Okasha, S. (2011). Optimal Choice in the Face of Risk: Decision Theory meets Evolution. *Philosophy of Science*, 78(1), 83 - 104.
<https://doi.org/10.1086/658115>

Early version, also known as pre-print

Link to published version (if available):
[10.1086/658115](https://doi.org/10.1086/658115)

[Link to publication record in Explore Bristol Research](#)
PDF-document

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Optimal Choice in the Face of Risk: Decision Theory meets Evolution

1. Introduction

The problem of how a rational agent should choose between risky options, or lotteries, is a famous one. The orthodox answer to this problem, of course, is given by expected utility (EU) theory, first made explicit by von Neumann and Morgenstern (1944). EU theory teaches us that so long as an agent's preferences over lotteries obey certain fairly intuitive axioms, then the agent behaves as if she is maximising the expected value of a utility function. The shape of the utility function then reflects the agent's attitude towards risk. EU theorists generally assume that agents have concave utility functions for wealth, as it is an established empirical fact that humans generally have risk-averse preferences, i.e. prefer \$x for certain to any lottery with expected monetary value of \$x.

The crucial axiom in EU theory is the famous independence axiom¹, the subject of much controversy. The axiom was originally attacked by Allais (1952), who regarded it as neither empirically plausible nor normatively compelling; subsequent experiments have found that the axiom is indeed systematically violated. This and other anomalies have prompted the development of various alternatives to EU, generically known as 'non-expected utility theories', which relax or replace the independence axiom, thus leading to maximisation of some quantity other than expected utility. These theories have a certain amount of empirical support, though none has achieved widespread acceptance.

Interestingly, the problem of optimal choice in the face of risk also arises in evolutionary biology.² A typical problem in this area is as follows. An animal can either forage in a resource-rich area where there is a high risk of predation, or in a resource-poor area where there is a lower risk. More resources mean greater survival, so more offspring, so higher Darwinian fitness. Which foraging strategy will be favoured by natural selection? Theoretical work shows that in many circumstances, strategies that reduce variability in resource acquisition will be selectively advantageous; such behavioural strategies are often called 'risk averse'. Empirical

¹ This says that if an agent prefers option a to b, then for all options c and probabilities p, the agent should prefer the lottery (a, c; p, 1-p), to the lottery (b, c; p, 1-p). (The lottery (a, c; p, 1-p) should be read as 'receive a with probability p, and c with probability 1-p').

² The biological literature often uses the term 'uncertainty' in lieu of 'risk', as for example in the title of Frank and Slatkin (1990). This does no harm, since the economist's distinction between risk (known chances) and uncertainty (unknown chances) has no obvious application to non-human animals.

work has confirmed the prediction that animals should often behave in a risk-averse way (cf. Seger and Brockmann 1987).

Surprisingly, there have been relatively few attempts to explicitly link the economic and the evolutionary literatures on risky choice, which have largely proceeded in parallel. (Notable exceptions include Robson 1996, Stearns 2000 and Orr 2007.) This is surprising because in other areas such as game theory, there *has* been significant cross-fertilisation between the economic and evolutionary discussions. The current paper takes a preliminary step towards redressing this situation, by exploring thematic and formal connections between rational choice and evolutionary optimality in risky situations.

My strategy will be to exploit an analogy between utility in rational choice and fitness in Darwinian evolution. (The former is the quantity that the rational agents try to maximise, the latter that natural selection tries to maximise.) This analogy has been noticed before, and has led many authors to posit a link between evolution and EU maximisation (Cooper 2001, Stearns 2000, Orr 2007, Gintis 2009). However I argue that the correct analysis of how evolution works in risky situations in fact suggests a link with *non*-EU theory. To prepare the ground for this argument, I turn first to a neglected though philosophically rich facet of the debate between EU and non-EU theory.

2. EU theory, Allais' critique, and the concept of risk aversion

On orthodox EU theory, an agent's attitude towards risk is fully captured by the shape of her utility function, as noted above. In particular, risk aversion is reflected by a concave utility function for wealth, i.e. diminishing marginal utility (Figure 1). The reason is obvious: iff the agent has risk-averse preferences, i.e. always prefers \$x to any lottery with expected value of \$x, and she maximises expected utility, then her utility function must be concave (by Jensen's inequality). So risk aversion and diminishing marginal utility come to exactly the same thing, in EU theory.

Figure 1: Concave Utility Function

Though the orthodox EU treatment of risk-aversion is familiar, it is not without its critics. The suspicion has often been voiced, originally by Allais (1952), that the orthodox treatment fails to recognise that there are two different reasons why an agent might exhibit risk-averse preferences. The first is their attitude towards wealth, the second their attitude towards risk. Intuitively these are different things, and only the latter counts as ‘real’ risk-aversion. To illustrate, consider an agent who prefers \$5 for sure over the gamble ‘\$10 if a fair coin lands heads, nothing otherwise.’ Such a preference *might* reflect the fact that \$10 brings the agent less than twice the utility of \$5; but it might not. Even if the agent’s utility function is linear in money, she might prefer the certainty of \$5 to the gamble if she wishes to avoid risk. But orthodox EU theory appears only to countenance the first explanation, given that it *defines* risk-aversion in terms of diminishing marginal utility. Or so the objection goes.

An amusing version of this objection was presented by Hansson (1988), who imagined a conversation between a professional gambler and an EU theorist. The gambler is offered a choice between one copy of a book for certain, and either three copies of the book or nothing on the flip of a fair coin. He chooses the first option. The utility theorist concludes that the gambler is risk-averse, for he has declined a gamble with expected value of 1.5 books. When presented with this reasoning, however, the gambler strongly denies being risk-averse – after all, he is a professional

gambler! Rather, he simply has no use for more than one copy of the same book, so sees no point in incurring the risk of getting nothing in return for a chance at winning three copies. This is meant to show that diminishing marginal utility and risk aversion are not the same thing.

Orthodox EU theorists regard this argument as confused, stemming from a mistaken conception of what a utility function is. On their view, an agent's utility function is not supposed to *explain* their preferences, but rather just to *represent* them. So they reject the idea that concave utility is one possible explanation of risk-averse preferences, to be contrasted with other explanations. There simply is no conceptual gap between concave utility and risk averse preferences.

