Optimal Choice in the Face of Risk: Decision Theory Meets Evolution*

Samir Okasha[†‡]

The problem of how to make optimal choices in the face of risk (or uncertainty) arises in both economics/decision theory and also evolutionary biology; in the former, 'optimal' means utility maximizing, while in the latter it means fitness maximizing. This article explores the links, thematic and formal, between the economic and evolutionary theories of optimal choice in risky situations, with particular reference to the relationship between utility and fitness. It is argued that the link is strongest between evolution and 'nonexpected' utility theory, rather than traditional expected utility theory.

1. Introduction. The problem of how a rational agent should choose between risky options, or lotteries, is famous. 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 maximizing the expected value of a utility function. The shape of the utility function then reflects the agent's attitude toward 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, that is, prefer $x for certain to any lottery with an expected value of $x.

Crucial to EU theory is the famous independence axiom, which is the

*Received May 2010; revised August 2010.

†To contact the author, please write to: Department of Philosophy, University of Bristol, 9 Woodland Road, Bristol BS8 1TB, United Kingdom; e-mail: Samir.Okasha@bristol.ac.uk.

‡Thanks to Ken Binmore, John McNamara, Jonathan Wright, Brian Skyrms, David Makinson, Peter Godfrey-Smith, Patrick Forber, Cedric Paternotte, and Jonathan Grose for discussion and comments and to two anonymous referees for constructive criticism. Research for this article was supported by grant H/F017502/1 from the UK Arts and Humanities Research Council.

Philosophy of Science, 78 (January 2011) pp. 83–104. 0031-8248/2011/7801-0005$10.00
Copyright 2011 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/658115 Published online by Cambridge University Press

subject of much controversy.[1] 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 'nonexpected utility theories', which relax or replace the independence axiom, thus leading to maximization of some quantity other than EU. These theories have a certain amount of empirical support, although 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 forage either 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 means greater survival and thus more offspring and higher Darwinian fitness. Which foraging strategy will be favored by natural selection? Theoretical work shows that in many circumstances, strategies that reduce variability in resource acquisition will be selectively advantageous; such behavioral strategies are often called 'risk averse'. Empirical 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 Kagel, Battalio, and Green 1995; Robson 1996; Stearns 2000; Orr 2007.) This is surprising because in other areas such as game theory, there has been significant cross-fertilization between the economic and evolutionary discussions. This article takes a preliminary step toward 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 maximize; the latter, that natural selection tries to maximize.) This analogy has been noticed before and has led many authors to posit a link between evolution and EU maximization (Stearns

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, e.g., in the title of Wilbur and Rudolf (2006). This does no harm since the economist's distinction between risk (known chances) and uncertainty (unknown chances) has no obvious application to nonhuman animals.

https://doi.org/10.1086/658115 Published online by Cambridge University Press

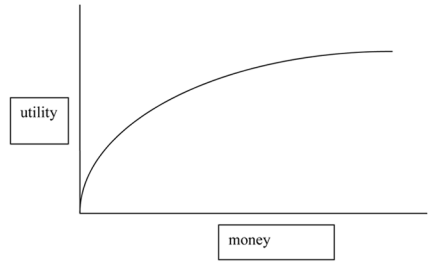


Figure 1. Concave Utility Function.

2000; Cooper 2001; 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 although philosophically rich facet of the debate between EU and non-EU theory.

2. EU Theory, Allais's Critique, and the Concept of Risk Aversion. On orthodox EU theory, an agent's attitude toward 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, that is, diminishing marginal utility (fig. 1). The reason is obvious: if and only if the agent has risk-averse preferences—that is, always prefers $x to any lottery with an expected value of $x—and she maximizes EU, 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.

Although 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 recognize that there are two different reasons why agents might exhibit risk-averse preferences. The first is their attitude toward wealth; the second, their attitude toward 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

https://doi.org/10.1086/658115 Published online by Cambridge University Press

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 an 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 he 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 his 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 à la von Neumann and Morgenstern. Thus, Savage cautioned against confusing von Neumann–Morgenstern utility with "the now almost obsolete notion of utility in riskless circumstances" (1954, 93), while Arrow called riskless cardinal utility "a meaningless concept" (1951, 425).

