

CRITICAL NOTICES

Transformative Experience

WILLIAM TALBOTT

L.A. Paul has written an important book.¹ In it she presents an innovative combination of a Sartrean conception of choice as authenticity with a Bayesian conception of choice as expected value maximization to serve as the basis for a unified normative decision theory. The challenge that motivates her account is to explain how to decide rationally in a special class of life-changing decisions, *transformative decisions*, which she defines and analyses. Paul's main argument is an argument for the following conditional claim: In our current state of empirical knowledge, if we want to be able to apply expected value theory to transformative decisions, which for Paul is a necessary condition for understanding ourselves as choosing rationally in such cases, we must understand our decisions as depending primarily on what she calls the *value of revelation*. I am going to question some of the steps of her argument, but none of my criticisms is intended to cast doubt on the importance of her contribution – not only to our understanding of rational choice, but also to our understanding of informed consent, and the philosophy of disability.

Paul's book is primarily an extended exploration of only one of the factors in rational choice, which she refers to as *first personal choosing* (124). Choosing is *first personal* when it is based on the *subjective value* of one's experiences – that is, 'the values of what it is like to have the experiences or of what it is like to be in these experiential states' (11). Because this sounds like a hedonistic theory of value, Paul repeatedly reminds the reader that, as she uses the term, subjective value is the value of veridical or, as she says, *lived* experience (11–12) and she insists that the subjective value of experience includes other cognitive elements in addition to its qualitative phenomenological character (27). Thus, Paul clearly excludes a purely hedonistic theory of value. The importance of first-personal choosing is the Sartrean element in her account.

Paul is not claiming that all rational considerations are first personal or that all value is subjective in her sense. She allows for both a first personal perspective and a third personal perspective on choice. The third personal perspective includes objective values, such as moral values, and objective, scientific evidence about what experiences are like (125). Her position is simply that, given our current empirical evidence, some cases are not rationally decidable from the third personal perspective alone; in some decisions, rational deliberation crucially depends on the first personal evaluation of

1 *Transformative Experience*. By L.A. Paul. Oxford University Press, 2015, x + 190 pp. £18.99 cloth, £11.99 paper.

the subjective values of the alternatives (25–26), or, at the very least, that it would be a disaster for our self-understanding (and, implicitly, for what it would be like to be us) if we were to completely replace the first person perspective with the third person perspective in rational choice (125).

Since subjective values depend on what it is like to have a certain (lived) experience, Paul uses Jackson's (1982) example of Mary, who has never seen any colour but black or white, to illustrate how someone could be unable to imagine what a certain experience would be like. Even if Mary has all the scientific knowledge we currently have about colour perception and even if she has plenty of reliable testimony about what it is like to see red, Mary cannot know or even imagine what it would be like to see red until she sees it (9–10). When she sees red for the first time, Paul says that Mary's experience will be *epistemically transformative*, because it will give her knowledge that could not have been acquired except by having that kind of experience (10).

Paul also identifies another important category of experiences – those that are *personally transformative* – that is, experiences that change you so radically that before having the experience, 'you can't know what it is going to be like to be you after the experience. It changes your subjective value for what it is like to be you, and changes your core preferences about what matters' (17). Paul refers to experiences that are both epistemically and personally transformative as *transformative experiences* (17). The main topic of Paul's book is a special category of choices – choices of whether or not to undergo a transformative experience, which she refers to as *transformative choices* (31).

Paul begins the book with a fanciful example of a transformative choice: Suppose you had to decide whether to undergo the experience of being transformed into a vampire (1–2). What is crucial about this transformation is that not only is it epistemically transformative, in that it involves new sensory experiences, but it is also personally transformative. We really have no way to imagine what it would be like to be a vampire or what it would be like to have the core preferences of a vampire (42–46).

The example of deciding whether or not to be a vampire is a fanciful one, but it is meant to focus our attention on issues that arise in many real-world examples of life-changing, transformative decisions, 'where we want to make a rational choice between relevant alternatives by determining how to act in a way that will have the highest expected subjective value' (25). Paul takes as given that 'the normative standard for rational decision-making is that the agent or decision-maker should choose the act that has the highest expected value' (21). She refers to this normative standard as 'expected value theory' (148). She insists that, for a decision to be rational, it must be the product of maximizing expected value, where the maximizing calculations themselves employ rational credences (subjective probabilities) *and* rational values – that is, both the credences and the values involved must be based on evidence (22–23).