This 'representationalist' viewpoint goes hand-in-hand with the idea that cardinal utility in a 'riskless' context makes no sense – a view explicitly endorsed by the architects of EU theory. An agent's riskless utility function for wealth would tell us how much utility the agent gets from the *definite receipt* of a given amount of money. EU theorists argue that the only justification for cardinal utility comes from the representation theorem of von Neumann and Morgenstern (1944), which considers an agent's preferences over risky gambles. In a purely riskless context, an agent's preferences over definite monetary sums could not be used to construct cardinal utility a la von-Neumann/Morgenstern. Thus Savage (1954) cautioned against confusing von-Neumann/Morgenstern utility with "the new almost obsolete notion of utility in riskless circumstances" (p. 93); while Arrow (1951) called riskless cardinal utility "a meaningless concept" (p. 425).

Maurice Allais (1952), an early critic of EU theory, explicitly defended the existence of riskless cardinal utility. He insisted that utility was psychological real, thus rejecting the view that an agent's utility function is a mere representation of her preferences. Today, Allais is best-remembered for his discovery of the 'Allais paradox', but his broader methodological critique of EU theory went much deeper. In particular, Allais' insistence on the distinction between risk-aversion and diminishing marginal utility, and his defence of the psychological reality of cardinal utility, raise a challenge for EU theory that is independent of the question of whether people in fact obey the EU axioms or not.

Allais argued that a risk-averse agent will not let his choice among gambles be decided by the expected utility criterion; rather, he will take account of the entire distribution of cardinal utilities, not just the expectation. In particular, the variance of

the utilities will be important to the agent – for being averse to risk, he will want to reduce the variation in utility that he receives. Crucial to this argument is Allais’ assumption that riskless cardinal utility makes sense. On the EU view, which rejects this assumption, the idea that an agent’s choice among gambles might be influenced by the variance of the possible utilities makes no sense – the utility function is constructed precisely so that the agent cares only about its expectation. For this reason, most defenders of EU theory rejected Allais’ position as a confusion.³

However, given his starting assumptions, Allais’ treatment of risk aversion is actually very natural.⁴ Suppose that an agent with existing wealth of \$a is faced with a choice between \$5 for sure (option A), and either \$9 or \$1 on the flip of a fair coin (option B). On Allais’ view, the agent will first covert the monetary sums into utilities, by application of her (riskless) utility function $u(x)$; therefore, option A yields $u(a + 5)$ for sure, while option B yields $u(a + 9)$ or $u(a + 1)$ with probability $\frac{1}{2}$ each. The agent’s risk attitude will then enter the picture. The agent *may* choose the option with the highest expected utility, i.e. $\max [u(a + 5), \frac{1}{2} u(a + 1) + \frac{1}{2} u(a + 9)]$, but only if she is risk-neutral. She may also attend to the variance in utility attaching to each option. Since option A has zero variance while option B has a positive variance, this will count against option B if the agent is risk-averse.⁵ So on Allais’ picture, there are two quite different reasons why an agent might prefer option A to B – diminishing marginal utility of money, reflected in the concavity of $u(x)$, and attitude towards risk, reflected in the agent’s attention to the variance as well as the expectation of $u(x)$.

What has all this got to do with evolution? In section 4, we shall see that Allais’ distinction between diminishing marginal utility and ‘real’ risk aversion is mirrored, remarkably, in evolutionary theory. Firstly, we turn briefly to non-expected utility theory.

3. Non-expected Utility Theory

³ See the example the papers by di Finetti and Morgenstern in Allais and Hagen (eds.) (1979).

⁴ Essentially, this is because by ‘utility function’ Allais did not mean a von-Neumann-Morgenstern utility function, as his critics assumed he did.

⁵ Allais (1952) did not offer precise account of how exactly the risk-averse agent would trade off variance in utility (or higher moments) against the expectation, a point which his critics were quick to pick up on. See the papers in Allais and Hagen (eds.) (1979)

Non-expected utility theory, developed in works by Kahneman and Tversky (1979), Quiggin (1982), Machina (1982), Yaari (1987) and others, was motivated primarily by the descriptive failure of EU theory – in particular, the fact that experimental subjects systematically violate the independence axiom. These theorists thus sought better descriptive models of how people actually make risky choices. But a secondary motivation behind non-EU theory was conceptual: dissatisfaction with the orthodox EU treatment of risk attitude, in particular the equation of risk aversion and diminishing marginal utility. Many non-EU theorists regard this equation as erroneous, following Allais.⁶

To see how non-EU theory handles risk attitude, it helps to explicitly contrast the functional forms of the maximands in EU and non-EU theory. Consider lotteries of the form $(\$x_1, \dots, \$x_n; p_1, \dots, p_n)$, i.e. ‘get a prize of $\$x_1$ with probability p_1 , of $\$x_2$ with probability p_2 ...etc.’. In EU theory, an agent evaluates the lotteries according to the criterion $\sum_1^n u(x_i)p_i$, where $u(x_i)$ is the utility of prize $\$x_i$; the lottery with the highest value of this expression is chosen. In non-EU theory, agents choose lotteries according to a different evaluation criterion. The precise criterion (maximand) differs among different versions of non-EU theory, but can often be expressed generically as $\sum_1^n u(x_i)w_i$, where w_i is the *decision weight* of prize i , with $\sum_1^n w_i = 1$. The decision weight of a prize need not equal the probability of getting it. This means that agents are maximising the weighted average of the utilities of the prizes, where the weights may diverge from the true probabilities.⁷

Different versions of non-EU theory specify the decision weights in the above expression differently. Generally, the decision weight of a prize depends on its probability, but also on other factors (e.g. the prize’s value, the cumulative probability of getting a prize at least as good etc.) Many non-EU theories have been formally axiomatized, and they enjoy a certain amount of empirical support. For example, the systematic violations of the independence axiom can be accounted for quite well by assuming that agents are maximising some sort of non-expected utility. Other

⁶ This point comes across clearly in Yaari (1987), Quiggin (1993) and Wakker (1994).

⁷ See Machina (2008) for a good survey of a number of different non-EU theories. Note that some non-EU theories specify evaluation criteria (maximands) that do not have the generic functional form discussed in the text.

empirical features of decision-making in risky situations, such as the general tendency towards risk-aversion, can also be captured by non-EU theory.⁸

Importantly, given the generic form of the non-EU maximand, a distinction between risk attitude and curvature of the utility function immediately arises. To see this, suppose that an agent's utility function $u(x)$ is linear in money. In EU theory, it follows that the agent is risk-neutral, i.e. indifferent between $\$x$ and any gamble with expected monetary value of $\$x$. In non-EU theory this does not follow, for the decision weights need not equal the probabilities. Depending on what the weighting function looks like, the agent may be risk-averse, risk-neutral, or risk-loving; nothing can be deduced about their risk attitude from the fact that their utility function is linear in money, nor vice-versa.

