)), \qquad G := \int_T^\infty g\exp(-rt)dt = \frac{g}{r}\exp(-rT),$$
$$U := \int_0^T u\exp(-rt)dt = \frac{u}{r}(1-\exp(-rT)), \qquad H := \int_T^\infty h\exp(-rt)dt = \frac{h}{r}\exp(-rT),$$

are the present values of the different kinds of lives and afterlives. Let p be the probability that god exists. After a few manipulations we conclude that being religious is the best reply if and only if

$$p > p^* := \frac{u+c}{g+h}(\exp(rT)-1).$$

Without discounting $(r=0)$ we are back at Pascal's wager: any strictly positive probability of god's existence $(p>0)$ rationalizes to be religious, because, in that case, $p^*=0$. But with discounting $(r>0)$, this is no longer true, because now $p^*>0$. This means that god has to be *sufficiently probable* in order for an individual to rationally choose to be religious. The reason why this happens is that, despite the fact that afterlife is by assumption eternal, *the slightest amount* *of discounting makes the present value of afterlife finite*.[3]

Actually, it is somewhat interesting to study how the threshold probability, p^*, changes with the remaining length of life. According to the above formula, young people (large T) would need better evidence for the existence of god in order to be religious than old people (small T), because p^* is decreasing in T. As death approaches $(T \to 0)$, the required probability vanishes $(p^* \to 0)$, and so eventually it becomes rational for everyone to be a theist. The reason for this effect is that the relative weight of life before death compared to potential afterlife eventually vanishes as life comes to an end.

Now, all of this is, I think, quite ridiculous. The wager is interesting for us not as an argument for religion, but because, to my knowledge, Pascal, who is one of the founding fathers of probability theory, is the first scholar to explicitly propose the expected utility of possible outcomes of a given choice as a decision rule. Thus, we conclude that this first component of expected utility theory is as old as probability theory itself.

20.3 Decreasing Marginal Utility

The second component – the assumption that marginal utility is a decreasing function – is the hallmark of the marginalist revolution that took place in 19th century economics, but which also bears fruit in other areas.

Fechner (1860), following the work of Weber (1851), developed a research program, which he called psycho-physics, that tried to relate stimulus to sensation in a quantitative fashion. By how much does the sensation of light or loudness of touch change as a result of brighter light, louder sound, or more pressure? He concluded from his experiments that Bernoulli's logarithmic specification, to which he refers (and which we discuss in the next section) was a generally valid principle: let x be stimulus and let u be sensation, then the Weber-Fechner law says that the *just noticeable difference* ("*eben merkliche Unterschied*"), that is, the smallest increase in stimulus, dx, that leads to a noticeable difference of sensation, du, is proportional to the level of the stimulus. Formally, $kdx = xdu$, or $u(x) = k \ln(x)$. A hundred years later, Stevens (1961) challenged the Weber-Fechner law and proposed, instead, a power specification, $u(x) = k(x - x_0)^b$.[4] To an economist, it is difficult to understand how one could make a big fuss about these specifications, since both specifications feature constant relative risk aversion, and economists are not interested in absolute utility scales. This is, of course, very different for psycho-

[3] Pascal argues as follows: one bets one certain life against one uncertain afterlife. But because afterlife (if it exists) is eternal, the payoff in afterlife swamps all other payoffs (Pascal 1670, § 233). Discounting invalidates this conclusion.

[4] This, in turn, has not passed unchallenged either, see Florentine and Epstein (2006).

physicists, who are looking for a *quantitative* relation.

