

debate about the foundations of the expected utility theory that the book tackles in a stimulating *and* very constructive way. Its reading will be valuable to every scholar – philosopher or economist – interested by decision theory.

**Mikaël Cozic**

***Paris-Est Créteil Val-de-Marne  
and Institut d'Histoire et de Philosophie des Sciences et des Techniques  
(ENS Ulm, Paris 1 and CNRS)***

#### REFERENCES

- Anscombe, F. J. and R. J. Aumann 1963. A definition of subjective probability. *Annals of Mathematical Statistics* 34: 199–205.
- DeGroot, M. 1970. *Optimal Statistical Decisions*. McGraw-Hill.
- Jeffrey, R. 1965/1983. *The Logic of Decision*, 2nd edn. Chicago: University of Chicago Press.
- Joyce, J. M. 1999. *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Peterson, M. 2004. From outcomes to acts: a non-standard axiomatization of the expected utility principle. *Journal of Philosophical Logic* 33: 361–378.
- Savage, L. J. 1954/1972. *The Foundations of Statistics*. New York: Wiley. [Second edition (1972) New York: Dover Publications.]

doi:10.1017/S0266267110000489

*Decision Theory and Rationality*, José Luis Bermúdez. Oxford University Press, 2009. 189 pages.

According to decision theory, as standardly conceived, the rational agent chooses from the acts that are open to her an act that maximizes expected utilities relative to her probabilities and utilities, where an act's expected utility is defined as the probability-weighted average of the utilities of the various possible outcomes of the act. In the most common conception of this theory, its purpose is to help us make the right decisions; it guides our actions. But it can also have a role in evaluating other people's choices – what one might call its 'normative role'. For instance, we may accuse someone of irrational behaviour because we have reason to believe that she chose in a way that failed to maximize expected utility relative to her probabilities and utilities. Furthermore, decision theory may play a role in predicting and explaining people's behaviour, as when we predict that Sarah will come to the party rather than stay home because only the former choice maximizes expected utility relative to what we believe to be her probabilities and utilities.

In his book, Bermúdez argues for two claims: first, that decision theory cannot play all three of the aforementioned roles (guidance, evaluation, prediction/explanation of actions), and second, that decision

theory cannot play any of those roles if it does not play all three of them. Jointly, these claims are meant to present us with a puzzle: there can only be one kind of decision theory – namely, one that plays the three different roles – and that kind of decision theory cannot exist. But then what kind of theory is the theory we customarily refer to as ‘decision theory’?

Bermúdez makes no attempt to solve the puzzle. If I am right, however, no solution is needed, for the puzzle is unreal. More specifically, I think that the first of Bermúdez’s claims may well be right, but that the second is false; at a minimum, Bermúdez’s arguments for the latter are quite unconvincing. Note that while the second of Bermúdez’s claims is quite bold and would, if true, have a broad impact on theorizing about decision theory, nothing much seems to follow from the first one; at least I know of no author who has defended, or is committed to, a contrary position.

Bermúdez’s argument for the first claim starts from the observation that decision theory must be more than the slogan that the rational agent maximizes expected utility. Most obviously, its central theoretical terms – probability and utility – stand in need of interpretation. Moreover, even if, in the process of deliberation, we take the set-up of the decision problem as given, in assessing someone else’s decision our evaluation of the set-up may partly determine our verdict as to whether or not the person acted rationally. In this connection, Bermúdez points to well-known evidence showing that people may assign different utilities to one and the same outcome depending on how that outcome is presented to them – whether, for instance, the results of a treatment option are specified in terms of five-year survival rates or in terms of the percentage of patients who have died five years after treatment. In some cases, this may be excusable – say, if it cannot be reasonably expected that someone can see through the presentations of the outcomes and realize that she is really dealing with one and the same outcome – but in other cases an agent is intuitively reproachable for not having identified two descriptions as specifying the same outcome. If it is to serve its normative role, decision theory must offer some story about the individuation of outcomes. And there is more still. As it stands, decision theory completely disregards the temporal aspects of decision making. However, we have clear intuitions about the rationality (or otherwise) of people’s choices in light of earlier choices they have made or plans they have adopted. A full-fledged decision theory should do justice to such intuitions.

Bermúdez’s claim, then, is that there is no way to satisfy all these desiderata – to provide an interpretation of the key theoretical notions, to extend decision theory as we currently know it with some story about how to individuate outcomes, and to incorporate in it the temporal dimension of decision making – which will allow the resulting theory to serve the three purposes mentioned in the first paragraph of this review.