This means that in non-EU theory, there are two distinct ways in which risk-averse behaviour (i.e. preferences) can arise. One is diminishing marginal utility of money, the other is so-called 'probabilistic risk aversion' (Wakker 1994), which occurs when the weighting function is such as to generate aversion to risk independently of the shape of the utility function. Therefore, the problematic equation of risk aversion with diminishing marginal utility, which Allais had objected to, is avoided. The intuitive idea that an agent's attitude towards risk is one thing, their attitude to wealth another, which EU theory cannot accommodate, finds a natural home in non-EU theory.

From a formal point of view, non-EU theory arises by replacing the independence axiom of EU theory with a weaker axiom, thus generating a maximand that is not the expectation of the utility function. So EU and non-EU theories have a similar formal structure. Despite this, it is tempting to see significant philosophical differences between them. In particular, the idea that diminishing marginal utility constitutes a possible *explanation* of risk-aversion, that EU theory rejects, fits well with non-EU theory. More generally, Allais' ideas that utility is psychological real, that riskless cardinal utility makes sense, and that a rational agent might consciously attend to the dispersion of utilities as well as their expectation, all make good sense from a non-EU perspective, though nothing in the mathematics strictly requires them (cf. Wakker 1994).

⁸ See Harless and Camerer (1994), Hey and Orme (1994) for discussion of the empirical support for and against the various non-EU theories.

In the following section, I argue that the opposition between EU and non-EU theories finds an intriguing parallel in evolutionary biology.

4. Evolution and Rational Choice

Evolution by natural selection is commonly viewed as a maximising process, in which ‘nature’ chooses the phenotypic variants that are best suited to the environment, or fittest. In this metaphorical sense, a ‘choice’ is made whenever natural selection operates. Moreover, and non-metaphorically, the actual choice behaviour of animals (including humans), is very likely the product of natural selection, at least in part. For both these reasons, it is natural to expect connections between choice theory and Darwinian evolution.

The notion of ‘choice behaviour’ as applied to animals requires only behavioural plasticity, not sophisticated cognition. Thus some amphibians and fish must choose whether to produce many large eggs or a few small ones; some plants must choose whether to germinate in a given year or wait for a better one; some animals must choose how long to forage in a given area before searching for a new area, and so-on. This is standard usage in behavioural ecology.

Clearly, many of the choices faced by animals involve an element of risk, in that the consequences of any particular choice for reproductive success are stochastic. For example, suppose an animal may adopt one of two possible foraging strategies, A and B. Strategy A is ‘safe’ – it guarantees the animal 5 units of food in a given time period. Strategy B is ‘risky’ – it brings either 9 units of food or 1 unit per time period, depending on whether the animal has to flee a (non-lethal) predator. Suppose that the probability of encountering a predator is $\frac{1}{2}$. So in effect, the animal must choose between 5 units of food for certain, and a lottery that brings either 9 units or 1 unit with equal probability. Note that both strategies yield the same expected amount of food, namely 5 units. Conceptually, this is similar to a standard rational choice problem. It is then natural to ask: which strategy will be favoured by natural selection, i.e. which is evolutionarily optimal? Will evolution favour risky or safe strategies?

A more general line of enquiry suggests itself. Given that animals face risky choices, and given that their choice behaviour is influenced by natural selection, we can ask what sort of choice behaviour we should expect to find. For example, will evolution produce creatures that obey expected utility maximisation? It is natural to think that the answer should be ‘yes’, on the grounds that evolution generally

produces ‘well-designed’ creatures who behave as if they are rationally pursuing a goal; and EU maximisation is arguably a canon of rationality. This suggestion is merely programmatic; however, a number of authors have tried to show that EU maximisation can be given an evolutionary foundation (e.g. Cooper 2001, Gintis 2009, Stearns 2000, Orr 2007).

I maintain that in fact, the link is strongest between evolution and *non*-expected utility theory, both formally and conceptually. This is so for two related reasons. Firstly, the distinction between diminishing marginal utility and ‘real’ risk aversion, which non-EU theory recognises, finds an analogue in evolutionary theory. Secondly, the characterization of evolutionarily optimal behaviour in the face of risk yields a maximand which is structurally similar to the generic non-EU maximand. I expand on both points below.

5. Evolution, Risk-Aversion, and Non-Expected Utility

Consider again the choice between the two foraging strategies A and B above, i.e. 5 units of food for sure (A), and either 1 or 9 units with probability $\frac{1}{2}$ each (B). Both strategies have the same expected value, but it does not follow that they are evolutionarily equivalent. On the contrary, under many conditions organisms using the risk-averse strategy (A) will enjoy a selective advantage, and thus ultimately dominate the population. Interestingly, this selective advantage may arise for two different, logically independent reasons (cf. Frank and Slatkin 1990, Robson 1996, Okasha 2007.)

The first reason is straightforward: reproductive output may scale concavely with food intake, i.e. additional food leads to additional offspring, but with diminishing returns (Figure 2). Empirically this is quite plausible, and is a common assumption in optimal foraging models. It implies that to maximise expected reproductive output, an animal should exhibit risk aversion in choosing between lotteries with food prizes. So an animal employing strategy A will have a higher expected reproductive output than one employing B; this constitutes a selective advantage and will lead A to dominate the population, *ceteris paribus*. This is the first and most obvious mechanism by which risk aversion may evolve by natural selection.

Figure 2: Concave Fitness Function

The second reason is more subtle; it arises because in certain evolutionary contexts, strategies that lead to a high variance in reproductive output are intrinsically disadvantageous (cf. Gillespie 1977). In such contexts, expected reproductive output is not the sole determinant of evolutionary success, and so is not the correct measure of Darwinian fitness; the variance must be taken into account too. This means that strategy A will have an evolutionary advantage over B, quite apart from considerations of diminishing returns. Even if reproductive output is linear in food intake, the advantage will accrue.

One way to understand this is to invoke the well-known ‘geometric mean principle’ (cf. Lewontin and Cohen 1969, Houston and McNamara 1999). Under certain sorts of environmental stochasticity (described below), the appropriate measure of Darwinian fitness is the geometric, not the arithmetic, mean of reproductive output. (This is intuitive since reproduction is a ‘multiplicative’ process.) In the above example, if predators are present in some time periods and absent in others, each with probability $\frac{1}{2}$ and independent across periods, then the geometric mean principle applies. Supposing for simplicity that x units of food implies x offspring, i.e. reproductive output is linear in food, it follows that strategy A has a

geometric mean output of 5 offspring per time period, while B has 3.⁹ Therefore type A will dominate the population.

