Risky Business: Evolutionary Theory and Human Attitudes towards Risk

A Reply to Okasha[*]

Armin W. Schulz

University of Wisconsin-Madison

In a recent article, Samir Okasha has given a novel and interesting explanation for the pervasiveness of risk-aversion in decision making – one based on evolutionary considerations.[1] This explanation furthermore has considerable intrinsic importance, as it shows how useful evolutionary approaches can be for the investigation of many long-standing issues in philosophy and in the (natural and social) sciences.

However, despite these successes, a closer look at the problem Okasha is trying to solve makes clear that, overall, he has not managed to adequately account for all the issues that need to be addressed. As this paper aims to show, by disregarding the fact that *risk-loving* behaviour, too, is widespread, Okasha's account ultimately fails to explain the nature of our attitudes towards risk. Further, while a non-adaptationist extension of his account may be available, it requires much further work to be made plausible.

The paper is structured as follows. Section I briefly revisits the argument of Okasha's paper. Section II raises an objection to Okasha's account and briefly suggests a possible reply to this objection. I conclude in section III.

I. Okasha's Argument for the Evolution of Risk-Aversion

The problem that Okasha takes to require a solution is the following. It is an empirical fact that people frequently behave in ways that suggest *risk-aversion*: they are not only concerned with the *expected value* of the outcomes of the options under consideration, but also consider the *variance* associated with these outcomes.[2] What requires explanation is the origin of this behaviour – why are people often willing to forgo potential, but uncertain, gains?[3]

In order to answer this question, it is best to begin by making explicit what it means to exemplify a specific attitude towards risk.[4] Consider a lottery L and a random variable X that ranges over the possible outcomes of that lottery. Then:

The agent is *risk-neutral* if and only if she is indifferent between receiving the expected payout of L with certainty and receiving L itself, i.e.: $E[U(X)] = U(E[X])$

The agent is *risk-averse* if and only if she (strictly) prefers receiving the expected payout of L with certainty to receiving L itself, i.e.: $E[U(X)] < U(E[X])$.

The agent is *risk-loving* if and only if she (strictly) prefers receiving L itself to receiving the expected payout of L with certainty, i.e.: $E[U(X)] > U(E[X])$.

Furthermore, distinguish three types of explanations of our attitudes towards risk: *purely psychological* ones (which appeal only to facts about our minds now), *social scientific* ones (which give particular prominence to the constraints put on us by our social environment), and *evolutionary* ones (which appeal to the circumstances under which our minds evolved).[5] Since

Okasha's interest is mainly in the last category, he groups the first two types of explanations together; and since nothing much hangs on this distinction in the present context either, I shall do the same.[6]

As Okasha further notes, a pivotal position in debates about the nature of our attitudes towards risk is held by a psychological / social scientific account that is based on the assumption that people have *declining marginal utility of money*. Making this assumption entails that the associated utility function is *concave*, and that weighted averages between two sums of money are associated with higher levels of utility than weighted averages of the utilities associated with the two sums – i.e. $E[U(X)] < U(E[X])$. Graphically, this can expressed as follows:

[figure 1]

However, as Okasha also makes clear, this type of approach is far from being universally accepted – in particular, it has been criticised for appearing to make numerous empirically false predictions.[7] For this reason, he is led to pursue an alternative explanation of risk aversion based on *evolutionary considerations*.

Specifically, he argues that, if the populations that hominids evolved in were finite or contained reproductive dependencies (so that the reproductive success of an organism in a population is correlated with that of the others), risk-averse behaviour can be fitter than risk-neutral behaviour.[8] To make this argument, Okasha appeals to the fact that *fitness* is not always captured well by the expected *absolute* number of offspring an organism has. Instead, it at times needs to be seen as its expected *relative number* of offspring.[9]

This can be illustrated with a case of the following sort: assume there are two types of organisms, A and B. Organisms of type A consistently double their numbers in every new generation, while organisms of type B either quadruple (with probability 1/2) or fail to reproduce at all (also with probability 1/2). Finally, assume that A and B start out being equally numerous, and that there are n individuals of each type. Then it holds that:

Expected *number* of offspring of type B = 1/2 (4n) + 1/2 (0n) = 2n = Expected *number* of offspring of type A

The expected *ratio* of A's to B's = 1/2 (n/6n) + 1/2 (n/n) = 1/12 + 1/2. = 7/12 > 1/2