Maurice Allais (1952), an early critic of EU theory, explicitly defended the existence of riskless cardinal utility. He insisted that utility was psychologically 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's insistence on the distinction between risk aversion and diminishing marginal utility and his defense 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.

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

https://doi.org/10.1086/658115 Published online by Cambridge University Press

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's 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's position as a confusion.[3]

However, given his starting assumptions, Allais's 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's 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 .5 each. The agent's risk attitude will then enter the picture. The agent may choose the option with the highest EU, that is, max [u(a + 5), .5 u(a + 1) + .5 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's 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 toward risk, reflected in the agent's attention to the variance as well as the expectation of the utilities associated with each option.

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

3. Nonexpected Utility Theory. Non-EU theory, developed in works by Kahneman and Tversky (1979), Machina (1982), Quiggin (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

3. See, e.g., the papers by di Finetti and Morgenstern in Allais and Hagen (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 a precise account of how exactly the risk-averse agent would trade off variance in utility (or higher moments) against the expectation, a point that his critics were quick to pick up on. See the papers in Allais and Hagen (1979).

https://doi.org/10.1086/658115 Published online by Cambridge University Press

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, . . . , $x_n; $p_1, . . . , p_n$), that is, 'get a prize of $x_1 with probability p_1, of $x_2 with probability p_2, and so on'. In EU theory, an agent evaluates the lotteries according to the criterion $\Sigma_i u(x_i)\, p_i$, 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 $\Sigma_i\, u(x_i)\, w_i$, where w_i is the *decision weight* of prize i, with $\Sigma_i\, w_i = 1$. The decision weight of a prize need not equal the probability of getting it. This means that agents are maximizing 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) 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 maximizing some sort of non-EU. Other empirical features of decision making in risky situations, such as the general tendency toward 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, that is, indifferent between $x and any gamble with an expected 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

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.

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

https://doi.org/10.1086/658115 Published online by Cambridge University Press

looks like, the agent may be risk averse, risk neutral, or risk loving; nothing can be deduced about his risk attitude from the fact that his utility function is linear in money or vice versa.

This means that in non-EU theory, there are two distinct ways in which risk-averse behavior (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 generates 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 toward risk is one thing, his 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, which EU theory rejects, fits well with non-EU theory. More generally, Allais's ideas that utility is psychologically real, that riskless cardinal utility makes sense, and that a rational agent might consciously attend to the dispersion of utilities as well as his expectation all make good sense from a non-EU perspective, although nothing in the mathematics strictly requires them (cf. Wakker 1994). 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 maximizing 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 nonmetaphorically, the actual choice behavior 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 behavior' as applied to nonhumans requires only behavioral 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 behavioral ecology.

https://doi.org/10.1086/658115 Published online by Cambridge University Press

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 that 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 (nonlethal) predator. Suppose that the probability of encountering a predator is .5. So in effect, the animal must choose between 5 units of food for certain and a lottery that brings either 9 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 favored by natural selection, that is, which is evolutionarily optimal? Will evolution favor risky or safe strategies?

A more general line of inquiry suggests itself. Given that animals face risky choices, and given that their choice behavior is influenced by natural selection, we can ask what sort of choice behavior we should expect to find. For example, will evolution produce creatures that obey EU maximization? 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 maximization is arguably a canon of rationality. This suggestion is merely programmatic; however, a number of authors have tried to show that EU maximization can be given an evolutionary foundation (e.g., Stearns 2000; Cooper 2001; Orr 2007; Gintis 2009).

I maintain that, in fact, the link is strongest between evolution and non-EU theory, both formally and conceptually. This is so for two related reasons. First, the distinction between diminishing marginal utility and 'real' risk aversion, which non-EU theory recognizes, finds an analogue in evolutionary theory. Second, recent work has shown that the characterization of evolutionarily optimal behavior in the face of risk yields a maximand that is structurally similar to the generic non-EU maximand. I expand on both points below.

