

UC Berkeley

UC Berkeley Previously Published Works

Title

Precis of Risk and Rationality

Permalink

<https://escholarship.org/uc/item/01z98751>

Journal

Philosophical Studies, 174(9)

ISSN

0031-8116

Author

Buchak, Lara

Publication Date

2017-09-01

DOI

10.1007/s11098-017-0904-7

Peer reviewed

Precis of *Risk and Rationality*

Lara Buchak

ABSTRACT

My book *Risk and Rationality* argues for a new alternative to the orthodox theory of rational decision-making. This alternative, risk-weighted expected utility maximization, holds that there are three important components involved in rational decision-making: utilities, probabilities, and risk-attitudes. This essay explains the basic outline of the theory and precisely how it differs from the orthodox theory. It also summarizes the main threads of argument in the book.

The orthodox theory of rational decision-making is known as expected utility maximization. Roughly, expected utility maximization holds that there are two important components to decision-making: how much an individual values the various outcomes he might obtain (his utilities) and how likely he thinks a given act is to realize these outcomes (probabilities). The value of an act is its expected utility: a weighted average of utility values, each utility value weighted by the probability that the act realizes it. A rational decision-maker will therefore prefer the act with the highest expected utility. In *Risk and Rationality*, I developed an alternative theory of rational decision-making, in the tradition of “rank-dependent” utility theory.¹ The key claim of this alternative, *risk-weighted expected utility maximization*, is that there are *three* important components in decision-making: an individual’s utilities, his probabilities, and his *attitude towards risk*.

To understand the basic outline of the theory and how it differs from expected utility maximization, consider a simple example of a decision between two acts. An act can be thought of as a function from states of the world to outcomes, where outcomes are described to include everything that the individual cares about. Consider an individual who faces a decision about whether to work at a risky start-up or instead take a dull but stable job; and let us assume for the sake of example that we know the probabilities

¹ Buchak (2013). Other rank-dependent theories include anticipated utility (Quiggin 1982), dual theory (Yaari 1987), Choquet expected utility (Schmeidler 1989, Gilboa 1987), and cumulative prospect theory (Kahneman and Tversky 1979, Tversky and Kahneman 1992). For a survey, see Wakker (2010).

he assigns to states and the utilities he assigns to outcomes.² If the start-up is amazingly successful (probability 0.01) then he will be very wealthy, feel a strong sense of personal accomplishment, have a short and pleasant workday, and be able to travel the world and enjoy the finer things in life (utility 7). If it is very successful (probability 0.29) then he will be fairly wealthy, feel accomplished, and have an enjoyable workday (utility 6). If it is moderately successful (probably 0.5), then he will have enough money to pay the bills, but the hours will be long and boring (utility 4). If it fails (probably 0.2), then he will have to leave and get an unpleasant, demanding job where he is merely scraping by (utility 1). On the other hand, if he takes the stable job, he will have enough money to pay the bills, but with long and boring hours (utility 4). We can represent these two options graphically, with the height of each bar representing the utility of each outcome, and the width of each bar representing the probability of each outcome:

