

## Book Review

*Transformative Experience*, by L. A. Paul. Oxford: Oxford University Press, 2014. Pp. 208.

My driving test is in two weeks' time. I am trying to decide whether or not I should practise. How should I make this decision rationally? The standard account runs as follows. For each of the two possible actions I might perform – namely, practising or not practising – I should begin by calculating the subjective expected utility of that action from my perspective. To do this, I consider each of the possible outcomes of the action. Let's keep it simple by assuming that there are just two – I pass or I fail. Then, I go through each of the outcomes, and take the subjective utility that I assign to that outcome occurring when I have performed the action in question, and I weight that utility by the subjective probability that I assign to that outcome occurring on the supposition that I perform that action; and I add up these weighted utilities. Thus, for instance, the subjective expected utility of practising is the sum of my utility for practising and passing, weighted by how likely I think it is that I would pass if I were to practise, and my utility for practising and failing, weighted by how likely I think it is that I would fail if I were to practise. Finally, I compare the subjective expected utilities of the various possible actions I might take. Rationality obliges me to pick an action with maximal subjective expected utility. This is *expected utility theory* (Briggs 2015; Steele and Stefánsson 2015).

Now, note two presuppositions of this standard account of rational decision-making.

- First presupposition: In order for an agent to follow this procedure, she must have access to her probabilities and to her utilities: the former measure the extent to which she believes a proposition; the latter measure the extent to which she values a situation. Thus, the procedure presupposes that I have access to how likely I think it is that I would pass if I were to practise, for instance, and how much I would value passing having practised.
- Second presupposition: In the description of this procedure, we have not specified which of her utilities the agent should weight by her probabilities in order to obtain her expected utility for a given action – is it her utilities at the time she makes the decision, or at some

other time, perhaps when the consequences of her decision will be making themselves felt? For instance, should I appeal to my current utilities, two weeks prior to my driving test, or should I appeal to the utilities I will have in two weeks' time, after the driving test? This is not specified in standard expected utility theory. So, if that account is to be well-defined, there must be a presupposition that my utilities remain unchanged over time – I value passing having practised exactly as much now as I always have and as I always will.

In her rich and thought-provoking recent book, *Transformative Experience*, L. A. Paul claims that there are decisions that we face in our lives that cannot be accommodated within the standard account of rational decision-making offered by expected utility theory because one or both of these presuppositions of that standard account are false in those cases.

Suppose I'm faced with the choice whether or not to eat a durian fruit. I've never eaten one before. How much I value the outcome in which I eat the durian is determined by a number of factors, including, for instance, the nutritional value of the fruit, the ethical implications of harvesting durians, and so on. But naturally it will be determined in large part by how the durian tastes to me – that is, it will be largely determined by *what it is like for me* to eat the durian fruit. Thus, in order to access my utility for that outcome, it seems that I must be able to access that first-personal experience. However, the thing about durians is that until you've tried them, you don't know what they will taste like for you – no description can do them justice; no testimony from those who have eaten them can teach you what you need to know. Thus, in the case of this decision, the first presupposition of standard expected utility theory is false: I do not have access to my utilities. Paul calls the experience of eating a durian fruit an *epistemically transformative experience* (2014, p. 15). Until you have had the experience, you cannot know what it is like to have the experience; eating one effects a transformation in your epistemic state that cannot be effected in any other way. Paul claims that standard expected utility theory cannot accommodate any decision in which one of the actions might give rise to an outcome that involves an epistemically transformative experience.