5. Evolution, Risk Aversion, and Nonexpected Utility. Consider again the choice between the two foraging strategies A and B above, that is, 5 units of food for sure (A) and either 1 or 9 units with probability .5 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

https://doi.org/10.1086/658115 Published online by Cambridge University Press

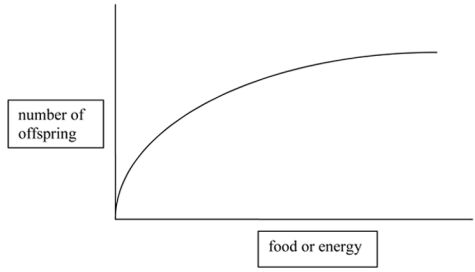


Figure 2. Concave Fitness Function.

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, that is, additional food leads to additional offspring but with diminishing returns (fig. 2). Empirically this is quite plausible and is a common assumption in optimal foraging models. It implies that to maximize 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.

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 .5 and independent across periods, then the geometric mean prin-

https://doi.org/10.1086/658115 Published online by Cambridge University Press

ciple applies. Supposing for simplicity that x units of food implies x offspring, that is, reproductive output is linear in food, it follows that strategy A has a geometric mean output of five offspring per time period, while B has three.[9] Therefore, type A will dominate the population.

These two evolutionary reasons why risk-averse behavior may be selectively favored 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; cf. figs. 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 behavior may be favored, over and above considerations of diminishing marginal returns. Of course, in one case the 'favoring' 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) emphasized that natural selection will penalize genotypes that lead to high variability in reproductive output. He showed that under fairly general assumptions, selection will favor the genotype that maximizes an expression of the form

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

where f is an increasing function. (The precise function depends on the exact pattern of stochasticity; see sec. 6.) Gillespie's formula highlights

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 randomized 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 favored by natural selection.

https://doi.org/10.1086/658115 Published online by Cambridge University Press

the potential trade-off 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's discussion of rational behavior 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 EU, if its variance was lower too. If 'reproductive output' is replaced with 'utility', Gillespie's formula captures precisely Allais's conception of what genuine risk aversion amounts to. So Allais's idea that risk attitude is not 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 behavior in the face of risk and non-EU theory—for, as we have seen, the latter recognizes Allais's 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 behavior of rats, violations of the independence axiom of EU theory have been discovered (see Kagel et al. 1995 for details). Moreover, the pattern of violations the rats exhibited was broadly similar to the violations that humans exhibit. This suggests a common evolutionary origin for the animal and human behaviors. Since non-EU theory was partly designed to account for violations of the independence axiom, and since the rats' choice behavior 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 behavior 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 McNamara (1995). Consider lotteries of the form

$$(1 \text{ unit of food } \ldots \text{ n units of food; } p_1 \ldots p_n), \text{ where } \sum p_i = 1.$$

(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 behavioral strategy, or genotype, of an animal. Let $v(x_i)$ equal the reproductive output, or number of offspring, an animal will get

https://doi.org/10.1086/658115 Published online by Cambridge University Press

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 behavior associated with the lottery) equals $\Sigma_i\ v(x_i)\ p_i$. One might think that selection will favor the behavior with the highest value of this quantity, but this ignores variance discounting. The behavior 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 $\Sigma_i\ v(x_i)\ w_i$, 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 maximizes the weighted average of the utilities of the prizes, where the weights do not equal the true probabilities. Similarly, evolution favors the lottery that maximizes the weighted average of reproductive output, where the weights are not 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. First, the distinction between 'genuine' risk aversion and diminishing marginal utility, which the non-EU theorists insist on, has an evolutionary analogue. Second, there is direct evidence that animals and humans violate the independence axiom, in similar ways. Third, Gillespie's formula for fitness in stochastic environments, which formalizes the notion of variance discounting, is strikingly analogous to Allais's conception of genuine risk aversion. Fourth, 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.