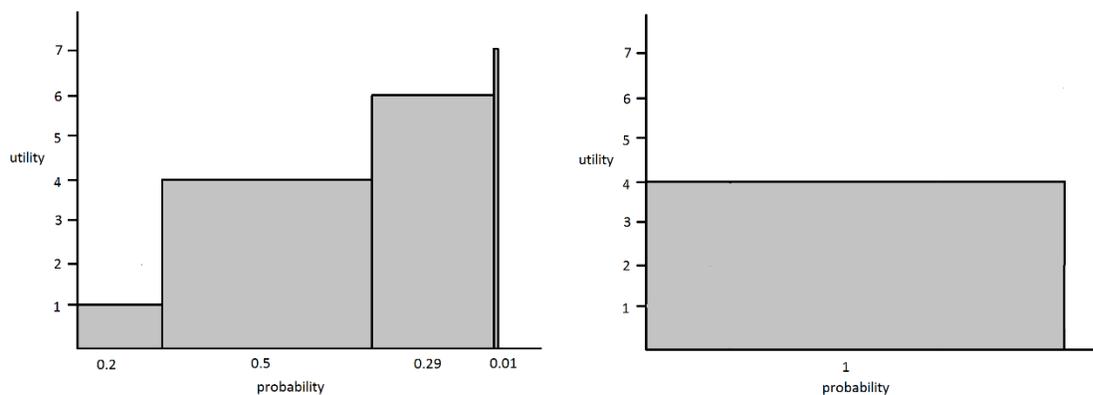


Figure 1: Choice between Working at a Risky Start-up and a Stable Job

Again, according to expected utility (EU) maximization, the utility value of each option is a weighted average: the utility value of each outcome is weighted by the probability of the states that realizes it, and the result is summed. So we have:

$$EU(\text{Start-up}) = (0.2)(1) + (0.5)(4) + (0.29)(6) + (0.01)(7) = 4.01$$

$$EU(\text{Stable job}) = (1)(4) = 4$$

Thus, the value of each option is the area under the curve in the above graph. Furthermore, the decision-maker should pick the option with the highest EU (in this case, he should go for the start-up).

² Example and graphs taken from Buchak (2017).

One way to think about what EU-maximization says (paying attention just to the graph on the left) is that the relevant considerations in evaluating an act are the four utility values of the outcomes the individual might receive, and each of these considerations gets a weight equal to the probability of the states that realize it. But we can instead hold that the relevant considerations are the *incremental utility benefits* the individual might receive: he will at least get utility 1; in 80% of the states, he will do better than this by at least utility 3; in 30% of the states, he will do better than this by at least utility 2; and in 1% of the states, he will do better than this by utility 1.

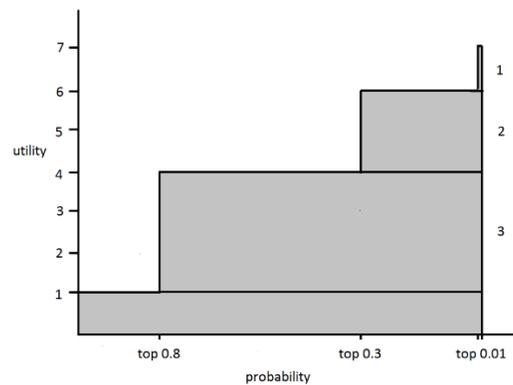


Figure 2: Expected Utility, Reconceptualized

In this reconceptualized graph, the height of each rectangle represents the difference between two adjacent utility levels (benefits that one might receive), and the width of each rectangle represents the probability of attaining *at least* the relevant utility level (the probability of receiving those benefits). As before, the area under the curve is the EU of the gamble. Thus, EU-maximization holds that the weight of each consideration of the form *I might obtain benefits of a certain size (in addition to whatever other benefits I obtain)* is the probability of obtaining those benefits.

So, for example, the value of working at the start-up and the stable job can be calculated:

$$\text{EU(Start-up)} = (1)(1) + (0.8)(3) + (0.3)(2) + (0.01)(1) = 4.01$$

$$\text{EU(Stable job)} = (1)(4) = 4$$

EU-maximization thus commits an individual to weighting these considerations to a specified degree. But notice that different individuals might in fact weight these considerations differently. One individual might care more about benefits that are guaranteed—or about what happens in worse states—and so a benefit that he will get in only 80% of the states will not factor very heavily into his evaluation of the

gamble, and a benefit that he will get in only 30% of the states will factor in even less. Another individual might care more about what happens in better states, and so a benefit that he will get in only 80% of the states will factor very heavily into his decision—nearly as heavily as a guaranteed benefit.

