Epistemically Transformative Experience*

Jane Friedman | jane.friedman@nyu.edu

04/15

1 Introduction

In her book, *Transformative Experience*, L.A. Paul argues that some decisions involve transformative experiences. Two kinds of transformative experiences are the focus of her book: *epistemically transformative* experiences (ETEs) and *personally transformative experiences* (PTEs).

– ETE: "When a person has a new and different kind of experience, a kind of experience that teaches her something she could not have learned without having that kind of experience, she has an epistemic transformation."

– PTE: "If an experience changes you enough to substantially change your point of view, thus substantially revising your core preferences or revising how you experience being yourself, it is a personally transformative experience."

Paul argues that decisions involving transformative experiences – of either sort – pose a problem for decision theory (DT) and for theories of rational decision making more generally. I want to focus my discussion on ETEs. My plan will be to get clear on Paul's argument that ETEs generate some sort of problem for theories of rational decision making and say a few things about it.

2 The argument

Paul's claim about ETEs is that they pose a problem for standard accounts of rational decision making, and in particular for DT. The very rough thought is this: You can't

*Presented at the Pacific APA in Vancouver, April 2015 as part of an *Author Meets Critics* session on L. A. Paul's book, *Transformative Experience*.

know "what it's like" to have an ETE. Some decisions though are (in effect) decisions to have ETEs. But if you can't know what it will be like to have the ETE, then how can you rationally decide to have it or not? Paul's thought is that in making decisions involving ETEs in the relevant ways, we are saddled with a sort of ignorance about the outcomes of our decisions that makes it difficult or even impossible for us to make those decisions rationally and authentically (more on 'authenticity' in a moment). But she thinks that these sorts of decisions are ones that we have to make all the time and they're some of the most important ones we make. This is obviously a jarring result: according to Paul, many of the decisions we've already made and will have to make in the future, cannot be made both rationally and authentically.

Let's try to get her argument out more carefully.

(1) At least some decisions involve our deciding to have experiences, i.e., choosing between outcomes that involve having experiences.

(2) In at least some of those cases the experiences we'll have are ETEs. Let's call this set of cases Δ.

(3) An act is rational only if one's preferences over the possible outcomes of that act are rational.

(4) One's preferences over the possible outcomes of some act are rational only if those preferences are all informed – based on information about those outcomes.

(5) If we try to inform our preferences in the "standard way", then our preferences over outcomes involving ETEs cannot be informed preferences.

(6) So, if we try to inform our preferences in the "standard way", then our preferences over outcomes involving ETEs cannot be rational preferences.

(7) So, if we try to inform our preferences in the "standard way", then we cannot act rationally in the cases in Δ.

(8) Non-standard ways of informing our preferences leave us with "inauthentic" preferences and so "alienate" us from our own choices and futures; we can get rational preferences in these ways, but they won't be authentic.

(9) So, in the in the cases in Δ we cannot act both rationally and authentically.

Paul argues that many of life's central and not-so-central decisions are implicated

here: decisions about whether to pursue certain career paths, have a first child, go to war, try new foods or drugs, and many more. So Δ is a set of cases that involve truly central life decisions, as well as many everyday ones. In part, Paul is claiming that perhaps none of us have ever made these decisions both rationally and authentically.

2.1 Some clarification: on Informed preferences and standard decision procedures

What are "informed preferences"? This isn't Paul's term, but I think it applies nicely to her cases. Let's say that a preference for an outcome A over another outcome B is an *informed preference* iff it's based on information about A and B – their natures, what they are like, what happens in them, etc. The basic thought is that informed preferences are preferences based on evidence.

(Note: the expression 'what they are like' is somewhat loaded in this context, but I don't mean to be using it in any special way at this point in the discussion. The outcomes the subject is choosing between will be some way or other – the relevant bit of the world will be some way or other and the subject will experience that bit of the world as being some way or other. All of this (not just the "phenomenal") is part of what the outcomes are like.)