What this shows is that, over time, organisms of type A are expected to increase in frequency *relative to* the number of organisms of type B, since the latter's chances of reproduction are more variable than the former's. This therefore makes clear that, from an evolutionary point of view, there is selection *against* such variability in reproductive success. In this manner, Okasha also comes to endorse Gillespie's measure of fitness, according to which it holds that (for an organism of type X and total population size N)

Fitness of organism type X ≈ E(number of offspring of type X) – 2[Var(number of offspring of type X)/N][10]

For what follows below, it is important to be clear about what this equation expresses: according to the Gillespie/Okasha measure, the fitness of an organism is its expected number of offspring, *corrected for* the riskiness of obtaining this outcome. That is, an organism's fitness is *lower* the *greater* the variance of the strategies involved: increases in the variance of the strategy *always* impact the fitness of the organism *negatively*.

In this way, Okasha seeks to explain modern attitudes towards risk as the product of selection pressures in the past: those of our ancestors that were more circumspect in their actions did better than those that took more risks. This, he submits, gives a better account of the phenomena than the traditional psychological and social scientific theories on offer. However, it turns out that the issues here are significantly more complex than he lets on.

II. Attitudes Towards Risk: Another Look

In order to see why Okasha's evolutionary argument does not give a completely successful account of human attitudes towards risk, it is necessary to start by reconsidering his conceptualisation of the *problem* to be solved. The real issue concerning human attitudes towards risk is not that people are *risk-averse*, but that they are risk-averse *and* risk-loving. More specifically, it is the fact that many people are both risk-loving and risk-averse *at the same time*. Explaining this *non-monotonicity* in people's attitudes towards risk is the real heart of the problem – making sense only of their risk-*aversion* merely covers half of the ground to be traversed.[11]

A useful illustration of why *this* is the problem to be solved is the following.[12] Many of people across the world take out some form of voluntary *insurance* if it is offered to them (and they are not financially constrained in various ways); however, equally, many people around the world *gamble* in some form or another. Moreover, often the very same people do both: they take out insurance and also play the lottery. Note, however, that taking out insurance only makes sense if the agent is risk-averse: in effect, it replaces a certain loss (the insurance premium) with an uncertain loss that could be either zero or a large amount.[13] On the other hand, buying a lottery-ticket only makes sense on the assumption of being risk-loving: it replaces the status-quo for a small chance of gain and a large chance of loss (i.e. the price of the ticket).[14] Thus, the fact that people frequently act in this way shows that, at least prima facie, what requires explanation is not just why people risk-averse, but why they are *risk-averse and risk-loving*.

It is important to note that this is a fundamentally different problem from the one that Okasha tries to solve: since explananda generally need not 'compose', explaining why people are risk-averse (but not why they are also risk-loving) *and* explaining why people are risk-loving (but not why they are also risk-averse) need not make for an explanation of why people are risk-averse *and* risk-loving.[15] That is, the point is not that there is something that is not explained by Okasha's account – it is rather that (at least potentially), *nothing* is. This is because, in general, if we want to explain some phenomenon N, then it will not be sufficient to start by explaining some other phenomenon M – and this is so *even if* it might appear that M is somehow 'a part of' N.

As an analogy, imagine two identical metal balls being released in a magnetic field, and that one of them accelerates downwards towards the ground, while the other one moves upwards towards one of the poles of the field. In this case, it will not do to explain the behaviour of the first ball as due to gravity, and the behaviour of the second due to the magnetic field – for what requires explanation is why the first ball behaved in one way and the second in another *in the same circumstances*. To be sure, gravity and the magnetic field will need to be cited in this explanation – but their effects cannot be separated out according to what *object* they concern. Instead, the *distribution* of the causal factors in *both* cases needs to be explained: why did gravity and electro-magnetism combine to yield one effect in one case and another in another?

In the present context, this implies that in order to explain compellingly why people are risk-averse, it may need to be taken into account that they are *also* risk-loving – in short, what requires explanation is the nature of their attitudes towards risk *in their entirety*. This makes for a fundamentally different problem than the one attacked by Okasha.