In the economic field, Dupuit (1844, 1853) was the first to derive from the general concept of decreasing marginal utility the idea of a decreasing demand function. By clearly distinguishing the utility generated by the last used unit from the total utility he also developed the concept of the consumer's rent. Without reference to Dupuit, Gossen (1854) deduced from the idea of decreasing marginal utility the conclusion that an individual would optimally allocate his income in such a way that the marginal contribution of money to utility would be equal for all possible uses of money. In other words, if p_i is the price of good i, and du_i is the marginal utility of good i for a given person, then du_i / p_i should be the same for all commodities i for a given person. This is Gossen's most significant "second law" and is the same as the first order condition of utility maximization subject to a budget constraint, assuming price-taking behavior. Yet, Gossen's work was without any consequence because no one read his book. This work may have passed by unnoticed due to poor marketing. His position as a retired public servant was probably not helpful either in promoting his notability amongst academics. Jevons reports that none of the academics of the time who thought they were proficient in German economics had heard of Gossen (see § 28 of the preface to the second edition of Jevons 1871). It was finally Jevons who discovered Gossens' book in 1878. He acknowledged that Gossen had preceded him, but it was Jevons' theory of exchange that influenced the discussion at the time.

Significant progress was achieved by Walras (1874) and by Edgeworth (1881). Walras analyzed a complete system of multiple markets, assuming price-taking behavior by each individual person. From the aggregation of individuals' budget constraints he derived the famous Walras' Law, stating that if $n-1$ markets are in equilibrium, then the *n*-th market is necessarily also in equilibrium. This was, of course, the foundation of general equilibrium theory. Edgeworth, on the other hand, analyzed multiple bilateral exchange. He realized that many allocations would be possible in equilibrium (the contract curve), but conjectured that as competition intensifies, the set of equilibria should shrink. The existence of a Walras equilibrium was later proved formally by Arrow and Debreu (1954), and the validity of Edgeworth's core convergence conjecture was established by Debreu and Scarf (1963).

All these authors shared a common device: they used abstract, unspecified utility functions.[5] Consequently, the resulting equilibria possessed only rudimentary structure. This lack of structure finally led the field into a dead end. All that economists were able to show was that an abstract economy had an abstract equilibrium, and that the equilibrium allocation would satisfy certain properties (such as efficiency). But, except for simple toy models, it was

[5] Jevons, however, fully acknowledged the need to be concrete: "We cannot really tell the effect of any change in trade or manufacture until we can with some approach to truth express the laws of the variation of utility numerically" (Jevons 1871, Chapter IV, § 105).

impossible in general to construct an equilibrium and see what it looked like. The Sonnenschein-Mantel-Debreu theorem (Sonnenschein 1973, Mantel 1974, Debreu 1974) can be seen as the tombstone of abstract general equilibrium theory. It says that general equilibrium theory is compatible with everything and therefore is not falsifiable. Consequently, it is not a scientific theory in the sense of Popper (1966). Scientific orthodoxy requires more structure and more concrete assumptions, which are, ideally, empirically supported.

20.4 Cramer and Bernoulli Knew it All

The combination of the two components discussed above produces the very powerful theory of expected utility as we know and use it today. It is surprising to realize that all of this was already known in the 18th century, long before the marginalist revolution in economics. In discussing the St. Petersburg paradox, Gabriel Cramer, in a letter written in 1728, proposed to evaluate gambles by considering the expected utility of the money gained, where the utility would be measured as the square root of the payout. Ten years later, Daniel Bernoulli proposed to use the logarithm. It is quite striking to read the few lines in which Bernoulli lays out the ideas of expected utility theory (I quote from the English translation):

> "If the utility of each possible profit expectation is multiplied by the numbers of ways it can occur, and we then divide the sum of these products by the total number of cases, a mean utility (moral expectation) will be obtained, and the profit which corresponds to this utility will equal the value of the risk in question" (Bernoulli 1954, § 4).
> "However, it hardly seems plausible to make any precise generalizations since the utility of an item may change with circumstances. Thus, though *a poor man generally obtains more utility than does a rich man from an equal gain*, it is nevertheless conceivable, for example, that a rich prisoner who possesses two thousand ducats but needs two thousand ducats more to purchase his freedom, will place higher value on a gain of two thousand ducats than does another man who has less money than he. Though innumerable examples of this kind may be constructed, they represent exceedingly rare exceptions. We shall, therefore, do better to consider what usually happens, and in order to perceive the problem more correctly we shall assume that there is an imperceptibly small growth in the individual's wealth which proceeds continuously by infinitesimal increments. Now it is highly probable that *any increase in wealth, no matter how insignificant, will always result in an increase in utility which is inversely proportionate to the quantity of goods already possessed*" (Bernoulli 1954, § 5, first emphasis added).