How do we base preferences on evidence? They aren't beliefs after all, how should this work? Aren't some preferences just brute – our matter of fact basic tastes, likes and dislikes etc.? Maybe. And perhaps those are arational. But when it comes to outcomes, the thought is, we need to know what they are, what they're like, what happens in them, to have informed preferences about them. If Ben and Jerry's introduces two new flavours: Abracadabra and Zoolander, and I have no idea what they've concocted with either, then I might in fact have preferences over the outcomes in which I eat each, but those won't be informed. If I find out that one is mint chocolate chip and the other is Oreo, then I'm in the position to have informed preferences over those outcomes (given what I already know and how I already feel about those flavours).

How much information do we need to have about outcomes in order that our preferences over them be informed in the way required for them to be rational? We'll come back to that question.

What about the "standard way" of informing our preferences? At least in the sorts

of cases in Δ, Paul thinks that there's a specific method that we want to or should use to inform our preferences. We can call this method *imaginative forecasting* (IF). When we use this method we imagine the various outcomes and what they'll be like and then base our preferences on the results of this act of imagination. She describes this as "running a mental simulation of what it would be like".

(4), (5) and (8) are are the most interesting premises here and so I want to discuss each in turn.

3 Premise (4)

(4) – one's preferences over the possible outcomes of some act are rational only if those preferences are all informed – is not uncontroversial, but I think that Paul is on to something interesting. What's the controversy?

First, even the claim that preferences are open to rational assessment is not obviously true, as we've seen. Certainly many of our preferences feel immovable in ways that also make them feel plausibly arational. In fact, this was Hume's thought and many have followed. Are preferences really the sort of thing that can be rationally assessed at all?

> 'Tis not contrary to reason to prefer the destruction of the whole world to the scratching of my finger. 'Tis not contrary to reason for me to chuse my total ruin, to prevent the least uneasiness of an Indian or person wholly unknown to me. 'Tis as little contrary to reason to prefer even my own acknowledg'd lesser good to my greater, and have a more ardent affection for the former than for the latter. (2.3.3.6)

Others have thought that some rational constraints should be placed on our preference sets if not our individual preferences. A "moderate Humean" thinks that rational preferences are properly arranged preferences. On this view, it may not make sense to talk about an individual preference as rational or irrational, but sets of them can be, e.g., a set of preferences might be irrational if the set includes a preference for A over B and a preference for B over A.

Paul's claim is stronger than this: it's the claim that for our preferences to be rational they need to be properly based in evidence. This is a bolder claim but I think it has some things going for it. Other things equal, there's something wrong with my Abracadabra-Zoolander preferences when I have no information about what

those flavours are. Also, it's worth thinking about the analogous claims about belief in epistemology. Few think that Humeanism about belief is plausible, but at least some think that epistemic rationality is a matter of belief coherence (moderate Humeanism about belief). But this feels like it leaves out something central about epistemic rationality. If coherence is all that matters, then the subject who followed their evidence and arranged their beliefs in response to it and is coherent at t is no more epistemically rational than the subject who was hit by a rock in the head and ended up with coherent beliefs as a result at t. But most of us think that that's not right. How you got your beliefs and what sustains them – what the bases of your beliefs are – matters. And I think it's plausible enough that similar claims could be made about preferences.

That said, the plausibility of Humeanism in the realm of preferences can make it feel as though preferences are importantly different from beliefs and perhaps shouldn't be subject to the same sorts of strong rational constraints.

Moderate Humeanism is typically the way when it comes to DT. For many, central bits of DT fall out of those sorts of coherence constraints on preferences. I don't think there is a ton of support for Paul's strong anti-Humeanism among DT-ists, and some will even opt for the bare Humean thing. DT tells you what to do when you have preferences over outcomes (and some relevant credences); it's not typically concerned with where those preferences came from. It's a theory that tells us how to act when we have certain kinds of beliefs and desires, rather than one that (also) tells us how to get certain kinds of beliefs and desires.

If 'rational preferences' are understood to be informed preferences, then I don't think there's any support for (3) in central bits of DT. DT isn't reserved for those who have based their preferences on complete information about outcomes, it's for anyone with beliefs/credences and desires/utilities (perhaps properly arranged).