These two evolutionary reasons why risk-averse behaviour may be selectively favoured are logically independent. The first involves a concave relation between resources (e.g. food or energy) and reproduction; the second arises because of the intrinsic disadvantage of having a high variability of reproductive output. So the former works by raising the expected output, the latter by reducing the variance of output. The former mechanism may be termed ‘diminishing marginal output’, the latter ‘variance reduction’.¹⁰

This distinction is analogous to the distinction discussed in section 2, between two reasons why a rational agent might exhibit risk-averse preferences, namely a genuine aversion to risk and mere diminishing marginal utility. For diminishing marginal output is analogous to diminishing marginal utility – both involve a concave relation between a resource (food or money), and some thing of value (utility or offspring). (Compare Figures 1 and 2.) It is then natural to suggest that ‘genuine’ risk aversion is analogous to variance reduction – both constitute an *intrinsic* reason why risk-averse behaviour may be favoured, over and above considerations of diminishing marginal returns. Of course, in one case the ‘favouring’ is done by the rational agent, in the other by natural selection.

This argument can be bolstered by considering John Gillespie’s well-known analysis of evolution in stochastic environments. In such environments, the reproductive output of a genotype (or strategy) is a random variable. Gillespie (1977) emphasised that natural selection will penalise genotypes that lead to high variability in reproductive output. He showed that under fairly general assumptions, selection will favour the genotype which maximises an expression of the form:

$$\text{Expected [reproductive output]} - f[\text{Var(reproductive output)}]$$

where f is an increasing function. (The precise function depends on the exact pattern of stochasticity; see section 6.) Gillespie’s formula highlights the potential trade-off

⁹ Recall that if two random variables have the same arithmetic mean, the one with the lower variance will have the higher geometric mean.

¹⁰ Frank and Slatkin (1990) recommend the term ‘risk aversion’ for the former and ‘bet hedging’ for the latter. However this recommendation has not been widely heeded; moreover, some authors use ‘bet-hedging’ to refer to one particular *way* of reducing variability in output, namely using randomised strategies. So I will continue to use ‘risk aversion’ to mean an animal employing a strategy like A over one like B; on this usage, the two mechanisms described above are alternative reasons why risk-aversion might be favoured by natural selection.

between the expectation and the variance of reproductive output – the type with the highest expected output may not be the fittest, if the variance is too high. Note the analogy between Gillespie’s formula and Allais’ discussion of rational behaviour in the face of risk. As we saw, Allais argued that the rational agent might care about the variance of the utilities, not just its expectation, and that a dislike of risk might lead them to prefer a lottery with lower expected utility, if its variance was lower too. If ‘reproductive output’ is replaced with ‘utility’, Gillespie’s formula captures precisely Allais’ conception of what genuine risk-aversion amounts to. So Allais’ idea that risk attitude isn’t captured by the shape of an agent’s utility function has a close evolutionary analogue.

This suggests a link between the evolutionary theory of optimal behaviour in the face of risk and *non*-EU theory – for as we have seen, the latter recognises Allais’ distinction between risk-aversion and diminishing marginal utility, unlike EU theory. The previous paragraph provides some evidence for such a link. Another, more direct piece of evidence is the fact that in experiments on the choice behaviour of rats, violations of the independence axiom of EU theory have been discovered; see Kagel, Battalio and Green (1995) for details. Moreover, the *pattern* of violations the rats exhibited were broadly similar to the violations that humans exhibit. This suggests a common evolutionary origin for the animal and human behaviours. Since non-EU theory was partly designed to account for violations of the independence axiom, and since the pigeon’s choice behaviour has presumably been fashioned by natural selection, this bolsters the link between evolution and non-EU theory.

As discussed in section 3, non-EU theory was motivated by two different concerns: first, to account for observed violations of the EU axioms, and second, to avoid the problematic equation of risk aversion with diminishing marginal utility. It is remarkable that both of these find analogues in evolutionary biology. Animals as well as humans violate the EU axioms, and the analysis of evolutionarily optimal behaviour in the face of risk yields an analogue of the distinction between ‘true’ risk aversion and diminishing marginal utility.

This argument can be bolstered further, and made more precise, by drawing on another well-known analysis of evolution in stochastic environments, due to J. McNamara (1995). Consider lotteries of the form:

(1 unit of food...*n* units of food; $p_1 \dots p_n$) where $\sum p_i = 1$.

(This should be read as ‘get 1 unit of food with probability p_1 , 2 units with probability p_2 etc.) Each lottery can be thought of as implemented by a particular behavioural strategy, or genotype, of an animal. Let $v(x_i)$ equal the reproductive output, or number of offspring, an animal will get from consuming x_i units of food; we assume that $v(x_i)$ is an increasing function of x . So the expected reproductive output an animal will get from choosing a given lottery (i.e. engaging in the behaviour associated with the lottery) equals $\sum_1^n v(x_i)p_i$. One might think that selection will favour the behaviour with the highest value of this quantity, but this ignores variance discounting. The behaviour with the highest expected output may have a high variance, which as we have seen will often constitute a selective penalty.

McNamara (1995) shows that the actual maximand of natural selection, in the above model, equals $\sum_1^n v(x_i)w_i$, where w_i is the *weight* of outcome x_i . This expression is the weighted average of reproductive output, but it is not the expectation, because the *weight of an outcome does not in general equal its true probability*. McNamara provides an explicit formula for calculating the weights; the details of his argument are explained in section 6. For the moment, what matters is the striking similarity between the maximand in this evolutionary model and the generic non-EU evaluation criterion of section 4. We saw that in non-EU theory, the preferred lottery is the one that maximises the weighted average of the utilities of the prizes, where the weights do not equal the true probabilities. Similarly, evolution favours the lottery that maximises the weighted average of reproductive output, where the weights aren’t the true probabilities. It is easy to see that with suitable weights, low-risk lotteries may do best according to this maximand, even if $v(x_i)$ is linear in x .

This shows that there really is a close connection between non-EU theory and evolutionary theory, in respect of the relation between diminishing marginal returns and risk-aversion. Just as concavity of the utility function is one reason, but not the only one, why low-risk lotteries may score highly on the non-EU evaluation criterion, so concavity of the reproductive output function $v(x_i)$ is one reason, but not the only one, why low-risk lotteries may score highly given the maximand of McNamara’s evolutionary model.