Paul points out that in ordinary cases of choices that are neither epistemically nor personally transformative, we have a first personal way of rationally assigning subjective values to the outcomes. We engage in a kind of *cognitive modelling* (26). We 'evaluate each possible act and its experiential outcomes by imagining or running a mental simulation of what it would be like' (26). There are two factors that make cognitive modelling impossible for transformative decisions: (i) At least one of the

outcomes involves epistemically transformative experience. By definition, it is impossible to cognitively model or mentally simulate an epistemically transformative experience before having the relevant kind of experience. (ii) At least one of the outcomes involves personally transformative experience. What it is like for you to experience a given outcome depends on your other psychological states, including your preferences. But a personally transformative experience alters core preferences in ways that, before the decision, make it impossible for you to know what it will be like to be you after the transformation. For these two reasons, it is impossible to use cognitive modelling or mental simulation to assign subjective values to all of the alternatives (31–32). But if at least one of the subjective values is undefined, then there is no way to calculate the expected value of all of the alternatives, and thus, no way to apply the formula for maximizing expected value (30–33). The key idea is that, unlike uncertainty about the outcomes, which can be modelled probabilistically, there is no way to model unknown values (32–33).

Paul's analysis of transformative decisions ultimately leads her to a point where she can see only three possible alternatives for making transformative choices:

- (1) Replace first personal with third personal reasoning. Notice that the impossibility of cognitively modelling transformational choices is a first personal limitation. If it were possible to assign subjective values on the basis of third personal information – for example, empirical social science or testimony from those who have undergone the relevant transformative experience – then applying normative expected value theory to transformative decisions would be possible. Paul recognizes this, but her Sartrean self argues that to lose the first personal perspective on choice would be 'disastrous' (87). She also has arguments that we simply could not rationally use the third personal information that is currently available to us to replace first personal cognitive modelling. I consider those arguments shortly.
- (2) Continue to make transformative choices from the first person perspective, but give up on making them rationally. She finds this alternative so unsatisfactory that it motivates her to find another alternative.
- (3) Reframe the choice in terms of the revelatory value, not the subjective value, of the outcomes – for example, as a choice between discovering and not discovering what it is like to have the experiences and preferences of a vampire (114, 118). As she points out, the choice cannot be evaluated as the difference in *subjective value* between discovering and not discovering what it is like to be a vampire, because, *ex ante*, the subjective value of discovering what it is like to be a vampire is just as unknowable as the subjective value of what it is like to be a vampire (93). The choice to become a vampire must be based on the value of revelation *for its own sake* (122–123).

Consider the first alternative, using third person information to evaluate the expected subjective value of transformative experiences without mentally simulating them. Paul does not claim that we can never do this. She does allow that there are some cases in which we can evaluate the subjective value of an outcome without being able to mentally simulate it – for example, the experience of being eaten by a shark (27) or of having your leg amputated without anaesthesia (28) or of being hit by a bus (110), to assign them very low subjective value. Call these cases *shark-like cases*. Paul

distinguishes the shark-like cases from the case of deciding whether to become a parent and the other transformative choices she discusses, because in the cases she is interested in ‘you are not sure how you would respond to an experience’ (28).

If I read her correctly, Paul is holding that you could rationally decide to avoid being eaten by a shark or avoid having your leg amputated without anaesthesia or avoid being hit by a bus *because of the very low expected subjective value of any of these outcomes*, even if each of these outcomes involves a transformative experience that you cannot mentally simulate *ex ante*. Third personal information is enough. And, if I read her correctly, Paul would also agree that you could rationally avoid a high probability of any of these outcomes because of its very low expected subjective utility, even if they all would involve transformative experiences that you could not mentally simulate *ex ante* (cf. 127). Again, third personal information is enough. What is the relevant difference between these kinds of transformative experiences and the example of choosing to become a parent and the other kinds of transformative experiences that Paul thinks cannot be rationally made on the basis of third personal information about the expected subjective values of the outcomes?

In the Afterword, Paul explores this issue in detail. I focus on what seems to me to be the main line of her argument: That for these choices the third personal information ‘isn’t personalized enough to make accurate individual-level decisions’ (126). She points out that statistical research in the social sciences, such as psychology, economics and sociology, ‘is often focused on broad demographic categories, for example, on what people of different socioeconomic classes, genders, races and cultures tend to do, often with further subdivisions based on general characteristics like age, health, background, life stage, leisure pursuits, personal abilities and personality traits’ (131). She thinks this information can be useful for government or other large-scale institutional planning and policy-making, but she thinks that it is much less useful for individuals, ‘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’ (132). From this she concludes that we cannot rationally replace first personal information with third personal information in the cases of interest (132).