Suppose next that I'm a socially anxious person. I prefer not to attend parties because I abhor getting trapped in a conversation filled with awkward silences. My doctor prescribes something for this: it will not make it any less likely that I'll find myself in such conversations, but it will make me loathe those awkward silences less. Now I receive an invitation to a party that will take place in a few months, once the doctor's remedy has taken effect. When performing my subjective expected utility calculation, should I appeal to my current utilities, which reflect my hatred of awkward silences, or should I appeal to the utilities I will have at the time of the party, which will reflect my newfound indifference to lengthy pauses in conversation and the embarrassed

shuffling of feet? Thus, in the case of this decision, the second presupposition of the standard account of rational decision-making is false: my utilities can change over time. Paul calls the experience of taking the doctor's remedy a *personally transformative experience* (2014, p. 16). Having the experience transforms a crucial part of the person you are, namely, your desires and preferences as encoded in your utilities. Paul claims that standard expected utility theory cannot accommodate any decision in which one of the actions might give rise to an outcome that involves a personally transformative experience.

Paul's particular interest in her book lies in major life decisions in which both of the presuppositions of expected utility theory are false. These are decisions in which one of the possible actions between which you are choosing may give rise to an outcome that includes an experience that is both epistemically and personally transformative for you. The book expands on Paul's *Res Philosophica* paper 'What You Can't Expect When You're Expecting' (2015), which raises this problem for deciding rationally whether or not to become a parent, and this example of a major life decision that might lead to an epistemically and personally transformative experience forms a central case in Paul's book as well. Paul claims that, as with eating a durian fruit, you can only know what it's like to become a parent by becoming a parent – that is, becoming a parent is an epistemically transformative experience. Moreover, as with taking the doctor's remedy for social anxiety, becoming a parent might change your utilities – that is, becoming a parent could be a personally transformative experience. But there are many other examples of such decisions. The first chapter of the book introduces this sort of decision beautifully, using the example of the wonderfully fantastical decision whether or not to become a vampire. The second chapter introduces the standard account of rational decision-making and explains Paul's objection to it by appealing to the vampire example. The third chapter describes further decisions that, like the decision whether or not to become a vampire or to become a parent, pose similar problems for the standard account: deciding whether to move to a different country, especially one whose cultural norms are far from those of your home country; deciding to commit to a long-term romantic relationship with someone, especially someone with preferences that are very different from yours; deciding whether or not to have certain risky treatments for a serious medical condition; and deciding which career to pursue. The fourth chapter introduces Paul's proposed solution to the problem created by these examples: you should make the decision based on whether or not you would like to have the sort of new experience afforded by the transformative experience in question. Thus, I should decide whether or not to become a parent by choosing whether or not I would like to discover what it is like for me to become a parent. In the fifth chapter, Paul wraps up by engaging fully with a rich array of objections to her proposals.

Thus Paul makes two claims about such major life decisions as becoming a parent:

- (1) The standard account of rational decision-making provided by expected utility theory is inapplicable because its two presuppositions are false, and so it is not possible to use it to make the decision rationally.
- (2) There is an alternative way to make the decision rationally, namely, choosing solely on the basis of whether or not I would like to discover what it is like to have the transformative experience in question.

I'd like to say something briefly about (2), and then I'd like to focus attention on (1).

Let's begin, then, with Paul's proposed solution to the problem that she has posed to expected utility theory. Of the decision whether or not to eat the durian, Paul says:

The relevant outcomes, then, of the [reconfigured] decision to have a durian are *discovering the taste of durian* versus *avoiding the discovery of the taste of durian*, and the values attached reflect the subjective value of making (or avoiding) this discovery, not whether the experience is enjoyable or unpleasant. (2014, p. 113)