To summarize, a number of lines of argument point to a link between evolutionary theory and non-EU theory. Firstly, the distinction between ‘genuine’

risk-aversion and diminishing marginal utility, that the non-EU theorists insist on, has an evolutionary analogue. Secondly, there is direct evidence that animals and humans violate the independence axiom, in similar ways. Thirdly, Gillespie's formula for fitness in stochastic environments, which formalises the notion of variance discounting, is strikingly analogous to Allais' conception of genuine risk aversion. Fourthly, McNamara's analysis shows that the evolutionarily optimal strategy, in a risky environment, has the same mathematical form as the rational choice optimum when calculated according to the non-EU evaluation criterion.

Finally, it is perhaps significant that the utility / fitness analogy works best on an Allais-like conception of utility, rather than the orthodox EU conception. Recall that Allais believed in 'riskless' cardinal utility, where EU theorists held that in the absence of risk, utility can only be ordinal. In a biological context, it is clearly untrue that in the absence of risk, Darwinian fitness is only ordinally measurable; on the contrary, fitness is usually treated as either cardinal or ratio-scale measurable, whether or not the evolutionary model incorporates stochastic elements.¹¹ So in this respect too, the link is strongest between evolution and non-EU theory.

6. *Evolution in Stochastic Environments*

The previous section mentioned a number of results from evolutionary theory, but did not explain their derivation. This section provides more detail, and shows how to reconcile the McNamara and Gillespie results.

Consider again the choice between the two foraging strategies A and B, i.e. 5 units of food of certain, versus 1 or 9 units with probability $\frac{1}{2}$ each. Assume that strategies are genetically hard-wired, and perfectly inherited. Assume further that reproductive output is linear in food, so A and B are in effect lotteries over offspring. Suppose a large population initially contains organisms of both types. How will it evolve? Both types have an arithmetic mean of 5 offspring per time period; but type B has a lower geometric mean, due its higher variance. Can we apply the geometric mean principle and conclude that type A is fitter, so will dominate the population?

It depends. We need to ask whether the risk faced by the type B organisms is 'aggregate' or 'idiosyncratic'.¹² If each type B organism faces an *independent* 50-50

¹¹ See Okasha (2009) for discussion of the appropriate measurement scale for biological fitness.

¹² 'Aggregate' versus 'idiosyncratic' risk is economic terminology; in biology, the same distinction is often captured by contrasting 'environmental' and 'demographic' stochasticity.

gamble on 1 or 9 offspring, i.e. there is a separate coin flip for each organism, then the risk is purely idiosyncratic. If on the other hand a single fair coin flip determines whether *all* the type Bs leave 1, or *all* leave 9, then the risk is purely aggregate.

The weather is a standard example of aggregate risk – a very harsh winter may kill all members of a population. So if there is a 5% chance of a very harsh winter in a given year, then all organisms face a 5% chance of leaving no offspring. By contrast, predation may well give risk to idiosyncratic risk. Each organism may have a 5% chance of getting killed by a predator in any year, and this chance may be independent across organisms. Purely aggregate and purely idiosyncratic risk are opposite ends of a spectrum; most real cases will lie somewhere in between.

To see how the aggregate / idiosyncratic distinction affects the applicability of the geometric mean principle, suppose firstly that the risk facing the type Bs is purely idiosyncratic. How will our population evolve? Over a single time period, each type A will leave exactly 5 offspring. Each type B will leave either 1 or 9 offspring, with probability $\frac{1}{2}$ each. Since the population is large, then the Bs will leave approximately 5 offspring per capita, by the law of large numbers. So there will be no (or minimal) evolutionary change – the A and B types are equally fit, leaving 5 offspring per capita per period. (Essentially, large population size cancels out the idiosyncratic risk.) The correct measure of fitness, in this model, is simply expected reproductive output – variance is not relevant. So the geometric mean principle does not apply. Note that this argument does not work if the population size is small.

Now suppose the risk is purely aggregate. As before, each type A leaves exactly 5 offspring. With probability $\frac{1}{2}$, all the Bs leave 9 offspring each, and with probability $\frac{1}{2}$, they all leave 1 each. So each period, the A's multiply their numbers by 5, and the Bs by either 1 or 9, with probability $\frac{1}{2}$ each. Assume that the aggregate risk faced by the Bs is independent across periods. Now, there *will* be evolutionary change – the As have an advantage. To see this, consider a 'typical' sequence of 10 years for the Bs: {9, 1, 9, 9, 1, 1, 9, 9, 1, 1}. The product of these numbers is far less than 5^{10} , so in the limit, the As will take over the population.¹³ The geometric mean is the right measure of fitness in this case, and the B type has the lower geometric mean output, owing to its greater variability.

¹³ Appealing to a 'typical' sequence is purely a heuristic; some simple limit calculations, omitted here, generate the result that the As will take over; see Houston and McNamara (1999) or Robson (1996).

A different (though equivalent) perspective on this case is useful (cf. Robson 1996). At any point in the future, the expected *number* of As and Bs in the population is the same. But this does not mean that the two types are equally fit, for what matters to evolution is relative success, not absolute success. And crucially, the expected *proportion* of type As exceeds the expected proportion of type Bs, at every point in time. For example, after one period the ratio of As to Bs will be either 5/9 or 5/1, with probability $\frac{1}{2}$ each; so the expected ratio is 2.7: 1. Essentially, the high variability in output of the type Bs reduces the expected fraction of the population that they will comprise. In the limit, this means that type A will dominate the population.

Gillespie (1977) provided two useful (approximate) formulae for determining the fitness of a type, under the extremes of purely idiosyncratic and purely aggregate risk. If the risk is purely idiosyncratic risk, a type's fitness is given by:

$$\text{Exp [reproductive output]} - \text{Var [reproductive output]} / N$$

where 'N' is the population size. As N gets larger and larger, the second term gets smaller and smaller, so the expectation becomes the major determinant of fitness and the variance irrelevant.¹⁴ This explains our first case above. If the risk is purely aggregate, then a type's fitness is given by:

$$\text{Exp [reproductive output]} - \text{Var [reproductive output]} / 2.$$

This shows that variability in output is heavily penalised when the risk is aggregate, as we saw in our second case.