Risk-weighted expected utility (REU) maximization accounts for this difference by assuming that each individual has a *risk-function* which “shrinks” or “stretches” the horizontal rectangles, by a factor which represents the degree to which an individual cares about what goes on in the relevant portion of states. Specifically, the risk function $r(p)$ measures the importance of the top p -portion of outcomes to an individual’s decision-making. If an individual’s risk-function is convex—meaning that as a benefit is obtained in a smaller and smaller portion of states, he gives that benefit proportionally less and less weight—then we say he is *risk-avoidant*. If an individual’s risk-function is concave—meaning that as a benefit is obtained in a smaller and smaller portion of states, he gives that benefit proportionally more and more weight—then we say he is *risk-inclined*. If an individual’s risk-function is linear—meaning that he weights benefits exactly in proportion to the portion of states in which they are realized—then we say he is *globally neutral*; globally neutral individuals maximize expected utility.

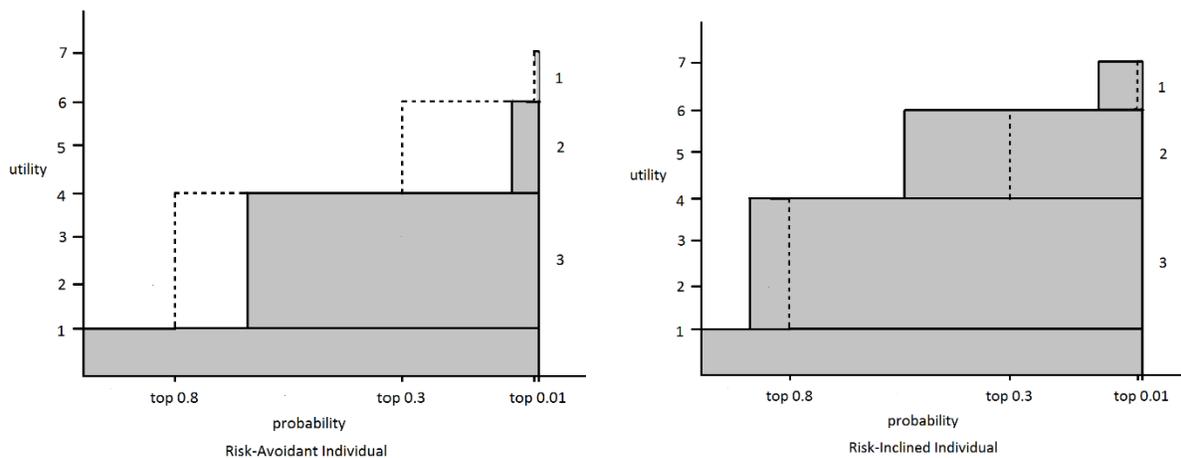


Figure 3: Risk-Weighted Expected Utility

The shaded area in the above graphs represents the value of working at the start-up, for a *risk-avoidant* individual with $r(p) = p^2$ (left) and a *risk-inclined* individual with $r(p) = p^{0.5}$ (right). We can calculate the value of working at the start-up and the stable job for each individual:

Risk-avoidant individual:

$$\text{REU}(\text{Start-up}) = (1)^2(1) + (0.8)^2(3) + (0.3)^2(2) + (0.01)^2(1) = 3.10$$

$$\text{REU}(\text{Stable job}) = (1)^2(4) = 4$$

Risk-inclined individual:

$$\text{REU}(\text{Start-up}) = (1)^{0.5}(1) + (0.8)^{0.5}(3) + (0.3)^{0.5}(2) + (0.01)^{0.5}(1) = 4.88$$

$$\text{REU}(\text{Stable job}) = (1)^{0.5}(4) = 4$$

The risk-avoidant individual will conclude that taking the stable job is preferable to working at the start-up, and the risk-inclined individual will conclude that working at the start-up is preferable.

We can note the general question for expected utility maximization (in the re-conceptualized form) and the general equation for risk-weighted expected utility maximization. Where $g = \{E_1, x_1; \dots; E_n, x_n\}$ is an ordered gamble that yields outcome x_i in event (set of states) E_i and $x_1 \leq \dots \leq x_n$:

$$EU(g) = \sum_{i=1}^n \left[\left(\sum_{j=i}^n p(E_j) \right) (u(x_i) - u(x_{i-1})) \right]$$
$$REU(g) = \sum_{i=1}^n \left[r \left(\sum_{j=i}^n p(E_j) \right) (u(x_i) - u(x_{i-1})) \right]$$

where u is a utility function, p is a probability function, and r is 'risk function' from $[0, 1]$ to $[0, 1]$, with $r(0) = 0$, $r(1) = 1$, and r non-decreasing.