Thus, the proposed solution has two parts: first, we reconfigure the decision so that it is a decision between discovering what it's like to eat durian and avoiding discovering what it's like to eat durian, rather than a decision between eating durian and not eating durian; second, when calculating our subjective expected utility for each of these new possible actions, the only factor that is taken to feed into the utilities that we assign to the outcomes is the value we attach to having a novel experience. The first part of the solution in fact does not change anything, since both of these decisions have the same outcomes. I can only discover what it's like to eat durian by eating it; and by eating it I discover what it's like to eat durian – indeed, this is what makes eating durian an epistemically transformative experience. So the outcomes of choosing to make that discovery are exactly the same as the outcomes of choosing to eat durian, and the decision problem has not been altered: the outcomes are the same and the problems with accessing our utilities for the outcomes remain in place. The second part of Paul's solution is intended to deal with the access problem. My worry about this part of the solution is that it seems to ignore the very problem that Paul has raised. Underpinning Paul's central objection to expected utility theory is the observation that my overall subjective utilities for the possible outcomes of my actions are determined by a number of factors, including, for instance, the moral value of the outcome, its aesthetic value, the extent to which it realizes my long-term goals, and so on. One of these factors will be the extent to which I value the novelty of any new experiences it affords me; and another will be the extent to which I value what it's like to experience that outcome. Now one of Paul's main objections

to expected utility theory is that, for many important decisions, there are possible outcomes of our actions such that we cannot know what it's like to experience them. If this is true, and if the nature of this first-personal experience is indeed one of the factors that determines the overall utility of that outcome, we cannot know what our overall utilities are for those outcomes. Paul's solution is that we should ignore this factor, along with all other factors except the extent to which the outcome affords the agent a new experience. The problem with this, however, is that it does not seem to solve the problem that Paul has posed any better than a solution that says, for instance, that we should make such major transformative life decisions by simply ignoring all factors that determine the overall utility of the outcomes except their moral value; or by ignoring all factors except their aesthetic value; or by ignoring only what it's like to have the transformative experience, which is, in any case, the only factor to which we do not have access.

I'd like to turn now to Paul's claim that decisions in which one of the possible actions might give rise to an epistemically transformative experience cannot be accommodated within expected utility theory. Consider again my decision whether or not to become a parent. To calculate the expected utility of becoming a parent, I must know the utility of the outcome in which I do; but to know that, I must know what it's like to be a parent; and since I've never been a parent before and since becoming a parent is epistemically transformative, I don't know what it's like. This, Paul argues, spells trouble for expected utility theory. The standard response to this sort of example runs as follows. (Paul Weirich, 2004, section 5.1.2, calls it *higher-order utility analysis*.) Originally, I took there to be just two outcomes relevant to this decision problem: I do not become a parent; I do become a parent. But, since I do not know my utility for the latter outcome, I expand the set of possible outcomes by fine-graining that outcome to give the following new enriched set: I do not become a parent; I become a parent and it is wonderfully fulfilling; I become a parent and it is quite enjoyable; I become a parent and it is quite unenjoyable; I become a parent and it is awful. Now, I know my utility for each of these outcomes, since I know the utility I assign to something being wonderfully fulfilling, or quite enjoyable, or quite unenjoyable, or awful. Finally, I assign subjective probabilities to each of these new outcomes (under the supposition of each of the possible actions). And that gives me all the ingredients I need to calculate my subjective expected utilities. If I have no information about the experience of becoming a parent, I might perhaps follow the Principle of Indifference and divide my credences equally over it being wonderfully fulfilling, quite enjoyable, quite unenjoyable, or awful (on the supposition that I become a parent). But I may have testimonial information from 100 people, 70 of whom found it wonderfully fulfilling, 15 of whom found it quite enjoyable, 10 of whom found it quite unenjoyable, and 5 of whom found it awful. That might lead me to assign a subjective probability of 70% to the outcome in which I become a

parent and find it wonderfully fulfilling (on the supposition that I become a parent), 15% to the outcome in which I become a parent and find it quite enjoyable (on the supposition that I become a parent), and so on. That is, I might take testimony about the extent to which *others* value the experience of becoming a parent as evidence concerning the extent to which *I* will value it.