Gillespie's formulae cover the two extreme cases, but it would be good to have an analysis that applies when the risk is partially idiosyncratic and partially aggregate. McNamara (1995) and Robson (1996) provide such an analysis. In their model, there is a set of possible environmental states S, and a probability distribution over S; the state varies from year to year, with independence across years. There are a number of different types of organism. The reproductive output of an organism depends on both its type and the state of the environment. Consider a given type, called 'A'. In any given environmental state s, each organism of type A faces an *independent* lottery over offspring. So for example, in state s, each type A organism might have a $\frac{1}{2}$ chance of leaving 4 offspring and an $\frac{1}{2}$ chance of leaving none. Let $r_A(s)$ equal the mean reproductive output of type A in environmental state s, averaging over idiosyncratic risk. Let $p(s)$ equal the probability that the environment is in state s.

¹⁴ Many evolutionary analyses assume an infinite population; in this case, the second term in Gillespie's formula goes to zero, under purely idiosyncratic risk.

One might think that natural selection will favour the type with the highest value of $\sum_s r(s)p(s)$, i.e. the arithmetic mean over states of the type's expected reproductive output. But this is incorrect. Since the environment state varies from year to year, with independence across years, the geometric mean over states is the relevant quantity. Maximising the geometric mean is equivalent to maximising its logarithm; since the logarithm of the geometric mean of a random variable equals the arithmetic mean of its logarithm, it follows that selection will maximise $\sum_s \log r(s)p(s)$. The type with the highest value of this expression will dominate the population. This expression thus provides a measure of Darwinian fitness, in a very general model which combines both idiosyncratic and aggregate risk.

Interestingly, maximisation of the above expression may require the type to play a mixed strategy, i.e. to randomise over pure strategies. This is quite intuitive. Suppose that each pure strategy leaves zero offspring in some environmental state, and that each state obtains with non-zero probability. Then, pure strategies are doomed to extinction in the long run. Randomising over pure strategies, or 'bet-hedging' as biologists call it, is the optimal thing to do in such a circumstance.

It is easy to see that the McNamara / Robson model is broadly compatible with Gillespie's formulae. Since the logarithmic function is concave, the maximand in the McNamara / Robson model, $\sum_s \log r(s)p(s)$, implies that a type with a high variance in reproductive output (across states) will be at a disadvantage compared to a type with the same expected value of $r(s)$ but less variance. Variance-discounting of the sort embodied in Gillespie's principles thus falls directly out of the McNamara/Robson analysis. In the case where all the risk is idiosyncratic, i.e. $r(s) = r$ is constant across all states s for each type, the McNamara/Robson model implies that the fittest type maximises $\log r$, or equivalently r itself, i.e. expected reproductive output. This is compatible with Gillespie's formula because McNamara/Robson are assuming a very large population, so the first term in Gillespie's formula will dominate.

McNamara (1995) provides an alternative, equivalent characterisation of the fittest type, in this model. He shows that the type which maximises $\sum_s \log r(s)p(s)$ will thereby maximise $\sum_s r(s)p^*(s)$, where $p^*(s)$ is a distortion of the true probability

distribution over states. The latter expression is thus the weighted average of expected reproductive success, where the weights differ from the true probabilities. The distorted probability distribution $p^*(s)$ is related to the true probability distribution $p(s)$ by the formula $p^*(s) \propto p(s) / \bar{r}(s)$, where $\bar{r}(s)$ is the average value of $r(s)$ in the whole population, i.e. average fitness. Essentially, therefore, fitness maximisation requires underweighting environmental states in which the population as a whole does well, and overweighting states in which it does badly, relative to the state's true probability of occurrence.¹⁵

7. Evolution of Irrationality?

So far, we have highlighted a thematic link between evolution and non-EU theory. Does it follow from this that we should expect the Darwinian process to produce humans with non-EU preferences, e.g. ones that violate the independence axiom? If so this would be a striking conclusion, for it would arguably amount to the 'evolution of irrationality', given the widespread view that the EU axioms are normatively compelling. It might also explain why those axioms are routinely violated by experimental subjects.

This issue was addressed by Robson (1996), who uses his evaluation criterion (above) to argue that in the presence of aggregate risk, evolution may indeed favour irrational behaviour; a similar argument is made by Houston, McNamara and Steer (2007). To understand Robson's argument, suppose that individuals face lotteries of the form (1 unit of food...n units of food; $p_1 \dots p_n$). Assume that food translates into reproductive output according to a positive function v , assumed linear for simplicity.¹⁶ So an individual's behaviour, i.e. choice of lottery, induces a gamble over numbers of offspring. Robson then applies his evaluation criterion to determine the evolutionarily optimal lottery. He then asks: suppose evolution were trying to design a human with preferences over lotteries, such that they always prefer lotteries which are nearer the evolutionary optimum. What sort of preferences would the human have? Would those preferences have an EU representation, or not?

¹⁵ This highlights the important fact that with aggregate risk, there is automatically a game-theoretic or strategic aspect to an individual's optimal choice of lottery; see Robson (1996), McNamara (1995) and Houston and McNamara (1999) for discussion of this point.

¹⁶ This assumption is for expository convenience only; in Robson's own argument it is not used.

The answer depends on the nature of the risk. With purely idiosyncratic risk, then EU maximisation falls directly out of evolutionary optimality, as Robson notes. In that case, by simply equating utility with number of offspring (or food, given our linearity assumption), we generate the result that an EU maximising individual will make the evolutionarily optimal choice. This follows from the fact that with purely idiosyncratic risk, evolution selects for the maximisation of expected reproductive output, as we saw above. Define utility as reproductive output, and EU maximisation will therefore coincide with evolutionary optimality.

What about purely aggregate risk? Robson does not discuss this case, but it is straightforward to show that if utility is equated with the *logarithm* of offspring number (or food, given our linearity assumption), then agents who maximise EU will make evolutionarily optimal choices. That this is so can be seen by inspecting the McNamara/Robson evaluation criterion $\sum_s \log r(s)p(s)$. With purely aggregate risk, $r(s)$ is simply the *actual* number of offspring left by an organism of a given type in environment s (rather than an average across idiosyncratic risk). Therefore, natural selection maximises the expectation of the logarithm of reproductive output. Define utility as the logarithm of reproductive output, and EU maximisation will again produce evolutionary optimality.

This result is interesting, in that it *appears* to supply an evolutionary foundation for the idea of logarithmic utility – an idea with a famous history in decision theory.¹⁷ A number of authors, including Sinn (2003) and Stearns (2000) have been impressed with this fact. However, the result holds only for the case of *purely* aggregate risk, which is a special circumstance – a point that Sinn and Stearns do not mention. With any combination of aggregate and idiosyncratic risk, the result does not go through.