https://doi.org/10.1086/658115 Published online by Cambridge University Press

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, that is, 5 units of food for certain versus 1 or 9 units with probability .5 each. Assume that strategies are genetically hardwired 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 five 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 and 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 gamble on one or nine offspring, that is, there is a separate coin flip for each organism, then the risk is purely idiosyncratic. If however a single fair coin flip determines whether all the type Bs leave one or all leave nine, 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 rise 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

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.

https://doi.org/10.1086/658115 Published online by Cambridge University Press

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 first 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 five offspring. Each type B will leave either one or nine offspring, with probability .5 each. Since the population is large, the Bs will leave approximately five 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 five 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 five offspring. With probability .5, all the Bs leave nine offspring each, and with probability .5, they all leave one each. So each period, the As multiply their numbers by 5, and the Bs by either 1 or 9, with probability .5 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.

A different (although 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 .5 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) formulas for de-

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 Robson (1996) or Houston and McNamara (1999).

https://doi.org/10.1086/658115 Published online by Cambridge University Press

termining 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

$$\exp[\text{reproductive output}] - \frac{\text{Var}[\text{reproductive output}]}{N},$$

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 becomes irrelevant.[14] This explains our first case above. If the risk is purely aggregate, then a type's fitness is given by

$$\exp[\text{reproductive output}] - \frac{\text{Var}[\text{reproductive output}]}{2}.$$

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

Gillespie's formulas 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 organisms. 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 .5 chance of leaving four offspring and a .5 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.

One might think that natural selection will favor the type with the highest value of Σ_s r(s) p(s), that is, the arithmetic mean over states of the type's expected reproductive output. But this is incorrect. Since the environmental state varies from year to year, with independence across years, the geometric mean over states is the relevant quantity. Maximizing the geometric mean is equivalent to maximizing 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 maximize Σ_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

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.

https://doi.org/10.1086/658115 Published online by Cambridge University Press

fitness, in a very general model that combines both idiosyncratic and aggregate risk.

Interestingly, maximization of the above expression may require the type to play a mixed strategy, that is, to randomize 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 nonzero probability. Then, pure strategies are doomed to extinction in the long run. Randomizing 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 formulas. Since the logarithmic function is concave, the maximand in the McNamara-Robson model, Σ_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 in which all the risk is idiosyncratic, that is, r(s) = r is constant across all states s for each type, the McNamara-Robson model implies that the fittest type maximizes log r, or equivalently r itself, that is, expected reproductive output. This is compatible with Gillespie's formula because McNamara and Robson are assuming a very large population, so the first term in Gillespie's formula will dominate.

McNamara (1995) provides an alternative, equivalent characterization of the fittest type, in this model. He shows that the type that maximizes Σ_s log r(s) p(s) will thereby maximize Σ_s log 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, that is, average fitness. Essentially, therefore, fitness maximization requires underweighting environmental states in which the population as a whole does well and overweighting states in which it does poorly, 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

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 McNamara (1995), Robson (1996), and Houston and McNamara (1999) for discussion of this point.

https://doi.org/10.1086/658115 Published online by Cambridge University Press

preferences, for example, 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 favor irrational behavior; 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 . . . p_n). Assume that food translates into reproductive output according to a positive function v, assumed to be linear for simplicity.[16] So an individual's behavior, that is, choice of lottery, induces a gamble over number of offspring. Robson then applies his evaluation criterion to determine the evolutionarily optimal lottery. He then asks, suppose evolution were trying to design humans with preferences over lotteries, such that they always prefer lotteries that are nearer the evolutionary optimum, what sort of preferences would the humans have? Would those preferences have an EU representation or not?

The answer depends on the nature of the risk. With purely idiosyncratic risk, EU maximization 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-maximizing individual will make the evolutionarily optimal choice. This follows from the fact that with purely idiosyncratic risk, evolution selects for the maximization of expected reproductive output, as we saw above. Define utility as reproductive output, and EU maximization 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 maximize EU will make evolutionarily optimal choices. That this is so can be seen by inspecting the McNamara-Robson evaluation criterion Σ_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 maximizes the expectation of the logarithm of reproductive output. Define utility as the logarithm of reproductive output, and EU maximization will again produce evolutionary optimality.