In the first quote, Bernoulli proposes Pascal's principle.[6] In the second quote, he first proposes the general principle of decreasing marginal utility, and then also proposes a specific functional form, namely $du = x^{-1}dx$, or in other words, $u(x) = \ln(x)$.

Bernoulli then goes on to explain that what matters is not the gain in the particular gamble, but the total wealth of the individual, where zero wealth is defined as the subsistence level:

> "[...] nobody can be said to possess nothing at all in this sense unless he starves to death. For the great majority the most valuable portion of their possessions so defined will consist in their productive capacity, this term being taken to include even the beggar's talent: a man who is able to acquire ten ducats yearly by begging will scarcely be willing to accept a sum of fifty ducats on condition that he henceforth refrain from begging or otherwise trying to earn money. For he would have to live on this amount, and after he had spent it his existence must also come to an end. I doubt whether even those who do not possess a farthing and are burdened with financial obligations would be willing to free themselves of their debts or even to accept a still greater gift on such a condition. But if the beggar were to refuse such a contract unless immediately paid no less than one hundred ducats and the man pressed by creditors similarly demanded one thousand ducats, we might say that the former is possessed of wealth worth one hundred, and the latter of one thousand ducats, though in common parlance the former owns nothing and the latter less than nothing" (Bernoulli 1954, § 5).

Bernoulli lays out a very modern concept of wealth here: wealth is not the stock of assets a person owns, but rather the ability to generate an income stream. More precisely, it is the amount that the individual would be willing to trade in exchange for his ability to generate future income. This is not exactly the net present value of lifetime income, because the individual's preferences are used to value risky cash flows at different points in time instead of market prices, but it does come close to it.

It is surprising that 110 years before Gossen wrote his little-read book, 120 years before Fechner formulated his law, and 130 years before Jevons formulated his theory of exchange, two great mathematicians had already formulated a superior decision-theory. I say 'superior' because Cramer's and Bernoulli's formulation contained both components – expected utility and decreasing marginal utility –, not the second component alone. This construction is capable

[6] Interestingly, Bernoulli does not refer to Pascal's wager, even though the idea is the same. I do not know whether Bernoulli was not aware of Pascal's "Pensées" (which seems hard to imagine), or whether it was just not as usual as it is today to explicitly acknowledge previous thinkers.

of dealing with situations involving risk. In fact, their theory was designed for this case. Gossen's, Fechner's and Jevons' theories are not capable of addressing the exchange of risky bets.

Bernoulli's work was known to Fechner and Jevons – they both refer to him –, but these scholars did not realize that Bernoulli's formulation was superior. Jevons writes:

> "The variation of utility has not been overlooked by mathematicians, who had observed, as long ago as the early part of last century - before, in fact, there was any science of Political Economy at all - that the theory of probabilities could not be applied to commerce or gaming without taking notice of the very different utility of the same sum of money to different persons. [...] Daniel Bernoulli, accordingly, distinguished in any question of probabilities between the *moral expectation* and the *mathematical expectation*, the latter being the simple chance of obtaining some possession, the former the chance as measured by its utility to the person. Having no means of ascertaining numerically the variation of utility, Bernoulli had to make assumptions of an arbitrary kind, and was then able to obtain reasonable answers to many important questions. It is almost self-evident that the utility of money decreases as a person's total wealth increases; if this be granted, it follows at once that gaming is, in the longrun, a sure way to lose utility; that every person should, when possible, divide risks, that is, prefer two equal chances of £50 to one similar chance of £100; and the advantage of insurance of all kinds is proved from the same theory" (Jevons 1871, chapter 4, § 125).

It is evident from this quote that Jevons perfectly appreciated some fundamental implications of expected utility theory, such as the soundness of diversification and the demand for insurance. He failed, however, to fully spell out these implications in any further detail. Had he combined expected utility theory with his own theory of exchange, he would have reached a theory of the exchange of risky gambles, and he might have become the founder of what we call finance today.