What does happen when aggregate and idiosyncratic risk are combined? In this case, EU maximisation cannot be so simply recovered. This should not be surprising, given McNamara's characterisation of the evolutionarily optimal choice as a weighted average of reproductive output, using *biased* probabilities. Another way to see this is to note that, in the combined case, the evaluation criterion $\sum_s \log r(s)p(s)$ is

¹⁷ Daniel Bernoulli (1738), one of the founders of EU theory, proposed logarithmic utility as a way of avoiding the famous St. Petersburg paradox.

an expectation over environmental states of $\log r(s)$, but $r(s)$ is *itself* an expectation (over idiosyncratic risk); so all the summation terms cannot be brought to the left-hand side. As Robson (1996) points out, this means that the evaluation criterion is not a function of the marginal probabilities facing a given individual – which immediately implies that we are outside the ambit of EU.

A simple example may help make this clear. Consider an individual trying to evaluate the lottery ‘9 offspring or 1 offspring, probability $\frac{1}{2}$ each’. The lottery that the individual faces may be ‘realised’ by many different combinations of aggregate and idiosyncratic risk; as specified, the lottery is compatible with the risk being purely idiosyncratic, purely aggregate, or some combination. And crucially, the McNamara/Robson evaluation criterion applies to these realizations, not to the marginal lottery. (It is impossible to apply the criterion directly to the lottery – there is not enough information.) Further, the different realizations will all receive different evaluations, i.e. have different consequences for Darwinian fitness. But from the point of view of the individual decision maker, they are all equivalent. So evolutionary optimality and individual rationality part ways.

A neat example from Robson (1996) highlights the consequences of this parting of ways. Consider two lotteries A and B. In A, then with probability $\frac{1}{2}$ everyone in the population leaves 9 offspring, and with probability $\frac{1}{2}$ everyone leaves 1. In B, then with probability 1 everyone in the population has a 50:50 chance of leaving 1 offspring or 8.5 offspring, with independence across all population members. Consider the situation from the perspective of a single individual in the population. That individual must surely prefer A to B, for A stochastically dominates B, i.e. for every number of offspring, the probability of having that number is higher under A than B. However, on the McNamara/Robson evaluation criterion, B scores higher than A.¹⁸ Natural selection will favour individuals who choose B over A, even though this choice seems irrational.

Robson uses this example to highlight a violation of stochastic dominance; but in fact, this also means that the independence axiom of EU theory is violated too – since stochastic dominance is a logical consequence of independence. This appears to show quite generally that evolution can lead to violations of EU maximisation and thus (arguably) to irrationality.

¹⁸ The evaluation of B is $\log [(\frac{1}{2}).1 + (\frac{1}{2}).8.5] = \log 4.75$. The evaluation of A is $\frac{1}{2} \log 9 + \frac{1}{2} \log 1 = \frac{1}{2} \log 9$, which is less than $\log 4.75$.

However, this conclusion rests on an implicit assumption, namely that an individual's utility function must depend only on their *own* reproductive output. *Modulo* this assumption, it is certainly true that no utility function exists such that an EU maximiser will be led to make evolutionarily optimal choices, except in the limiting cases of purely idiosyncratic and purely aggregate risk. But what if we relax this assumption? Curry (2001) has argued that EU maximisation can then in fact be recovered; his argument draws on Grafen (1999), which itself draws on McNamara (1995).

The Curry/Grafen point is simple. It derives from the fact, discussed above, that the McNamara/Robson evaluation criterion can equivalently be characterized as $\sum_s r(s) p^*(s)$, where $p^*(s)$ is a biased probability distribution. Recall that $p^*(s)$ is given by the relation $p^*(s) \propto p(s) / \bar{r}(s)$, i.e. the biased distribution shifts probability mass onto states where the population as a whole does badly (has low $\bar{r}(s)$), and away from states where the population as a whole does well, relative to the true distribution. Re-arranging, the criterion can be written as $\sum_s r(s) / \bar{r}(s) p(s)$. This is an expectation over environmental states of the quantity, $[r(s) / \bar{r}(s)]$, which is the *relative* reproductive output of a given type in state s , i.e. its output divided by average population output in that state. So the maximand of natural selection is expected *relative* fitness.

Crucially, this means that the evaluation of a lottery *is* after all a function of the marginal probabilities facing an individual, so long as we mean the marginal probabilities of having a given *relative* reproductive output, rather than absolute output. Moreover, the criterion is linear in those probabilities – so by suitable choice of utility function, EU maximisation must be restorable. The choice is not hard to find: simply let an individual's utility depend on their relative number of offspring (relative to the population average). If the individual obeys EU maximisation, they will then be led to make evolutionarily optimal choices.

It may seem puzzling that the same evolutionary model can apparently imply that EU maximisation is satisfied *and* that it is violated. How can this be? In fact there is no contradiction here. When Robson argues that EU maximisation is violated, and when Curry argues that it is not, they are (in effect) using different state spaces to set up the decision problem. Robson is assuming that the basic prizes, from which the

lotteries are constructed, are specifications of how many offspring an individual leaves. It then follows that the evolutionarily optimal preferences over lotteries will violate the EU axioms. Curry is assuming that the basic prizes are specifications of how many offspring an individual has *relative to the population average*. It then follows that the evolutionarily optimal preferences over lotteries will satisfy the EU axioms. In a way, this is an instance of the familiar moral that apparent irrationalities of choice can often be removed by enlarging the state space.

What then of the original question: will evolution tend to produce creatures that obey EU maximization? No definitive answer emerges. Curry's argument suggests that in theory the answer is 'yes'; but in practice, it may be extremely hard to produce people with utility functions that depend suitably on relative reproductive output. This is because one's relative output depends not only on how many offspring one has oneself, but also on how many others in the population have, so is extremely hard to keep track of. It may be that producing creatures that care about their absolute number of offspring (or things that promote it such as food, sex etc.), is the best that evolution can do, even though it would be better to produce organisms who cared directly about their relative number. If that is so, then we should expect to see violations of EU maximisation.