However, she does not acknowledge that her way of dealing with this problem would have profound implications for rational choices of all kinds, both transformative and non-transformative. Suppose you are diagnosed with a form of cancer for which there are two treatments, one with a 40% survival rate and one with a 10% survival rate. Suppose that the two treatments cannot be administered simultaneously and that there is only time to administer one treatment before the cancer will be too far developed for further effective treatment. Suppose finally that the 10% who would survive with the second treatment are not a subset of the 40% who would survive with the first treatment, but that there is no way of determining which, if either, treatment your cancer would be responsive to.

Paul’s argument would imply that in this case, you could not use the statistical information to rationally decide which treatment to choose. You might as well flip a coin to decide which treatment to receive. I think this is a mistake. I think that, in the absence of other relevant information, rationality requires you to choose the 40% option. Paul can correctly point out that, in the example as I have described it, it may well be true that, because of the particular characteristics of your tumour, if you

choose the 40% option, it will be ineffective, but if you choose the 10% option, it will be effective, even though it has a lower survival rate than the 40% option. To me, that is not a reason for thinking that the choice of the 40% option is not rational, but only for acknowledging that rational choices sometimes have worse outcomes than irrational ones. This sort of example leads me to conclude provisionally that Paul's argument does not rule out rationally choosing on the basis of third personal information based on 'impersonal' statistics.²

Even if we suppose that third personal information could be used to make transformative decisions, Paul would not look favourably on basing transformative decisions on third personal rather than first personal information. As I have already mentioned, she believes that doing so would be 'disastrous' (87), for it would 'do great violence to our ordinary way of thinking about deliberation' (128). This is the Sartrean part of her view.

To dramatize this concern, Paul imagines a woman, Sally, deciding whether or not to become a parent, who consults the empirical evidence on self-reports of levels of satisfaction by parents and non-parents. As it happens, there is a fair amount of evidence of this kind in the psychological literature. In a footnote, Paul reviews this evidence and concludes that most studies show that parents have a lower level of subjective well-being than non-parents (86, n. 46). She concludes from this that if Sally were to estimate the expected subjective values of the options solely on the basis of the empirical research, the research would indicate that she should not have a child (86–87).

I think that this is a mistake. To isolate the mistake, I want to stipulate away any ambiguity in the scientific literature. Let's suppose that the empirical studies were unequivocal that, based on self-reported levels of satisfaction, non-parents had significantly higher levels of satisfaction than parents. Paul clearly thinks that the advocate of using third person testimony would be committed to holding that this evidence would rationally require ranking the expected subjective value of a childless life above the expected subjective value of life as a parent.

I agree with Paul that this result would be disastrous. She uses it as a kind of *reductio* of the proposal to use third personal testimony from parents and non-parents to evaluate the subjective values of parenthood and non-parenthood. But her conclusion follows only if we understand subjective values in terms of self-reported feelings of satisfaction. Nothing in Paul's account of subjective value suggests such a hedonistic interpretation of subjective value, and some of what she says actually excludes it. She rules out hedonistic evaluation when she stipulates that subjective evaluations are evaluations of 'lived' experience – that is, veridical experience – and when she insists that the subjective value of experience is not solely a function of its qualitative elements (27) – that subjective value 'is distinct from merely valuing happiness or pleasure and pain' (178).

Since Paul explicitly rejects a hedonistic interpretation of subjective value, it is surprising that she would think that studies on self-reports of levels of satisfaction

2 In a footnote, Paul sets aside examples that involve physiological reactions, because she claims that 'when the relevant data exist, [such examples] admit of more precise prediction and assessment' (132, n. 8). I think this is a mistake. There are many medical situations in which we are almost wholly ignorant of the factors that determine whether or not a treatment will be successful. The science of designing a treatment based on the precise signature of a particular tumour is in its infancy.

would or even could settle the question of what it is rational to do. Perhaps she understands satisfaction non-hedonistically – for example, as some sort of measure of preference satisfaction. She seems to think that this non-hedonistic measure would be an acceptable proxy for subjective value – for example, when she says that if empirical work could provide us with the relevant individual-level probabilities for the outcomes, ‘along with how satisfied we’d be with them, we might be able to simply replace our ordinary deliberative procedure with the expected values specified by these results’ (127).