Concerning this more general problem, there are again – in principle – three types of explanations that can be distinguished: psychological ones, social scientific ones, and evolutionary ones. Again, though, separating the first two is not greatly relevant for present purposes, where distinguishing evolutionary from non-evolutionary explanations is sufficient. Equally, it is still the case that the non-evolutionary side – either in the form of a psychological or a social scientific theory – holds the central position in the literature.[16]

Among the most prominent explanations of this type are the following: (1) It may be that people are risk-loving when there is a chance of winning a significant amount, and risk-averse when it comes to losing small to medium-sized amounts (the Friedman-Savage hypothesis).[17] (2) It may be that they are risk-loving when the alternative to the gamble is a certain *loss*, but risk-averse when it is a certain *gain* (the Kahneman-Tversky hypothesis).[18] (3) There might be external constraints that make it appear *as if* the individuals are risk-averse and risk-loving, when they are in fact risk-neutral.[19] Graphically, these hypotheses yield more complex utility functions (compared to figure 1) – e.g. the famous S-shaped ones:

[figure 2]

[figure 3]

Now, despite the fact that Okasha has misconstrued the nature of the problem to be solved, he is correct in noting that it is very controversial how successful the non-evolutionary explanations are in dealing with the empirical phenomena. For example, there seem to be many cases of risk-loving behaviour that concern medium-sized losses (thus contradicting (1) above), concerning certain gains (thus contradicting (2)), and concerning cases where external constraints are absent (thus contradicting (3)).[20] Again, this fact – as well as the intrinsic interest of other explanations – might suggest looking for an evolutionary account as an alternative.

However, at this point the two scenarios depart. Unlike in Okasha's original conceptualisation of the problem, it becomes immediately clear that his adaptationist account is in no position whatsoever to say anything about the *general* problem concerning our attitudes towards risk: for if there is not only risk-aversion to contend with, but also risk-love, his adaptationist explanation cannot even get off the ground.

To see this, note that (as made clear earlier) according to the Gillespie/Okasha measure, the fitness of an organism depends negatively on the variance with which it reproduces. This though means that structurally, Okasha's argument as it stands cannot deal with risk-loving behaviour, for according to this view, choosing more risky strategies is never adaptive. It does not matter if the potential gain is huge and the potential costs are negligible, or if the payoffs involved are small – a risk-averse (i.e. low-variance) strategy is always fitter than a risk-loving (i.e. high-variance) one.

This might appear slightly counterintuitive at first: it may seem as if adaptationist explanations for risk-love ought to exist as well, since it is relatively easy to construct scenarios according to which strategies with a high variance in reproductive success are fitter than less variable ones.

In order to understand why this is misleading, one has to be careful in distinguishing 'variability' from 'riskiness': what the Gillespie/Okasha measure of fitness shows is that more variable strategies are never fitter than less variable ones *if the two have the same mean*. If the former has a higher mean than the latter, however, then it *can* be selected after all.[21] Importantly, though, what the definitions of the attitudes towards risk in section I show is that at stake here are instances of *mean-preserving* spreads of reproductive strategies – and it is these that are being selected against according to the Gillespie/Okasha measure of fitness. Furthermore, it is because people *in fact* pursue strategies of exactly this kind that Okasha's adaptationist account cannot provide an explanation of our attitudes towards risk – a fortiori, therefore, it cannot give an explanation that in some way *improves on* the non-evolutionary ones already available.

At this point, there are various paths open along which the ensuing debate could proceed. On the one hand, it may be necessary to accept the conclusion that evolutionary accounts are altogether unhelpful in this context. On the other, it may be possible to construct an extension of Okasha's account that is able to handle the issues after all. Before concluding, it is useful to briefly consider the outlines of such an extension.

This extension begins by noting that a form of Okasha's approach may be salvaged by retreating not from its appeal to the theory of evolution per se, but only from the *adaptationist* version it embodies. That is, one may respond to the above argument by claiming that there *are* useful evolutionary explanations of our attitudes towards risk – it is just that these do not appeal to the attitudes having been adaptive. Instead, these explanations take the attitudes to be the product of *other* forms of evolution: for example, they may be the upshot of entirely non-selective processes (as e.g. drift or mutation), or the by-product of the evolution of some other (psychological) trait.

Presently, it is enough to note that, while there is no reason to dismiss these sorts of non-adaptationist evolutionary explanations out of hand, a lot of work needs to be done in order for them to be plausible and detailed enough to be insightful. Specific evolutionary models must be proposed that are empirically well-grounded, and which have the evolution of risk-love and risk-aversion as their outcome. While not impossible, constructing such models is no trivial task. In short: non-adaptationist evolutionary accounts of our attitudes towards risk have the *potential* to be useful additions to the explanation of this part of our psychology; whether they can actually *fulfil* that potential, though, is yet to be determined.