A loose analogy of the situation is this. Consider the one-shot Prisoner's dilemma played among (human) relatives. If an individual cares only about their own payoff, then defection is the rational choice. But the evolutionary optimum may be to cooperate, because of standard kin selection considerations. Should we then expect evolution to produce irrational behaviour? Perhaps, but alternatively evolution may simply produce people who care about the welfare of their kin; this would re-align individual rationality and evolutionarily optimality. And evolution seems to have actually achieved this. In the case of risk, individual rationality and evolutionary optimality might also be restored by suitably modifying people's preferences, i.e. making them care about their relative fitness. But this will be much harder for evolution to pull off – for keeping track of relative fitness is at the very least cognitively demanding, and probably outright impossible for most organisms.

If this is correct, then it is quite plausible that evolution will tend to produce creatures that violate EU maximisation. Of course, this does not show that the specific patterns of EU violation discovered by behavioural economists can be given an

evolutionary rationale; whether that can be shown is an open question, requiring further work. But the foregoing considerations do suggest that it is a real possibility.

8. Conclusion

This paper has explored connections, both thematic and formal, between the economic and evolutionary theories of choice in the face of risk. That such connections exist is not surprising, given that optimization is central to both bodies of theory, but there have been relatively few attempts to explore the connections in detail. I have argued that the link is strongest, surprisingly, between evolution and non-EU theory; and in particular that there is a remarkable evolutionary analogue of the distinction between diminishing marginal utility and ‘real’ risk aversion, a distinction that cannot be drawn in orthodox EU theory. Further work will be needed to establish whether the specific versions of non-EU theory currently in vogue, e.g. prospect theory, rank-dependent utility etc. can be given an evolutionary foundation. Whatever the answer, I hope to have shown that the conceptual connections between rational choice theory and evolution are interesting and worthy of exploration.

References

- Allais, M. (1952) 'Le comportement de l'homme rationnel devant le risque, critique des postulats et axiomes de l'école Américaine', *Econometrica* 21, 503-46.
- Arrow, K. (1951), 'Alternative approaches to the theory of choice in risk-taking situations', *Econometrica* 19, 404-437.
- Allais, M. and Hagen, O. (1979) (eds.) *Expected utility hypothesis and the Allais paradox*, Dordrecht: Reidel.
- Bernoulli, D. (1738) 'Specimen theoriae novae de mensura sortis', translated as 'Exposition of a new theory on the measurement of risk', *Econometrica* 22, 23-36.
- Cooper, W. S. (2001) *The Evolution of Reason*, Cambridge: Cambridge University Press.
- Curry, P. (2001) 'Decision making under uncertainty and the evolution of interdependent preferences,' *Journal of Economic Theory*, 98, 357-69.
- Frank, S.A. and Slatkin, M. (1990) 'Evolution in variable environments', *American Naturalist* 136, 2, 244-60.
- Gillespie, J. (1977) "Natural selection for variances in offspring number: a new evolutionary principle", *American Naturalist* 111, 1010-14.
- Gintis, H. (2009) *The Bounds of Reason*, Princeton: Princeton University Press.
- Grafen, Alan (1999) 'Formal darwinism, the individual-as-maximizing-agent analogy, and bet-hedging' *Proceedings of the Royal Society B* 266: 799-803.
- Harless, D. W. and Camerer, C. F. (1994) 'The predictive utility of generalized expected utility theories', *Econometrica* 62: 1251-1289.
- Hey, J. D. and Orme, C. (1994) 'Investigating generalisations of expected utility theory using experimental data', *Econometrica* 62, 1291-1326.
- Houston, A. I. and McNamara, J. N. (1999) *Models of Adaptive Behaviour*, Cambridge: Cambridge University Press.
- Houston, A.I., McNamara, J. N. and Steer, M. D. (2007) 'Do we expect natural selection to produce rational behaviour?', *Philosophical Transactions of the Royal Society of London B*, 362, 1531-43.
- Kagel, J. H., Battalio, R. C. and Green, C. (1995) *Economic Choice Theory: an experimental analysis of animal behaviour*, Cambridge: Cambridge University Press.
- Kahneman, D. and Tversky, A. (1979) "Prospect theory: an analysis of decision under risk", *Econometrica* 47, 263-91.

- Lewontin, R. and Cohen, D. (1969) 'On population growth in a randomly varying environment', *Proceedings of the National Academy of the Sciences USA* 62, 1056-60.
- Machina, M. (1982) 'Expected utility' analysis without the independence axiom', *Econometrica* 50, 277-323.
- Machina, M. (2008) 'Non-expected utility theory', *The New Palgrave Dictionary of Economics*, 2nd edition, eds. S. N. Durlauf and L. E. Blume, Palgrave Macmillan.
- McNamara, J. N. (1995) 'Implicit frequency dependence and kin selection in fluctuating environments', *Evolutionary Ecology* 9, 185-203
- Okasha, S. (2007) 'Rational choice, risk aversion and evolution', *Journal of Philosophy* 104(5), 217-235.
- Okasha, S. (2009) 'Individuals, groups, fitness and utility', *Biology and Philosophy* 24: 561-84
- Orr, H A. (2007) 'Absolute fitness, relative fitness, and utility' *Evolution* 61, 2997-3000.
- Quiggin, J. (1982) 'A theory of anticipated utility', *Journal of Economic Behavior and Organisation* 3(4), 323-43.
- Quiggin, J. (1993) *Generalized Expected Utility Theory: The Rank-Dependent Model*, Amsterdam: Kluwer.
- Robson, A. (1996) 'A biological basis for expected and non-expected utility', *Journal of Economic Theory* 68, 397-424.
- S. Hansson (1988) "Risk aversion as a problem of conjoint measurement", in P. Gärdenfors and N. Sahlin, eds., *Decision, Probability and Utility*, Cambridge: CUP
- Savage, L. J. (1954), *The Foundations of Statistics*, New York: Wiley.
- Seeger, J. and Brockmann H. J. (1987) 'What is bet-hedging?' *Oxford Surveys in Evolutionary Biology* 4:182-211.
- Sinn, H. (2003) 'Weber's law and the biological evolution of risk preferences', *Geneva Papers on Risk and Insurance Theory*, 28(2), 87-100.
- Stearns, S. C. (2000) 'Daniel Bernoulli (1738): evolution and economics under risk', *Journals of Biosciences*, 25, 3, 221-8.
- von Neumann, J. and Morgenstern, O. (1944) *Theory of Games and Economic Behaviour*, Princeton: Princeton University Press.
- Wakker, P. P. (1994), 'Separating marginal utility and probabilistic risk aversion', *Theory and Decision* 36, 1-44.

Yaari, M.E. (1987) 'The dual theory of choice under risk', *Econometrica* 55, 95-115.