

# EPISTEMICALLY TRANSFORMATIVE EXPERIENCE\*

Jane Friedman, New York University<sup>†</sup>

## 1 INTRODUCTION

In her book, *Transformative Experience*, L.A. Paul argues that in deciding how to act we are sometimes also deciding whether or not to be transformed. And that these sorts of transformations or ‘transformative experiences’ throw a spanner in the works of central accounts of rational decision-making. Two kinds of transformative experiences are the focus of Paul’s book: *epistemically transformative experiences* (ETEs) and *personally transformative experiences* (PTEs). Here is how she initially fleshes out each of these,

**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. (TE: 11)<sup>1</sup>

**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. (TE: 16)

Paul argues that decisions to have or not have these sorts of transformative experiences (call these *potentially transformative decisions*) pose a problem for standard accounts of rational decision-making, including decision theory. The discussion that follows will focus on ETEs and the difficulties Paul thinks they make for our accounts of rational decision-making. Some of the main examples of ETEs in the book: having a first child, taking a novel career path, being/becoming a vampire, tasting a strange new food, and Mary’s seeing red for the first time. In the next section I’ll try to get clearer on how I take Paul’s argument to work. If the argument succeeds

---

\*This paper is an edited version of a talk I gave as part of an Author Meets Critics session on L.A. Paul’s book, *Transformative Experience* at the 2015 Pacific APA. The paper was edited some in 2022.

<sup>†</sup> ✉ jane.friedman@nyu.edu

<sup>1</sup>All references to TE are to Paul (2014).

then I agree with Paul that ETEs make a sort of trouble we should be deeply worried about. In sections 3, 4, and 5 though I'll press on three key premises of her argument and suggest some ways that each can be denied, presenting a few off-ramps for those (like me) worried about Paul's conclusions.

## 2 THE ARGUMENT

Paul claims that ETEs pose a problem for standard accounts of rational decision-making, decision theory included. The rough thought is this: You can't know 'what it's like' to have an ETE — ETEs are, by definition, types of experiences you've never had before. Some decisions though are (in effect) decisions about whether to have an ETE or not (e.g. Should I try this new drug?). But if you can't know what it will be like to have the ETE, then how can you make a rational decision about whether or not to have it? Paul's thought is that in making potentially epistemically transformative decisions, we are saddled with a sort of ignorance about the (values of the) possible outcomes of our decisions that makes it difficult or even impossible for us to make those decisions rationally and authentically. But she thinks that these sorts of decisions are ones that we do and must make all the time and that 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 me try to flesh out my understanding of Paul's argument in a bit more detail. First, let's say that  $\Delta$  is a set of possible cases in which subjects are making decisions which have the following property: at least one possible outcome of each decision is epistemically transformative for the decider. When the set of possible outcomes of a decision has a member that is epistemically transformative, I'll call that outcome set *potentially epistemically transformative*. Second, Paul claims that a decision is rational only if it is based on or emerges from a set of rational preferences. From here I take her to make the following key claims,

- (**INF**) One's preferences over the possible outcomes of some decision are rational only if those preferences are all 'informed' — based on information about (the value of) those outcomes.
- (**STAN**) When the set of possible outcomes of one's decision is potentially epistemically transformative (i.e. for every case in  $\Delta$ ), the standard method for informing one's preferences with respect to that set of outcomes fails.

So, take a case in  $\Delta$ . You're deciding whether to  $\varphi$  or not. To make that decision you're considering the possible outcomes of doing  $\varphi$  and of not doing

$\varphi$  and considering which outcomes you'd prefer to which others. But (STAN) says that the standard way of informing your preferences over those possible outcomes won't work (given that at least one of those possible outcomes is epistemically transformative). This means (given (INF)) that you cannot arrive at rational preferences in this case if you try to use this standard decision-making method. And that means that if you stick to that standard method, you will not be able to act rationally in this case (nor will anyone in any of the cases in  $\Delta$ ).