20.5 Unbounded Utility

Menger (1934) pointed out that the utility function must be bounded, for otherwise it may fail to yield finite expected utility with some distributions, and thus there may be no maximizer. Menger also pointed out that, for the same reason, the logarithmic or the square root functions do not really resolve the St. Petersburg paradox. Because these utility functions are unbounded, one can always find a distribution of the payoff that yields infinite expected utility. In order to prevent this, given arbitrary distributions of the payoff, the utility function itself has to be bounded. Pascal's wager suffers, of course, precisely from the fact that the construction yields infinite expected utility, and thus can lead to unconvincing conclusions.

Arrow (1965) concluded that relative risk aversion must approach a value smaller than one as wealth approaches zero, and must approach a value greater than one as wealth grows indefinitely, if the utility function is bounded. Thus, relative risk aversion must be globally increasing, although it can be locally decreasing: "Thus, broadly speaking, the relative risk aversion must hover around 1, being, if anything, somewhat less for low wealths and somewhat higher for high wealths" (Arrow 1965, p. 37). Essentially, any bounded utility function must hover around the logarithmic function, although the logarithmic function itself is not a valid utility function because it is unbounded. Arrow is again very clear:

> "[...] if, for simplicity, we wish to assume a constant relative risk aversion, then the appropriate value is one. As can easily be seen, this implies that the utility of wealth equals its logarithm, a relation already suggested by Bernoulli" (Arrow 1965, p. 37).

Ten years before Menger noted the need for a bounded utility function, Charles Jordan also argued for a bounded utility function on different grounds. He explicitly refers to the psycho-physics literature and then argues:

> "[...] while accounting for the threshold of sensations, it (Bernoulli's specification) asserts that there is no upper limit for them. The sensations increase, it states, indefinitely with the stimuli. But we know that this is not true [...]" (Jordan 1924, § 12).

He then proposes an alternative specification,

$$u(x) = \lambda \left[1 - \frac{x_0 + a}{x + a} \right].$$

$u(x_0) = 0$, and Jordan interprets x_0 as the "threshold of [...] cautiousness". x_0 and λ can be understood just as normalizations, with λ being a scaling factor and x_0 determining the absolute level of utility. Unlike psycho-physicists, economists have no interest in absolute utility levels or scales. Moreover, absolute and relative risk aversion are unaffected by these two coefficients, so for economic applications we may just as well set $\lambda = 1$ and $x_0 = 0$.

Jordan's specification is interesting because the range of this utility function is bounded (it is the unit interval). Relative risk aversion is also bounded and is monotonically increasing from 0 to 2.[7] Relative risk aversion at the point $x = a$ is 1. Thus, Jordan's utility function "hovers around the logarithmic function" in the sense of Arrow, because it features bounded relative risk aversion around unity. To my knowledge, this utility function has not been used by economists, despite its interesting properties.

The axiomatization of the theory that was provided by Von Neumann and Morgenstern (1944) as an appendix to their work on game theory does not allow for outcomes that are associated with infinite levels of utility. Thus, the theory is not strictly compatible with Cramer or Bernoulli, or with Pascal. But that is exactly why it is immune to the absurdity of Pascal's wager. Some generalizations are possible, if we restrict the distributions of the random variables. Ryan (1974) and Arrow (1974) work out cases where utility functions that are unbounded above may still be admissible. However, they show that one still needs either a lower bound on the utility or on the first derivative of utility, so either $u(0)$ or $u'(0)$ must be finite. Both conditions are not met by the popular constant relative risk aversion specification that we routinely use in economics and finance. Strictly speaking, these specifications are not covered by the theory.

20.6 Back to the Roots

Today, we very often use the constant relative risk aversion specification, i.e. the power or the log function. It is not without irony that the field has found that the original specification that was proposed by Cramer (power function) and by Bernoulli (log) are actually quite useful. These are also the specifications that have shaped psycho-physics, with the Weber-Fechner law being Bernoulli's specification, and Steven's formula being a generalization of Cramer's proposal. Jevons had called Bernoulli's specification an "assumption of an arbitrary kind" (see quotation above), but even if this choice was arbitrary in the sense of not being founded upon experiments, it still demonstrates great intuition or insight.