I think that it is a mistake to think that satisfaction can serve as a proxy for subjective value. Satisfaction, even when it is understood non-hedonistically, is an impoverished basis for evaluating the subjective value of a human life.

To understand how Sally might rationally choose to have a child even if the third personal evidence is overwhelming that non-parents report higher levels of satisfaction than parents, I have to say something about what is involved in living a meaningful or worthwhile life. As I use the terms, meaningfulness or worthwhileness is a hybrid objective/subjective value. That it is partly objective is illustrated by the fact that, as I use the terms, no life lived on an experience machine would score very high on meaningfulness or worthwhileness, no matter how meaningful or worthwhile it seemed to be to the person who was experiencing it. That it is partly subjective is illustrated by the fact that the belief that one is living a meaningful or worthwhile life makes an important difference to what it is like from the inside. But that difference cannot be captured by any usual measure of life satisfaction.³

A life that seems to be aimed at something of genuine value and importance can at times generate deep satisfaction, but it also can and typically does present frustrations and obstacles that call forth great exertions; it can require great personal sacrifice; it can and often does produce great regrets; and, in many cases, it includes great suffering. The sense of value and importance of a life does not typically make those experiences pleasant or satisfying; it makes their being unpleasant or unsatisfying seem less significant. There are many different goals that one can pursue that give a sense of value and importance to one’s lived experience. Having or adopting a child is one way to do so that is accessible to almost everyone.

When parents say that having children was a life-transforming experience or that it was the most important decision they ever made or that it was the most personally rewarding thing they have ever done, there is a temptation to understand them as making claims about how satisfying their lives are. But if you investigate their lived experience, you will probably find something not so satisfying – for example, that it involved staying up all night with a vomiting child, only to be stressed out the next day because they were not prepared for an important meeting at work; or that it involved giving up on obtaining a promotion because of childcare demands; or that they lost years of a normal sex life because they were often too busy and tired to make time.

I suspect that most parents who read Paul’s discussion will find themselves nodding in agreement as she describes the life-transforming aspects of becoming a parent for the first time. But even before the glow of holding one’s child for the first time wears off, new parents typically begin to appreciate something that they could not fully

3 For a fuller argument for understanding the meaningfulness of a life as an objective/subjective hybrid, see [Wolf \(2010\)](#).

appreciate before – the awesome responsibility they have just acquired for another human being. One big discovery that most parents only make after becoming parents is that, no matter how old your children are, you never stop worrying about them.

Once you know what it is like to be a parent, if you are not self-deceived, you may come to realize that if you had not become a parent, you could have easily had a life with more satisfaction. If you are lucky, when you become a parent you will gradually forget how enjoyable it was to have a normal sex life or how good you felt when you were more physically fit because you worked out every other day or how you felt when you could run a 6-min mile. If you don't forget, you make compromises, because some things are more important than satisfaction. Think about what it is like to find something of seeming value and importance to pursue. When you do, maximizing the felt-satisfaction of your life can seem trivial by comparison.

This is not to say that a life of seeming value and importance cannot be extremely satisfying or that in such a life feelings of satisfaction do not matter at all, but only that to think that you were living the life in order to maximize self-reported satisfaction would be to seriously misunderstand what typically motivates such a life. Suppose Dan and Susan are considering whether or not to have children in order to maximize their self-reported satisfaction. I would hope that the scientific evidence would lead them to decide not to do it. Suppose that they are unaware of the scientific evidence and decide to go ahead. Becoming a parent may be a transformative experience, and thus it may not be possible *ex ante* for them to evaluate the subjective values of the alternatives; but after they have a child they will have no difficulty at all in evaluating the subjective value of the alternatives. Suppose they discover that parenthood is not nearly as satisfying as they thought it would be. Suppose they come to realize that their earlier life – when they travelled, went to fine restaurants, and had a regular sex life – had a much higher level of satisfaction. Clearly, maximizing expected satisfaction would require them to give the child up for adoption. Anyone who would make such a calculation is unfit to be a parent.

But for someone who likes children and wants to live a life that contributes to something of importance – which is my take on Paul's example of Sally (87–88) – it is hard to think of something more important than playing a central and irreplaceable role helping to form a new person and helping that person herself to develop into someone who is able to live a life that contributes to something of importance. For this Sally, the scientific evidence on self-reports of satisfaction would be largely irrelevant. What would be relevant would be the reports of other parents who tried to describe to her, not only the satisfactions and dissatisfactions of life as a parent, but also the sense of the significance of the life-long journey they have embarked on.