This result is interesting, in that it appears to supply an evolutionary

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

https://doi.org/10.1086/658115 Published online by Cambridge University Press

foundation for the idea of logarithmic utility—an idea with a famous history in decision theory.[17] A number of authors, including Stearns (2000) and Sinn (2003) 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 maximization cannot be so simply recovered. This should not be surprising, given McNamara's characterization 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 $\Sigma_s \log r(s)\, p(s)$ is 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 'nine offspring or one offspring, probability .5 each'. The lottery that the individual faces may be 'realized' 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, that is, 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, there is probability .5 everyone in the population leaves nine offspring, and probability .5 everyone leaves one. In B, there is probability 1 everyone in the population has a 50 : 50 chance of leaving one 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, that is, for

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

https://doi.org/10.1086/658115 Published online by Cambridge University Press

every number of offspring, the probability of having at least that number is higher under A than B. However, on the McNamara-Robson evaluation criterion, B scores higher than A.[18] Natural selection will favor 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 maximization and thus (arguably) to irrationality.

However, this conclusion rests on an implicit assumption, namely, that an individual's utility function must depend only on his own reproductive output. Modulo this assumption, it is certainly true that no utility function exists such that an EU maximizer 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 maximization 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 Σ_s log 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), that is, the biased distribution shifts probability mass onto states where the population as a whole does poorly (has low $\bar{r}$(s)) and away from states where the population as a whole does well, relative to the true distribution. Rearranging, the criterion can be written as Σ_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, that is, 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 maximization must be restorable. The choice is not hard to find: simply let an individual's utility depend on his relative number of offspring (relative to the population average). If the individual obeys EU maximization, he will then be led to make evolutionarily optimal choices.

It may seem puzzling that the same evolutionary model can apparently

18. The evaluation is B is log [(.5) × 1 + (.5) × 8.5] = log 4.75. The evaluation of A is .5 log 9 + .5 log 1 = .5 log 9, which is less than log 4.75.

https://doi.org/10.1086/658115 Published online by Cambridge University Press

imply that EU maximization is satisfied and that it is violated. How can this be? In fact there is no contradiction here. When Robson argues that EU maximization 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 it 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 maximization.

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 his 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 behavior? Perhaps, but alternatively evolution may simply produce people who care about the welfare of their kin; this would realign 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, that is, 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 maximization. Of course, this does not show that the specific patterns of EU violation discovered by behavioral economists can be given an evolutionary rationale; whether that can be

https://doi.org/10.1086/658115 Published online by Cambridge University Press

shown is an open question requiring further work. But the foregoing considerations do suggest that it is a real possibility.

8. Conclusion. This article has explored connections, both thematic and formal, between 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, for example, prospect theory, rank-dependent utility, and so on, 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, Maurice. 1952. "Le Comportement de L'homme Rationnel devant le Risque, Critique des Postulats et Axioms de L'école Américaine." *Econometrica* 21:503–46.

Allais, Maurice, and Ole Hagen, eds. 1979. *Expected Utility Hypothesis and the Allais Paradox*. Dordrecht: Reidel.

Arrow, Kenneth. 1951. "Alternative Approaches to the Theory of Choice in Risk-Taking Situations." *Econometrica* 19:404–37.

Bernoulli, Daniel. 1738. "Specimen Theoriae Novae do Mensura Sortis" [Exposition of a new theory on the measurement of risk]. *Econometrica* 22:23–36.

Cooper, William S. 2001. *The Evolution of Reason*. Cambridge: Cambridge University Press.

Curry, Philip. 2001. "Decision Making under Uncertainty and the Evolution of Interdependent Preferences." *Journal of Economic Theory* 98:357–69.

Frank, Steven A., and Montgomery Slatkin. 1990. "Evolution in Variable Environments." *American Naturalist* 136:244–60.

Gillespie, John. 1977. "Natural Selection for Variances in Offspring Number: A New Evolutionary Principle." *American Naturalist* 111:1010–14.