Can we try to inform our preferences in a non-standard way and open up the possibility of acting rationally in the cases in  $\Delta$ ? Yes, but doing so raises a different problem. Here Paul claims,

**(NON-ST)** For every case in  $\Delta$ : Non-standard ways of informing one's preferences render those preferences inauthentic and alienate decision-makers from their own choices and futures.

This means that for Paul we can act rationally in the cases in  $\Delta$ , only if we act inauthentically and in a way that is troublingly alienating. So there is no way to act both rationally and authentically in the cases in  $\Delta$ . 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  $\Delta$  is a set of cases that involve truly central life decisions, as well as many everyday ones. We have all already made decisions like these and so part of what Paul is claiming is that perhaps none of us have ever made these decisions both rationally and authentically.

Let me say a bit more to clarify some key bits of these three claims ((INF), (STAN), (NON-ST)) and after that I'll discuss each claim in more detail, raising concerns about each.

#### ON INFORMED PREFERENCES AND STANDARD DECISION PROCEDURES

Paul doesn't talk about 'informed preferences' per se, but I think the term 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.<sup>2</sup> The basic thought is that informed preferences are preferences based on evidence.<sup>3</sup> For example, if Ben and Jerry's introduces two new flavours:

<sup>2</sup>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.

<sup>3</sup>Harsanyi (1996) thinks of informed preferences as the hypothetical preferences we would have if we had all the relevant information about the outcomes and made use of that information. But he also wants to use the term for actual preferences that 'agree with' these hypothetical ones. I'm not using the term to refer to something hypothetical but to characterize particular kinds of actual preferences.

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 (sufficiently) 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).

Paul’s claim is that rational decision and rational action require (sufficiently) informed preferences. 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 or to ground rational action and decision-making? More on that to come.

What about the ‘standard way’ of informing our preferences? At least in the sorts of cases in  $\Delta$ , Paul thinks that there’s a specific method that we want to use to inform our preferences. I’m going to call this method *imaginative forecasting*. When we use this method we imagine the various possible outcomes and what they’ll be like and then base our preferences on the results of this act of imagination. Paul describes this as ‘running a mental simulation of what it would be like’ (TE: 26). I’ll say more about imaginative forecasting in section 4. For now though, Paul’s claim is that the ‘natural’ and ‘ordinary’ way of informing our preferences in cases like the ones at issue goes by way of imagining what the outcomes will be like. And so the thought is that when some of those possible outcomes are epistemically transformative, that procedure will fail exactly because we have never had that sort of experience before and so can’t know what it will be like.

What I want to do now is go into a bit more detail about each of the three named claims from above, and raise a few concerns about each.

### 3 (INF)

(INF) says that one’s preferences over the possible outcomes of one’s decision are rational only if those preferences are informed preferences. This strikes me as a fairly plausible claim, but it’s certainly not an uncontroversial one.

First, even the claim that preferences are open to rational assessment is not obviously true. Certainly many of our preferences feel immovable in ways that also make them feel plausibly arational — just brute facts about our tastes. Are preferences really the sorts of things that can be rationally assessed at all? Recall Hume’s famous claim,

’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 choose 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 the latter. (Hume (2007): 2.3.3.6)

Let's call someone who says that preferences are not the sorts of things that can be rational or irrational at all — that they are simply outside the scope of rational assessment — a *Strict Humean* about preferences. We can then say that a *Moderate Humean* about preferences is, well, more moderate. In particular, our Moderate Humean agrees with the Strict Humean that preferences are not rationally evaluable individually, but disagrees with the Strict Humean by allowing for the rational evaluation of sets of preferences. Our Moderate Humean allows that improperly arranged preference sets — incoherent preference sets — can be irrational (and coherent ones rational). On this view, one's preference for *A* over *B* is never irrational on its own (nor ever rational on its own), but a set of preferences might be irrational if the set included (say) both a preference for *A* over *B* and a preference for *B* over *A*.