I don't mean to give the impression that parents live their lives constantly aware of the importance of what they are doing. Sometimes, when they are feeling especially proud of their child, they will reflect with satisfaction on the importance of what they are doing. But most of their lives will not be like that. The conviction that what they are doing is important will manifest itself not as a feeling of satisfaction, but in other ways – for example, as a reservoir of determination and energy that will push them to continue trying to do more than it is humanly possible to do.

Maybe, when you are considering becoming a parent, what you want to know from your friends and acquaintances is not how satisfying their lives are, but rather: If they had it to do over, would they choose to have children? Remember, they have been through the transformative experience, so they have some way of comparing the

subjective value of the two options. If their answer is yes, and if their reasons for saying yes indicate that they regard the life of a parent as a worthwhile or meaningful one, why isn't that relevant evidence? Wouldn't it be useful for those anticipating transformative decisions (e.g. careers) to have reliable empirical information giving the percentage of those who made the relevant choice (e.g. a particular career) and later regret choosing it? Couldn't such evidence, if one acquired it, be grounds for rationally choosing not to pursue a career that one had antecedently been interested in?

When we consider the possibility of evaluating the subjective value of a life, at least in part, on the basis of the significance that it is thought to have, we see that there is also a potential problem with Paul's normative standard of rational choice, expected value maximization. Living a worthwhile or meaningful life probably does not involve maximizing anything, surely not meaningfulness or worthwhileness. Meaningfulness or worthwhileness seems to be a holistic property of a life, not anything like a sum of units of meaningfulness to be maximized. If the goal is to live a worthwhile life, then it may be more important to minimize the probability of living a non-worthwhile life than to maximize worthwhileness (whatever that might mean).

Paul's own way out of the conundrum of how to rationally make transformative choices at times sounds similar to my proposal in terms of living a life felt to be meaningful or worthwhile (93), but there is a big difference. Because Paul's argument leads her to conclude that in transformative decisions (e.g. the decision to become a parent) we do not have third personal evidence that would enable us to rationally evaluate the expected subjective value of the options, she has to propose a way of deciding rationally that completely ignores the expected subjective value of the options. Her proposal is to substitute the value of discovery (revelation) *for its own sake* (122–123). But this proposal quickly generates problems.

To see why, consider another example of a transformative decision that Paul discusses, marriage. Suppose that Ken proposes to Gina as follows: 'I hope you will decide to marry me, but only if it is rational for you to make that choice. Since you agree with Paul that marriage is a transformative choice and that you cannot rationally decide on the basis of what it will be like to be married to me, to make a rational choice you must decide on the basis of revelation for its own sake. You know yourself that you place a high value on discovery, so your only rational choice is to decide to marry me.'

I think it is clear that Gina could have all sorts of third personal evidence that would make it rational for her to decide not to marry Ken, even though she would have to admit that marrying him would not be like getting hit by a bus, in that being married to him *could* turn out to be a worthwhile and satisfying life. Still, it might be a bad gamble and it might be rational for her to keep looking for a more compatible mate. One moral of this story is that it is probably always irrational to decide on the basis of revelatory value alone, without at least some reason for believing that it is unlikely that the revelation will be the discovery that you have chosen a miserable, mostly meaningless existence. For the reasons discussed above, I think we can and do use third personal information to rationally make such judgments, even in transformative choices. Perhaps the most implausible implication of Paul's view is that we cannot.

In conclusion, let me emphasize that none of my criticisms is meant to cast doubt on what I said at the outset: *Transformative Experience* is an important book. Paul's

introduction of subjective value (in her sense) into the theory of rational choice is an extremely important contribution to the literature, especially when it is recognized that subjective value is not limited to measures of satisfaction. Her definition and analysis of transformative decisions is an equally important contribution, even if Paul is mistaken to think that the rationality of such decisions can be fully analysed in terms of revelatory value. Paul shows the value of her account by using it to open up insightful new ways of thinking about such important questions as informed consent, sensory modifying technologies, marriage, parenthood and career choice. Her discussion of cochlear implants by itself is worth the price of the book. Almost anyone who makes choices or theorizes about them will find this book to be a revelation.

University of Washington Seattle
WA 98195–3350, USA
wtalbott@uw.edu

References

- Jackson, F. 1982. Epiphenomenal qualia. *Philosophical Quarterly* 32: 127–36.
Wolf, S. 2010. *Meaning in Life and Why It Matters*. Princeton: Princeton University Press.