One reason Bermúdez offers in support of this claim is that in its action-guiding role, decision theory must appeal to strictly subjective factors: a person's choices can only be informed by what is accessible to her. For the very same reason – Bermúdez claims – we will also want to invoke only such factors in predicting what a person will choose, or in explaining her choices. By contrast, to answer the question of what a person ought to have chosen, we may well want to appeal to things that were beyond the person's ken. Or we may want to label the utilities a person assigns to a set of outcomes as irrational, because (say) they fail to reflect what is good for her in some objective sense. Similarly in normatively assessing a person's choice, we may criticize how she individuated the outcomes of the acts open to her, or how she reflected (or failed to reflect) upon them. But from an action-guiding perspective, all that counts is how the person *actually* individuates and understands the outcomes. And again, according to Bermúdez, such factors are only relevant when it comes to predicting or explaining the person's choice. Finally, Bermúdez considers all the popular ways of extending decision theory so as to incorporate the temporal dimension of decision making, and he argues that on none of these extensions can decision theory play all three roles. In a nutshell, the point is that in deliberation we are solely guided by our current utilities and probabilities – even though these may be assigned to propositions about future or past preferences – while in assessing, predicting, or explaining a person's choices, we will also want to attend to how, for instance, her utilities changed over time.

There is much that I find puzzling in Bermúdez's defence of his first claim. For instance, with respect to Bermúdez's argument that we will want to supplement decision theory with some story about how to individuate outcomes, it struck me that the issue of idealization is hardly ever brought up in the book. To be sure, the empirical evidence showing that people tend to evaluate outcomes differently depending on how these outcomes are presented to them is undeniable. How relevant this evidence is to decision theory is another matter. Note that Bermúdez could have made a point about probabilities very similar to the one he makes about individuating outcomes. As probability theory is normally set up, interpreting its axioms as principles of doxastic rationality – as we do in the epistemology underlying decision theory – requires that we assign probability 1 to every logical truth, whether or not it is recognizable as such. What is more, the theory (thus understood) demands that we assign the same probability to necessarily equivalent sentences. So, for instance, we must assign the same probability to

Clark Kent is Superman

that we assign to

Clark Kent is Clark Kent.

For, assuming that 'Superman' is a proper name (at least we can pretend it to be), these sentences are necessarily equivalent – they express the same proposition – and such sentences must receive the same probability. But we need no experimental evidence to be convinced that people – even people we deem pretheoretically rational – frequently violate this requirement, or the requirement to assign probability 1 to all logical truths. It is thus generally recognized that imposing the axioms of probability theory as norms of doxastic rationality involves a great deal of idealization. There is nothing wrong with that a priori: in science, idealization is the rule rather than an exception. Nevertheless, if we are willing to tolerate the said idealization with respect to probabilities, why not tolerate a similar idealization with respect to how people understand outcomes? Surely, logically omniscient agents are able to see through the guises in which outcomes are presented to them.

On the other hand, if Bermúdez is bothered by the fact that decision theory is idealized in various ways, then I do not understand why he does not pay more attention to the literature aimed at 'de-idealizing' (or 'concretizing') decision theory. Most notably, Weirich's carefully crafted work on decision making of 'nonideal agents in nonideal circumstances' (as the subtitle of Weirich 2004 reads), which deals with the question of how to relax the various idealizations inherent in decision theory, is referred to once (in a footnote), but a discussion of this work is missing in the book.

What further puzzled me about Bermúdez' case for his first claim is his implicit assumption that it must be exactly the same theory that is used for, for instance, action-guidance and normative assessment. The following strikes me as a *prima facie* plausible alternative story: decision theory in its current form is, in the first instance, an action-guiding theory. But we also invoke it to assess other people's decisions: when we wonder whether a person chose rationally, we consider whether or not she maximized expected utility relative to her probabilities and utilities. If the answer to this question is positive, then that is still not enough to conclude that the person acted rationally. We may think that she did not act rationally if we have reason to believe that the way she set up the decision problem was somehow faulty or that her degrees of belief were, although coherent (and thus probabilities), crazy, or that her utilities were crazy. Only if we are justified in thinking that no faults of these kinds underlie the person's decision may we conclude that she acted rationally. Note that in this picture, it is not one theory that serves two purposes. Rather, there is one theory which serves one purpose (the action-guiding purpose) and another broader theory which encompasses the first, that serves the other, normative purpose. The broader theory may well appeal

to factors that are not accessible to the agent, but may be compatible with the claim that the action-guiding theory need not appeal to such factors. Perhaps there are reasons to dismiss this picture as inadequate, but it would have deserved some discussion in the book.