Paul considers this response but rejects it on two grounds. First, she worries that, for the sorts of cases that particularly interest us, the available data are in some way insufficient to allow us to assign the sorts of subjective probabilities to these enriched outcomes that can underpin rational decision-making (2014, pp. 131-6). Second, she worries that, even if the data were sufficient to do this, any decision that appeals to subjective expected utility calculations that incorporate subjective probabilities that are based on the first-personal experience of other people and not of the person making the decision are thereby rendered suboptimal because the decision maker is alienated from her decision (2014, pp. 124-31). I'll focus on Paul's first reason for rejecting this response.

Paul's first worry, which concerns the data on which we hope to base our credences about our own utilities, is a version of the *reference class problem* from statistics:

Empirical disciplines such as psychology, economics, and sociology collect data on large groups, and subdivide these groups based on externally determined criteria, giving us information about how different types of people tend to respond to different experiences. [...] [F]rom the personal point of view, when an individual is making decisions for herself, this information is, at best, only partially useful. This is because the data just do not give us the kind of fine-grained information about how a person who is just like us, a person with just our particular blend of personal abilities and personality traits, our likes and dislikes, our work ethic and neuroses, and so forth, is most likely to respond to a particular experience. (2014, pp. 131-2)

Now it is of course correct to say that the sort of empirical survey data we are considering will never tell you with *certainty* how you will respond to a particular experience, such as becoming a parent; indeed, it will not necessarily tell you how you are *most likely* to respond, since it may be that there are two or more ways of responding that have joint highest frequency amongst the members of the sample. But the fine-graining strategy above allows us to calculate subjective expected utility without either of these pieces of information. It simply requires a rational assignment of subjective probabilities to the fine-grained outcomes. Perhaps Paul thinks that even that is impossible. If so, I think that sets the bar for rational subjective probabilities too high. I think it's quite likely that there are no rationally *required* subjective probabilities in these cases: given the survey data described above, I'm not *obliged* to assign probability 70% to the outcome in which I become a parent and find it wonderfully fulfilling, because I may think that I belong to a smaller reference class amongst which the frequency of that response is

quite different. But it seems that that assignment is at least rationally *permissible*. Indeed, if it were not possible to assign rational subjective probabilities on the basis of the sort of sociological data in question here, it would not be only the sorts of potentially transformative decisions that Paul considers that would fall outside the scope of expected utility theory. We make decisions about the medical treatment we will administer to a particular patient, the educational strategies we will employ with a particular child or student, the incentives we will give to a particular employee, and so on, all on the basis of the sort of empirical survey data that Paul rejects as a basis for decision-making about whether to become a parent. In those cases, the surveys don't report how much utility the subjects of the survey assign to particular experiences; rather, they report the extent to which the subjects respond to the treatment, or the educational strategy, or the incentives in question. But the quality of the data is the same; they are purely statistical data; they are based on a reference class that is much broader than the narrowest relevant reference class to which the individual about whom you are making the decision belongs. If there is no way to use these data to assign rational subjective probabilities that might underpin expected utility calculations, then the scope of expected utility theory is much more limited than we initially thought.

To conclude, I'd like to turn to Paul's claim that decisions involving actions one of which might give rise to a personally transformative experience cannot be accommodated within expected utility theory (for related worries, see Ullmann-Margalit, 2006 and Bykvist, 2006). The objection, if you recall, is that if one of the possible outcomes of a decision involves an experience that will lead the decision maker to change her utilities, it is not clear whether she should make the decision based on her subjective expected utilities at the time of the decision or at a later time, when her utilities might be different. Here, I think, we come to a serious problem for expected utility theory. I think it is best appreciated by thinking of my decisions at a particular time as decisions I make on behalf of a group of agents, namely, my future selves. If my utilities do not change over time, all of those future selves have the same utility function and so I should just choose an action with maximal subjective expected utility relative to that utility function. But if my utilities do change over time, those future selves have different utility functions. So we face a problem of *attitude aggregation*. How should I combine the attitudes of my future selves – their preferences, utilities, etc. – with the attitudes of my current self to give a single attitude on which I can base my decision? The difficulty that Paul has raised for expected utility theory can be seen as a particular instance of the general problem of aggregating the attitudes of a group of agents into a single attitude. Taking our lead from work on aggregating the doxastic attitudes of a group of agents, we might hope to solve Paul's problem as follows: when we wish to aggregate the subjective probabilities of a group of agents, the natural strategy is to take a weighted average of them (this is known as linear pooling, McConway 1981); perhaps we should