III. Conclusion

I have tried to make clear that Okasha's argument, while immensely valuable for showing the usefulness of an evolutionary perspective on the investigation of human attitudes towards risk, fails to be fully compelling. In particular, I hope to have made clear that his focus on risk-aversion alone cannot do justice to the complexity of people's decisions in the face of uncertainty: since people are risk-loving as well as risk-averse, any explanation of their attitudes towards risk must take both aspects into account. Since Okasha's treatment is unable to do this, it falls short of what is required. While there is a possible reply available to this argument – in the form of a distinction between evolving and adapting – it requires significant further work to be made plausible.

Footnotes

[*] Many thanks to Elliott Sober, Samir Okasha, and the editors of this journal for useful comments on previous drafts of this paper.

[1] Samir Okasha, “Rational Choice, Risk Aversion, and Evolution,” *The Journal of Philosophy*, CIV, 5 (2007): 217-235.

[2] In fact, disregarding the variance is also theoretically dubious (as is shown e.g. by the St. Petersburg Paradox).

[3] See Okasha (*op. cit.*), pp. 217-18. Note that Okasha frames this as an issue about the plausibility of Rational Choice Theory (which assumes *risk-neutrality* and thus clashes with the empirical findings of risk aversion). However, since this raises difficult questions about the interpretation of Rational Choice Theory that do not bear on the actual argument at stake, it is better to frame the problem as a purely explanatory one, as is done here.

[4] In what follows below, it is easiest to understand ‘payoff’ as money. However, other interpretations are possible, as long as the payouts are not directly in utility-terms (the last would trivialise the issue). For similar definitions of the attitudes towards risk, see Andreu Mas-Colell, Michael Whinston, and Jerry Green, *Microeconomic Theory* (New York: Oxford University Press, 1995).

[5] It is not clear that these approaches necessarily need to conflict, but it is helpful for now to follow Okasha in assuming that they do. Section III discusses the reasonableness of this assumption further.

[6] For an example of a social scientific explanation of risk-aversion, see Olivier Mahul, “The Output Decision of a Risk-Neutral Producer under Risk of Liquidation,” *American Journal of Agricultural Economics*, 82, 1 (2000): 49-58.

[7] See e.g. Matthew Rabin, "Risk-Aversion and Expected Utility Theory," *Econometrica* 68, 5 (2000): 1281-1292, and Matthew Rabin and Richard H. Thaler, "Risk Aversion," *The Journal of Economic Perspectives*, 15, 1 (2001): 219-232. It is, however, worthwhile to note that Rabin's work is still somewhat controversial and that some of the assumptions he makes seem fairly strong. In particular, he assumes that utility is concave *throughout*, which, as will become clear below, does not appear altogether plausible. For further discussion, see also Martin M. Monti, Simon Grant, Daniel N. Osherson, "A note on concave utility functions," *Mind & Society*, 4 (2005): 85–96.

[8] In what follows, I shall largely ignore these two assumptions, as they do not play any role in my argument. For infinite populations, the benefits that risk-averse populations have over risk-neutral ones go to zero, but they never become *negative*. The same holds if it is assumed that individuals reproduce independently from one another. See also Okasha (*op. cit.*), pp. 225-226.

[9] For more on the concept of fitness in evolutionary biology, see Elliott Sober, "The Two Faces of Fitness", in Rama Singh, Diane Paul, Costas Krimbas, and John Beatty, eds., *Thinking about Evolution: Historical, Philosophical, and Political Perspectives* (New York: Cambridge University Press, 2001), pp. 309-321.

[10] See J. Gillespie, "Natural Selection for Variances in Offspring Number: A New Evolutionary Principle," American Naturalist, CXI (1977): 1010–14 and Okasha (*op. cit.*), p. 228.

[11] This non-monotonicity has been subject of much speculation in the economic literature ever since Friedman & Savage's and Kahneman & Tversky's classic pieces – see Milton Friedman and Leonard Savage, "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy*, 56, 4 (1948): 279-304, and Daniel Kahneman and Amos Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica* 47, 2 (1979): 263–291. More recent

discussions of these issues are Roger Hartley and Lisa Farrell, "Can Expected Utility Theory Explain Gambling?," *The American Economic Review* 92, 3 (2002): 613-624; Joseph G. Eisenhauer, "How Prevalent are Friedman-Savage Utility Functions?," *Briefing Notes in Economics*, 66 (2005): 1-9; Monti et al (*op. cit.*); Robert M. Wiseman and Luis R. Gomez-Mejia, "A Behavioral Agency Model of Managerial Risk Taking," *The Academy of Management Review*, 23, 1 (1998): 133-153.