Still, for all I know, Bermúdez is right that there is no one decision theory – equipped with an interpretation of the notions of probability and utility and motivated stories about the individuation of outcomes as well as about rationality over time – that serves the purposes of action-guidance, normative assessment, and explanation and prediction of behaviour. That this would be bad, as Bermúdez claims, is something I find hard to believe – and something Bermúdez's arguments fail to show.

Here, I want to focus on Bermúdez's argument that if decision theory serves the action-guiding purpose, then it must also serve the purposes of prediction and explanation. The argument involves what Bermúdez terms the 'belief-desire law', that is, the thesis 'that people generally act in ways that seem to them best to satisfy their desires in the light of their beliefs' (p. 165). The crux of the argument, then, is this:

If the belief-desire law is, broadly speaking, true ... and if the expected utility principle is really a regimentation of the belief-desire law, then it follows that agents must, broadly speaking, be maximizers of expected utility. But this, of course, entails that decision theory must be explanatory and predictive – that it should serve not just as a tool for deliberation, but also as a tool for explaining why people choose the way they do and predicting how they will choose. (p. 165f)

The argument is clearly a non-sequitur. Suppose that people are indeed expected utility maximizers. How does it follow from this that decision theory will help us in predicting people's choices? If we have no clue about someone's probabilities and utilities, then knowing that she is an expected utility maximizer will not help us predict what she will do. Similarly, how does it follow, making no assumptions about explanation, that decision theory will serve our explanatory purposes? If, as some hold, all explanation of behaviour must refer to brain states and processes – even if deliberating agents do not think about such states and processes – then it is obviously consistent to claim that decision theory is an action-guiding theory while turning to some other (e.g. neurophysiological) theory for explaining, and perhaps also predicting, people's behaviour.

While, to my mind, Bermúdez has failed to show that decision theory faces a puzzle, the book contains much useful discussion of problematic aspects of decision theory as we know it. For instance, how we are to interpret the notion of utility is still an open question. And we are still awaiting a convincing theory about the rationality (or otherwise) of changes of utility assignments. The book further contains valuable discussion of relevant literature by Broome, Pettit, Schick and others. Also,

Chapter 1 offers a clear and concise introduction to decision theory. For these reasons, the book is well worth reading, even though Bermúdez does not substantiate the book's grand claim.

**Igor Douven**

***University of Groningen***

#### REFERENCE

Weirich, P. 2004. *Realistic Decision Theory: Rules for Nonideal Agents in Nonideal Circumstances*. Oxford: Oxford University Press.

doi:10.1017/S0266267110000490

*The Social Epistemology of Experimental Economics*, Ana Cordeiro dos Santos. Routledge 2010. xi + 210 pages.

Experimental Economics is a discipline that has reached a critical point in its development. It has reached maturity in the sense that it is widely accepted within the economics discipline and its results inform large areas of economic analysis. However, it is still unclear what the results of given experiments actually mean and how far we can trust them. There has been a substantial amount of methodological analysis but much of it has been defensive, aimed at justifying experiments to economists. Furthermore, much of the extant methodological analysis is philosophically outdated since it relies on a theory-testing paradigm that has little relation to how economists (and other scientists) do experiments and how results are established.

For these reasons Ana Cordeiro dos Santos' book is a welcome newcomer to the debate. It uses current ideas in the philosophy of experiment and tries to extend them to account for the social aspects of experimental debate. While this approach has its problems, I believe that most of the fundamental points in the book are basically correct and any methodological debate should take account of its arguments.

According to Santos, experimental facts are generated through two general processes: a material process and a social process. The material process is the confrontation of the experimenter's ideas with evidence generated through experimental procedures. This is actually a complex three-way process involving the 'phenomenal model' – the experimenter's picture of how the world works, the 'instrumental model' – the experimenter's view of how the experimental apparatus performs and the 'material procedure' – the actual, physical experiment and its results. The latter includes the experimental subjects as well as the instruments used such as computers and instructions to the subjects.