And of course there's something to this thought. If you strongly prefer A to B, even if those preferences aren't properly based, there is a sense in which it's rational to act in accordance with those preferences if they are relevant to the decision you face; they are your preferences after all (it would be weird were you to act against them instead). And if you do act in accordance with those preferences, your behaviour will be explicable in their light. But while we have this very basic notion of rationality that is respected whenever someone acts on their beliefs and desires (no matter where those came from), I think many of us, Paul included, think that there must be more to rationality than that.

I'm not sure that DT is on board though and so if Paul is accusing DT of making it that many of our central decisions are being irrationally made, we should be worried about whether that's right at this point. DT is far more permissive than the theory Paul is promoting now. I'll come back to this later once we hear more about why certain kinds of decisions are always irrational.

4 Premise (5)

So (4) is a strong claim, but I think not implausible. Although what comes next surely puts more pressure on it. (5) says that if we try to inform our preferences in the "standard way", then our preferences over outcomes involving ETEs cannot be informed preferences.

What does Paul want to say in defence of (5)? Why can't we have informed preferences over outcomes involving ETEs when we use IF? We can start with the idea that when it comes to ETEs we are ignorant in particular ways. Let's call what we're missing with respect to ETEs, *exact phenomenal information*. We don't have that sort of information until we've had the relevant type of experience: information about "what it's like" to be on LSD, taste a new fruit, see a new colour, be blind, have a child, and so on. Of course, there is a great deal that we can know in advance about these ETEs and what they will be like (in the broad sense), but until we've experienced them we can't have certain kinds of phenomenal information: exact phenomenal information. By calling this phenomenal information 'exact' I mean to be highlighting the fact that for most ETEs we do have plenty of phenomenal information. We typically know something or quite a lot about the phenomenology of a new experience (perhaps unlike Mary in her black and white room). We've had other experiences of a similar type and so we almost always have some phenomenal information – we know something about what these experiences will be like. But there's still something we're missing: the information or knowledge that we can only get from having the experience. I'm calling that missing piece 'exact phenomenal information/knowledge'.

So, in the cases in Δ we can say that we lack exact phenomenal information wrt at least one of the relevant outcomes. The question now is why we need just that information to have informed preferences in those cases. Obviously there's a lot of information about the relevant outcomes that we can have/know. We know a lot about what life is going to be like when we have kids: we know a lot about the moral implications and financial costs, many of us have plenty of people to

talk to and have spent plenty of time with kids and parents before we have them. We also have lots of phenomenal information in many of these cases: we've had experiences at least somewhat like these. There is without a doubt a lot we know about the outcomes involving ETEs; we are not entirely in the dark. There are obviously some epistemic blindspots, one of which is this exact phenomenal aspect. But we have so much information, why isn't it enough?

Paul thinks that in the cases in Δ our lack of exact phenomenal information makes it that we can't properly imagine the various outcomes and so can't properly use IF when trying to inform our preferences and make our decisions. To use IF, to imagine each outcome with sufficient richness to get informed preferences over those outcomes as a result, we'd need that exact phenomenal information.

(Note: this sort of discussion is further evidence that DT isn't the issue here. DT is largely a criterion of rightness, not a decision procedure. Using heuristics so as to achieve the outcome with the greatest expected utility is just fine by the lights of DT (as good as doing an explicit EU calculation). And even if DT has some implications about how to decide, it certainly does not require that we make our decisions using this imagination-based method.)

I'll come back to the thought that we do or should use IF in these cases in the next section, but even so far we might wonder why imaginative forecasting should fail so badly given just our lack of exact phenomenal information. We have lots of information, including phenomenal information. It's hard to see quite why the information we have is not enough.

Perhaps this is something. Paul says a lot about the information we acquire from friends and other sources. She worries that we can't know for sure that it applies to us, and that maybe we're different from the others. Other people are madly, obsessively in love with their kids, but maybe that won't happen to us. But these feel like some sceptical 'maybes' and 'mights' creeping in. Of course there is no internal guarantee that the information we have is not misleading evidence or that we know the things we take ourselves to know, but giving into those sorts of concerns leads nowhere but scepticism. Short of that, doubt and uncertainty are perfectly compatible with rational opinion and action and so why not rational preferences? This is one lesson of a good deal of late 20th and early 21st century epistemology, and it seems to me to apply straightforwardly here.