That is the basic formal outline of REU theory. *Risk and Rationality* has four main sections. The first section motivates the claim that REU-maximization better captures ideal instrumental ('means-ends') reasoning than EU-maximization. One way this is done is by examining examples of preferences that EU-maximization cannot reconstruct but that REU-maximization can reconstruct. It is then argued that these preferences are prima facie rational in the sense that individuals who have them can be seen as displaying consistent values and consistently reasoning to realize them; furthermore, the reasoning underlying the preferences is well-captured by REU-maximization. Finally, it is argued that EU-maximization leaves something out, in the sense that an agent cannot reason towards realizing his values without determining his risk-attitude, even if this determination turns out to be the one recommended by EU-maximization. A key example in this section is the Allais paradox, which both commentators will describe in detail.

The second section of the book is a *representation theorem*. What representation theorems for EU-maximization show is that if an individual's preferences obey certain axioms, then he can be interpreted as maximizing EU with respect to a unique probability function and a utility function that is unique up to positive affine transformation. I prove a representation theorem for risk-weighted expected utility maximization, which shows that if an individual's preferences obey certain strictly weaker axioms, then he can be interpreted as maximizing REU with respect to a unique probability function, a utility function that is unique up to positive affine transformation, and *a unique risk function*. This allows us to do a number of things. First, we needn't adhere to a realist interpretation of the three components, but can instead see them as arising from, or being restatements of, preferences themselves. Second, even if we hold that the components are psychologically real, we have a way to discover them from (more easily observable) preferences. Finally, we can locate the debate between REU-maximization and EU-maximization at the level of which set of axioms are the correct constraints on rational preferences, rather than at the level of which formalized aggregation method is correct. The third section of the book is an argument that REU-maximization is not just EU-maximization with more complex outcomes. Neither of these sections will be discussed at length by the commentators.

The final section of the book argues that REU-maximizers are rational. Since EU-maximization is a special case of REU-maximization, the goal is primarily to rebut the presupposition that only EU-maximization is rational—and, in particular, to defend other REU-maximizers against the charge that they fall afoul of some criterion of rationality. Singled out for discussion in this symposium is the charge that REU-maximizers have preferences that make them subject to sure loss or willing to forgo sure gain. I argue that none of the arguments in favor of the unique rationality of EU-maximization is successful, and thus that REU theory correctly describes rational agents.

References

- Buchak, L. (2013), *Risk and Rationality* (Oxford: Oxford University Press).
- Buchak, L. (2017), 'Taking Risks behind the Veil of Ignorance,' *Ethics* 127: 610–644.
- Gilboa, I. (1987), 'Expected Utility with Purely Subjective Non-Additive Probabilities,' *Journal of Mathematical Economics* 16: 65-88
- Kahneman, D. and A. Tversky (1979), 'Prospect Theory: An Analysis of Decision under Risk,' *Econometrica* 47: 263-291.
- Quiggin, J. (1982), 'A Theory of Anticipated Utility,' *Journal of Economic Behavior and Organization* 3: 323-343.

Schmeidler, D. (1989), 'Subjective Probability and Expected Utility without Additivity,' *Econometrica* 57.3: 571-587.

Tversky, A. and D. Kahneman (1992), 'Advances in Prospect Theory: Cumulative Representation of Uncertainty,' *Journal of Risk and Uncertainty* 5: 297-323.

Wakker, P. (2010), *Prospect Theory for Risk and Ambiguity* (Cambridge: Cambridge University Press).

Yaari, M.E. (1987), 'The Dual Theory of Choice under Risk,,' *Econometrica* 55.1: 95-115.