[12] See also Friedman & Savages (*op. cit.*) and Hartley and Farrell (*op. cit.*).

[13] The economics of the insurance-industry are complex – the above is just an illustration for expository purposes. For a good introduction to this issue, see Georges Dionne and Scott E. Harrington (*eds*.), *Foundations of Insurance Economics - Readings in Economics and Finance* (New York: Kluwer Academic Publishers, 1992).

[14] This is so under the assumption that people do not obtain utility from playing the lottery (as opposed to just from receiving its outcomes). Relaxing this assumption makes for another potential psychological explanation of risk-love – see also below.

[15] This is also why it matters that the same people are acting in ways that suggest both risk-love and risk-aversion: this shows that these are not two disjoint phenomena concerning two different types of psychologies (say), but rather two aspects of the same psychology.

[16] Again, this strict division is somewhat artificial and more an expository device at this point of the argument. See also note 5.

[17] See Friedman & Savage (*op. cit.*) and Hartley and Farrell (*op. cit.*). The latter piece also makes clear that Rabin and Thaler's objection (*op. cit*., note 1, pp. 222-223) to this type of utility function appears overly quick.

[18] See Kahneman & Tversky (*op. cit.*)

[19] See e.g. Mahul (*op. cit.*) and Wiseman et al. (*op. cit.*).

[20] For more on this, see the literature cited in note 11. Note also that some of the difficulties here come from the fact that there is empirical support for both (1) and (2), but that the two are at least partly in conflict with one another.

[21] For example, this is what drives the model in Eddie Dekel and Suzanne Scotchmer, "On the Evolution of Attitudes towards Risk in Winner-Take-All Games", *The Journal of Economic Theory* 87 (1999): 125-143.