The Sonnenschein-Mantel-Debreu result was a wake-up call. Economics is now more interested in concrete models. In that sense, economics has moved closer to psycho-physics. This is also demonstrated by the fact that experiments have become a widely used method in economics in the more recent past. In this sense, the program to use experiments in economics could be labelled "econophysics", though the term seems to be taken already (Mantegna and Stanley 1999).

The move away from abstract theories that have too little structure to yield interesting (falsifiable) results, and towards more concrete models is also a move back to the roots, so to speak, because when economists use expected utility

[7] In fact, relative risk aversion is proportional to the utility level, $-xu''(x)/u'(x) = 2u(x)$.

theory today, they are, it seems to me, closer to Cramer and Bernoulli than to Gossen and Jevons.

References

Arrow K J (1965) Aspects of the theory of risk-bearing. Yrjö Jahnsson Foundation, Helsinki

Arrow K J (1974) The use of unbounded utility functions in expected-utility maximization: response. Quarterly Journal of Economics 88, pp. 136–138

Arrow K J, Debreu G (1954) Existence of an equilibrium for a competitive economy. Econometrica 22, pp. 265–290

Bernoulli D (1954) Exposition of a new theory on the measurement of risk. Econometrica 22, pp. 23–36 (Originally published in 1738 as 'Specimen theoriae novae de mensura sortis', St. Petersburg)

Debreu G (1974) Excess demand functions. Journal of Mathematical Economics 1, pp. 15–21

Debreu G, Scarf H (1963) A limit theorem on the core of an economy. International Economic Review 4, pp. 235–246

Diderot D (1875-1877) Pensées philosophiques, Vol. 1 (LIX)

Dupuit J (1844) De la mesure de l'utilité des travaux publics. Annales des Ponts et Chaussées 8

Dupuit J (1853) De l'utilité et de la mesure. Journal des Économistes 12, pp. 1–27

Edgeworth F Y (1881) Mathematical psychics. Kegan Paul & Co., London

Fechner G T (1860) Elemente der Psychophysik. Breitenkopf und Härtel, Leipzig

Florentine M, Epstein M (2006) To honor Stevens and repeal his law (for the auditory system). Invited talk, International Society for Psychophysics

Gossen H H (1854) Entwickelung der Gesetze des menschlichen Verkehrs, und der daraus fließenden Regeln für menschliches Handeln. F. Vieweg, Braunschweig

Jevons W S (1871) The theory of political economy, 1st edn. Macmillan and Co. (2nd edn published 1879)

Jordan C (1924) On Daniel Bernoulli's 'moral expectation' and on a new conception of expectation. The American Mathematical Monthly 31, pp. 183–190

Mantegna R N, Stanley H E (1999) An introduction to econophysics. Cambridge University Press, Cambridge (UK)

Mantel R R (1974) On the characterization of aggregate excess demand. Journal of Economic Theory 7, pp. 348–353

Menger K (1934) Das Unsicherheitsmoment in der Wertlehre – Betrachtungen im Anschluß an das sogenannte Petersburger Spiel. Zeitschrift für Nationalökonomie 5, pp. 459–485

Pascal B (1670) Pensées. Republished several times, for instance 1972 in French by Le Livre de Poche, Paris and 1995 in English by Penguin Classics, London

Popper K R (1966) Logik der Forschung. Mohr Siebeck, Tübingen

Ryan T M (1974) The use of unbounded utility functions in expected-utility maximization: comment. Quarterly Journal of Economics 88, pp. 133–135

Sonnenschein H (1973) Do Walras' identity and continuity characterize the class of community excess demand functions? Journal of Economic Theory 6, pp. 345–354

Stevens S S (1961) To honor fechner and repeal his law. Science 133, pp. 80–86

Von Neumann J, Morgenstern O (1944) Theory of games and economic behavior. Princeton University Press, Princeton

Walras L (1874) Éléments d'économie politique pure, ou théorie de la richesse sociale. Corbaz, Lausanne

Weber E H (1851) Die Lehre vom Tastsinne und Gemeingefühle auf Versuche gegründet. F. Vieweg, Braunschweig