Gintis, Herbert. 2009. *The Bounds of Reason*. Princeton, NJ: 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.

Hansson, Bengt. 1988. "Risk Aversion as a Problem of Conjoint Measurement." In *Decision, Probability and Utility*, ed. Peter Gärdenfors and Nils-Eric Sahlin, 136–58. Cambridge: Cambridge University Press.

Harless, David W., and Colin F. Camerer. 1994. "The Predictive Utility of Generalized Expected Utility Theories." *Econometrica* 62:1251–89.

Hey, John D., and Chris Orme. 1994. "Investigating Generalisations of Expected Utility Theory using Experimental Data." *Econometrica* 62:1291–1326.

Houston, Alasdair I., and John McNamara. 1999. *Models of Adaptive Behaviour*. Cambridge: Cambridge University Press.

Houston, Alasdair I., John McNamara, and Mark Steer. 2007. "Do We Expect Natural Selection to Produce Rational Behaviour?" *Philosophical Transactions of the Royal Society of London* B 362:1531–43.

Kagel, John H., Raymond Battalio, and Leonard Green. 1995. *Economic Choice Theory:*

https://doi.org/10.1086/658115 Published online by Cambridge University Press

An Experimental Analysis of Animal Behaviour. Cambridge: Cambridge University Press.
Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47:263–91.
Lewontin, Richard, and Daniel Cohen. 1969. "On Population Growth in a Randomly Varying Environment." *Proceedings of the National Academy of Sciences of the USA* 62:1056–60.
Machina, Mark. 1982. "'Expected Utility' Analysis without the Independence Axiom." *Econometrica* 50:277–323.
———. 2008. "Non-expected Utility Theory." In *The New Palgrave Dictionary of Economics*, ed. Steven N. Durlauf and Lawrence E. Blume. 2nd ed. Basingstoke: Palgrave Macmillan.
McNamara, John N. 1995. "Implicit Frequency-Dependence and Kin Selection in Fluctuating Environments." *Evolutionary Ecology* 9:185–203.
Okasha, Samir. 2007. "Rational Choice, Risk Aversion and Evolution." *Journal of Philosophy* 104:217–35.
———. 2009. "Individuals, Groups, Fitness and Utility: Levels of Selection Meets Social Choice Theory." *Biology and Philosophy* 24:561–84.
Orr, H. Allen. 2007. "Absolute Fitness, Relative Fitness, and Utility." *Evolution* 61:2997–3000.
Quiggin, John. 1982. "A Theory of Anticipated Utility." *Journal of Economic Behavior and Organisation* 3:323–43.
———. 1993. *Generalized Expected Utility Theory: The Rank-Dependent Model*. Amsterdam: Kluwer.
Robson, Arthur. 1996. "A Biological Basis for Expected and Non-expected Utility." *Journal of Economic Theory* 68:397–424.
Savage, Leonard J. 1954. *The Foundations of Statistics*. New York: Wiley.
Seger, Jon, and H. Jane Brockmann. 1987. "What Is Bet-Hedging?" *Oxford Surveys in Evolutionary Biology* 4:182–211.
Sinn, Hans-Werner. 2003. "Weber's Law and the Biological Evolution of Risk Preferences." *Geneva Papers on Risk and Insurance Theory* 28:87–100.
Stearns, Stephen C. 2000. "Daniel Bernoulli (1738): Evolution and Economics under Risk." *Journals of Biosciences* 25:221–28.
von Neumann, John, and Oskar Morgenstern. 1944. *Theory of Games and Economic Behaviour*. Princeton, NJ: Princeton University Press.
Wakker, Peter P. 1994. "Separating Marginal Utility and Probabilistic Risk Aversion." *Theory and Decision* 36:1–44.
Wilbur, Henry M., and Volker H. W. Rudolf. 2006. "Life-History Evolution in Uncertain Environments: Bet-Hedging in Time." *American Naturalist* 168:398–411.
Yaari, Menachem E. 1987. "The Dual Theory of Choice under Risk." *Econometrica* 55:95–115.

https://doi.org/10.1086/658115 Published online by Cambridge University Press