In fact, Paul herself thinks that sometimes we can have preferences over outcomes involving ETEs. In these sorts of cases, I take it that she does think that the other

sorts of information we have suffice and that either we don't need to imaginatively forecast or that that method works well in some cases involving ETEs. For instance, say my two options are to be mauled by a shark or not. There I think she wants to say that I can rationally prefer not to be mauled (and as a result of imagining the outcomes). That is, I have enough information to have imaginatively informed preferences in that case. I don't have exact phenomenal information there though if I've never been mauled, although I do have some general relevant phenomenal information, as I do in the other cases in Δ. So what's going on in this shark case? Presumably, I know that the mauling will come with experiences I don't want, e.g., extreme pain. But as before, maybe they won't. Maybe I'm the kind of person who will experience great ecstasy when my body is traumatized in that way; I've never had that exact type of experience, after all. Yes, maybe. But that's not a 'maybe' we respect here, so why are we respecting them in the other cases? We can gather plenty of information that will most likely apply to us, so why isn't it enough? (Of course, we might have set our preferences wrong in the end, but that doesn't mean that they weren't rational to begin with.)

We might wonder about how much we should care about this exact phenomenal information at all. Yes, how I'm going to feel when I have babies matters to the value of that outcome to me. But so do other sorts of considerations: whether I want to give life to someone, whether we should be making more people, whether I'll be a better person if I have a child, whether I can afford it, and so on. Sometimes it's hard to see how my lack of knowledge about exactly how having a baby is going to feel should be so important in the face of all of these other considerations. And it's certainly unclear why in the face of all of these considerations, plus the fact that I have some phenomenal information, the lack of exact feeling information should prevent me from having informed preferences via IF.

5 Premise (8)

Paul does think that we can get informed preferences in other ways. If we use a method other than IF we might be able to have rational preferences. Basing our preferences on moral information or testimony (etc.), can get us rational preferences and thereby put us in the position to make rational decisions. She thinks that we don't want to make the decisions in these other ways, and that we standardly use IF. Is this an empirical claim? I think I need to hear more about exactly what this method is and why we should think it's the one we want to use. Even if it's

true that I need to figure out what the relevant outcomes are going to be like (for myself) in order to have rational preferences over them, why does that have to go by way of imagination rather than just plain old thought and reason?

Paul does want to say that we have reason to prefer IF in these cases at least. What's that? Other forms of decision-making are "inauthentic" and "alienating". She groups these other methods under the umbrella of "third-personal" decision-making. This sort of decision making is a matter of ignoring exact phenomenal information or maybe even ignoring phenomenal information altogether. Crucially though, in these cases we don't use IF (whether we use phenomenal information or not). She wants to say that in these cases our preferences wouldn't be authentic.

It's hard to pin down just what's inauthentic about (say) acting on my deeply held preferences though, regardless of how those were formed. If I badly want to have a child – whether those preferences have been with me for as long as I can remember, or they came to me the other day after meeting my nephew, or I got them based on the empirical data – isn't my decision authentic? It's at least mine, based on my wants and preferences and so I'm not sure why it's alienating to have those preferences or decide based on them. As far as I understand though, this is what Paul wants to say. Paul wants to say that if I come to prefer the outcome in which I have a child in some way other than by imaging what that outcome will be like for me in sufficient richness then my preference isn't authentic, and I'd be alienated from my decision – even if it's based on some of my most deeply held desires, the ones I most strongly identify with. I'm not clear though on the sense in which acting on deeply held desires and basic preferences is inauthentic or alienating.

6 Where does this leave us?

I want to bring some of the thoughts here together and draw some further conclusions.