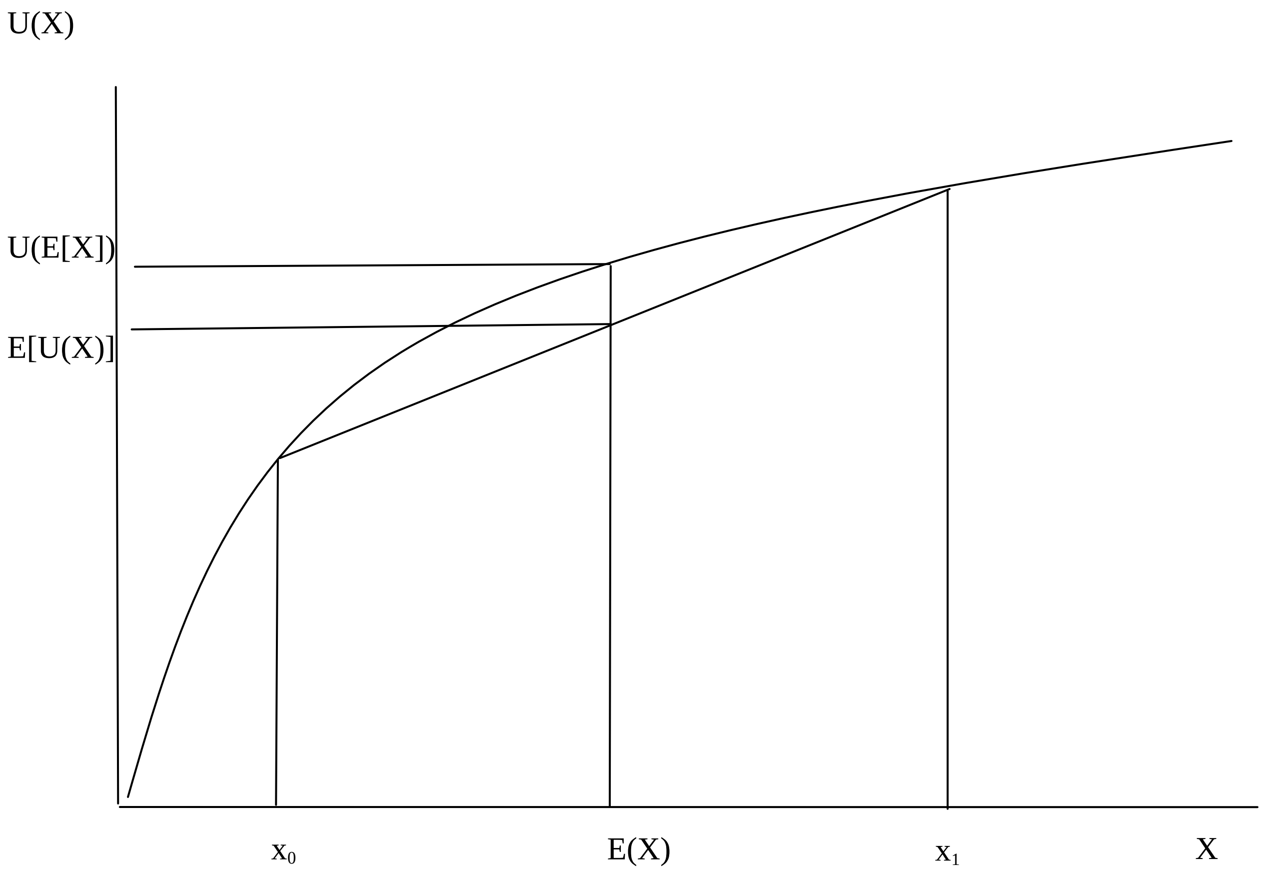

[Figure 1]