aggregate the utility functions of a group of agents in the same way, by taking a weighted average; thus, perhaps we should choose whether or not to become a parent, for instance, by maximizing subjective expected utility relative to a utility function that is obtained by taking a weighted average of the utility functions of my future selves (Pettigrew 2015, p. 773). The problem with this proposal is that it is not only our utility functions that might be changed by a personally transformative experience: our attitude to risk might also be transformed; even our attitude to maximizing vs satisficing may be transformed. Now, if it were only our utilities that were transformed, we might use the weighted average strategy to effect a compromise between our current and our future selves. But it is not at all clear how we might effect a compromise between a risk-seeking current self and a risk-averse future self, nor between a maximizing current self and a satisficing future self. This is analogous to a problem that faces accounts of moral uncertainty – if one of the moral theories we consider possible is a maximizing theory and the other is a satisficing theory, for instance, it is not clear how we could effect a compromise between them when we assign positive subjective probability to each (Hedden 2015, section 2.2). An interesting direction for future research prompted by Paul's book would be to explore how the different problems of attitude aggregation in different parts of formal epistemology, decision theory, and theory of action might inform one another.

Even at the time of writing, only a year after publication, Paul's book has proved enormously influential, and rightly so. The challenges that it presents to accounts of decision-making are deep, yet presented in a way that makes them accessible to a wide range of researchers and to those outside the academy. It has already succeeded in creating fascinating and productive discussions amongst philosophers of psychology, theorists of value, ethicists, formal epistemologists, political philosophers, philosophers of race, and feminist philosophers; moreover, it has brought researchers in psychology and economics into the debate as well. This is a rare achievement and one for which this book should be widely applauded.

## References

- Briggs, Rachael 2015, 'Normative Theories of Rational Choice: Expected Utility', *The Stanford Encyclopedia of Philosophy* (Winter 2015 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2015/entries/rationality-normative-utility/>>.
- Bykvist, Krister 2006, 'Prudence for Changing Selves', *Utilitas*, 18(3), pp. 264–83.



- Hedden, Brian 2015, 'Does MITE Make Right?: Decision-Making Under Normative Uncertainty', *Oxford Studies in Metaethics* (Oxford: Oxford University Press).
- McConway, Kevin J. 1981, 'Marginalization and Linear Opinion Pools', *Journal of the American Statistical Association*, 76(374), pp. 410–4.
- Paul, L. A. 2014, *Transformative Experience* (Oxford: Oxford University Press).
- 2015, 'What You Can't Expect When You're Expecting', *Res Philosophica*, 92(2), pp. 1–24.
- Pettigrew, Richard 2015, 'Transformative Experience and Decision Theory', *Philosophy and Phenomenological Research*, 91(3), pp. 766–74.
- Steele, Katie and H. Orri Stefánsson 2015, 'Decision Theory', *The Stanford Encyclopedia of Philosophy* (Winter 2015 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2015/entries/decision-theory/>>.
- Ullmann-Margalit, Edna 2006, 'Big Decisions: Opting, Converting, Drifting', *Royal Institute of Philosophy Supplement*, 81(58), pp. 157–72.
- Weirich, Paul 2004, *Realistic Decision Theory: Rules for Nonideal Agents in Nonideal Circumstances* (Oxford: Oxford University Press).

Department of Philosophy  
 University of Bristol  
 Cotham House, Cotham Hill  
 Bristol. BS6 6JL  
 Richard.Pettigrew@bristol.ac.uk  
 doi:10.1093/mind/fzw014

RICHARD PETTIGREW