It's important to feel the weight of Paul's claims and conclusions here, I think. She thinks there is one way that we all want to (and need to if we want to be authentic) make the relevant decisions, and ETEs bar us from rationally making the decisions in those ways. Since most of us have made many decisions of the relevant sort, Paul seems to be claiming that most of us have failed to make those decisions rationally (and will continue to?). Or if we did manage to make them rationally, we were inauthentic or alienated from those decisions. We are in a quite bad situation when

it comes to these decisions according to Paul then.

Paul sometimes claims that it is DT or EU theory that's to blame for leaving us in this predicament. I want to make clear that it is not. DT and EU theory do not have the result that these decisions are irrational, it's Paul's additional constraints on rational decision making and preferences in combination with her strong claims about informed preferences and imaginative forecasting that get us this result. If it seems like a bad result, then it's not DT to blame, but Paul.

Why isn't DT to blame? DT instructs us to act on our credences and utilities. If we have determinate preferences (that are in keeping with some coherence constraints), then it's rationally permissible for us to act in accordance with them. If you strongly prefer having a child to not, DT (other things equal) says do it. It's Paul telling you that you're either irrational or inauthentic if you do it, not decision theory. If we have no determinate preferences, DT is silent; it simply isn't so pervasive to give us verdicts about how to act in these cases.

Although Paul talks as if DT is to blame for our predicament – that we cannot choose both rationally and authentically – I don't think that's quite what she intends. Maybe the thing to say is this. DT can be very permissive. It typically allows that most any preference-based decision can be rational. And perhaps this fails to capture how we think we really ought to make many decisions. There is no demand that our preferences be informed, no constraints on which decision procedures we should use, and no way to accommodate the true complexity of these sorts of life decisions. This is a legitimate complaint. DT is highly idealized and as such may well be stripping away all sorts of important aspects of real life decision making.

That said, the additional constraints that Paul imposes on decision making in the relevant cases are leaving us in the position of not being able to decide both rationally and authentically. Like Paul, I think that rationality should demand something of us and that it's not reducible to mere post hoc explicability. And I think that there's something right about the thought that information matters, that basing matters, not just when it comes to our beliefs, but our preferences too. And it's also true that there is some information that we lack about the outcomes in the cases in Δ. But my own feeling is that it's hard to see why that information – exact phenomenal information – matters so profoundly to our rational decision making, even in the relevant cases. We have a lot of information about the relevant outcomes (phenomenal and otherwise), and so our preferences can certainly be pretty well informed. I find it hard to see why the information couldn't put us in the position to decide both rationally and authentically. A lot seem to hang for Paul on our need to imag-

inatively forecast in the relevant cases. But it's not clear to me that we do have that need, nor that our lack of exact phenomenal information renders the method largely idle.

Where are we left in the end? My own feeling is that it depends in part on how bad one thinks the result that we have to choose between rationality and authenticity when deciding to have children, take new career paths or make a whole range of everyday life decisions is. I guess I think it's very bad, reductio bad even, and so something that got us there should go. A theory of rational decision making shouldn't have the result that it is so hard to make authentic rational decisions (and for some of our most important life decisions, too). On Paul's account, rational decision making is itself alienating in the cases in Δ.

At some point Paul suggests that we can restore rational and authentic decision making in the cases in Δ if we make decisions just based on our preferences for having new experiences. So I can rationally and authentically decide to have a child based on my (informed) preference for trying new things over maintaining the status quo. But this feels like further evidence that we've gone down a bad path here somewhere. Because I've liked new experiences in the past, it's rational for me to prefer having a child to not. But all of my moral knowledge, my empirical research, my talking to friends, my spending time with kids and parents, my soul searching and more, can't make it rational? Again, we're in reductio territory for me here. Evidence gathering and rational inquiry can't make my decision authentic and rational, but my fetish for novelty can?

In the end I think reflecting on ETEs is obviously helpful for thinking about rational decision making and preferring. If we want many of our central decisions to come out both rational and authentic at least some of the time, then we'll need to say either that preferences involving ETEs can be informed and authentic (so that either IF can work fine despite the lack of exact phenomenal information or that other methods can lead us to both rational and authentic preferences), or that we can make authentic and rational decisions without informed preferences.