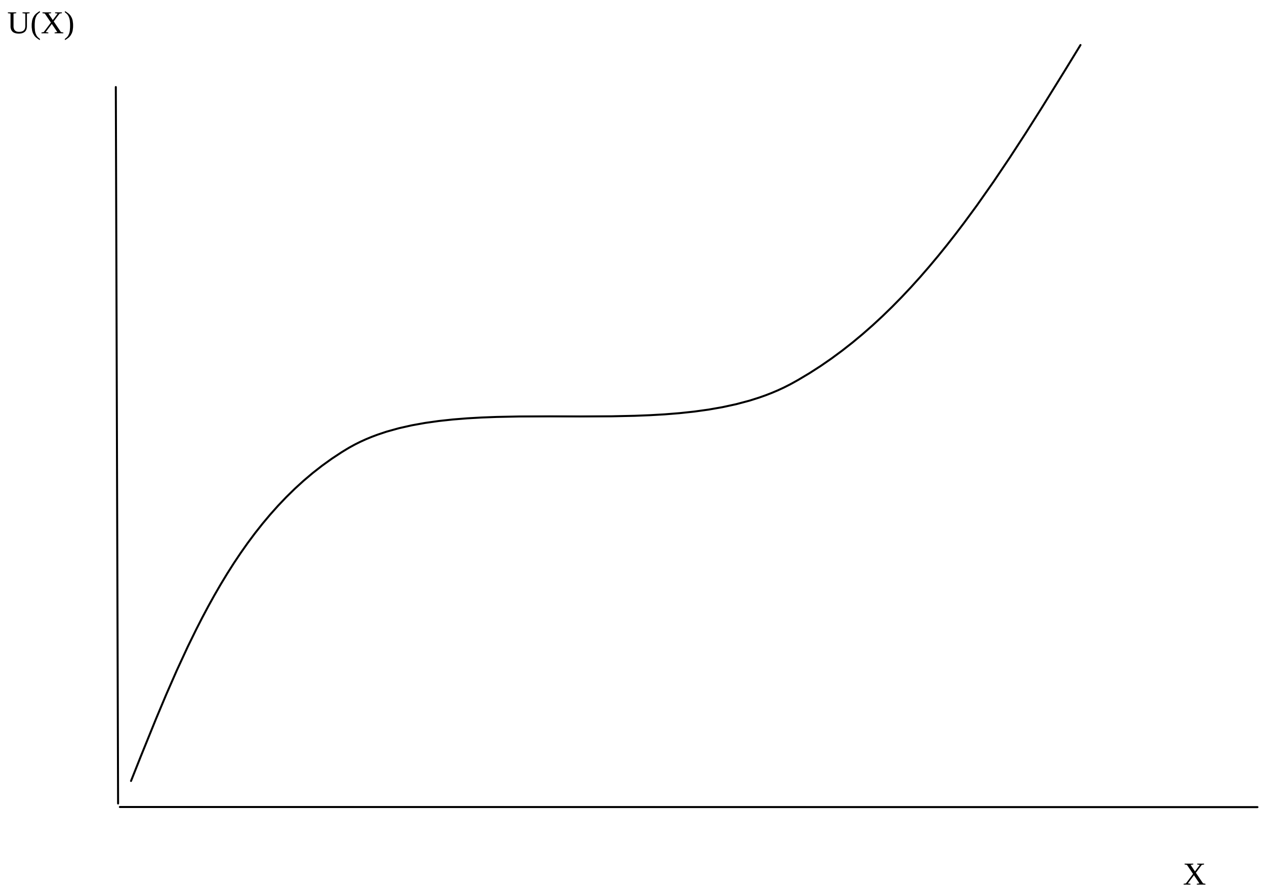

[Figure 2]

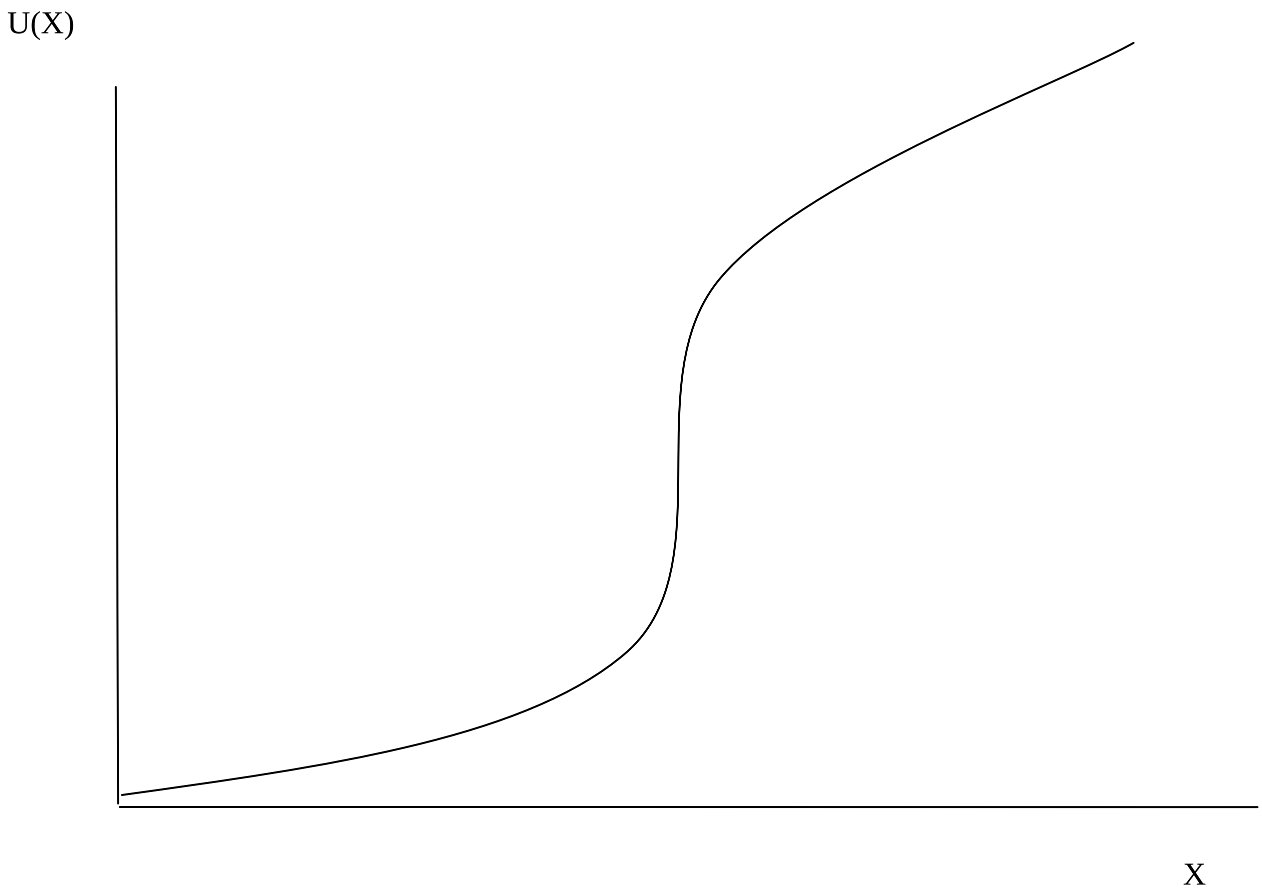

[Figure 3]