Paul's position on rational preferences is further away from Strict Humeanism than Moderate Humeanism is. Paul wants to say that individual preferences (e.g. a preference for *A* over *B*) are indeed rationally evaluable. And the basis for that evaluation is the extent to which those preferences are based on one's evidence — the extent to which those preferences are informed. Although more demanding than Moderate Humeanism, this stronger position has some intuitive appeal. Other things equal, there's something wrong with or at least strange about 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 Strict Humeanism about belief is plausible, but at least some think that epistemic rationality is a matter of belief coherence (this is a sort of Moderate Humeanism about belief). But this sort of coherentism feels as though 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 so has a coherent set of beliefs 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. 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 — matter. And I think it's plausible enough that similar claims could be made about preferences.

Moderate Humeanism is typically the way when it comes to decision theory though. For many, central elements of decision theory fall out of those coherence constraints on preferences.<sup>4</sup> Standard decision theory tells

---

<sup>4</sup>A standard account of decision theory (e.g. Savage (1954)) takes preferences over acts to be the basic or fundamental notion. A 'representation theorem' can then be used to show that subjects with well-ordered preferences can be 'represented' as having a unique

us what to do when we have (coherent) preferences over acts or outcomes and is 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 we ought to go about getting those beliefs and desires.

And a theory like this has a certain kind of appeal to be sure. If you strongly prefer  $A$  to  $B$ , even if that preference isn't properly based, there is some sense in which it's rational to act in accordance with that preference if it is relevant to the decision you face — it's your preference after all (it would be weird were you to act as if you preferred  $B$  to  $A$  instead). And if you do act in accordance with that preference, your behaviour will be explicable in its light. But while we have this very basic notion of rationality that is respected whenever someone acts on their (well-ordered) preferences, I think many of us, Paul included, think that there is more to rationality than that.

But part of what this means is that while (INF) may have some intuitive support, it doesn't have much decision theoretic support. And this means that as an objection to decision theory in particular, Paul's argument here is going to miss its mark some. If rational preferences for the decision theorist are simply well-ordered preferences, then they will reject (INF) and whatever problem Paul is pointing to for decision-making that goes by way of ETEs. Rational preferences needn't be informed preferences for the decision theorist, and so the fact that ETEs pose a problem for informing preferences is neither here nor there for them.

Of course, Paul's target isn't merely decision theory per se, but a more general notion of rational decision-making. Still, part of the claim I'm making here is general. Paul's account of rational decision-making is very demanding. Acting on your preferences can easily be irrational for her, even if those preferences cohere perfectly, and even if they are long-standing and deeply held. Choosing in line with the things you want most can be practically irrational for Paul. So while I'm sympathetic to some of Paul's more demanding constraints here, it is worth recognizing just how demanding they are.

#### 4 (STAN)

(STAN) says that when the possible outcomes of your decision are potentially epistemically transformative, the standard method for informing your preferences with respect to those outcomes fails. Here is my understanding of Paul's claim here. Say you're deciding whether to  $\varphi$  or not. Call the set of possible outcomes of this decision  $O_\varphi$ . And say that one of the outcomes in  $O_\varphi$  is epistemically transformative:  $o_e$ . You don't know what it will be like

---

credence function and a unique utility function. Probabilities and utilities are effectively constructed from preference sets.

to have the experiences that  $o_e$  involves. So if you try to (rationally) assign some subjective value to  $o_e$  by imagining what it will be like to have the experiences it involves, you won't be able to. But this means that this standard method can't help you (rationally) determine whether  $o_e$  is more or less valuable for you than any other outcome in  $O_\varphi$ . And that leaves your entire preference ordering for  $O_\varphi$  indeterminate or uninformed. In effect, if you try to use imaginative forecasting on  $O_\varphi$ , you'll fail to get a well-informed preference ordering since you won't be able to figure out where  $o_e$  should fit in that ordering. The top? Bottom? Somewhere in the middle?

At the heart of (STAN) is a claim about how the 'standard method' of informing our preferences fares when outcomes are potentially epistemically transformative. I've called this method 'imaginative forecasting', and here's a bit more about why Paul thinks we want to use it in the cases in  $\Delta$  and how it works. First, on the sorts of decisions subjects in  $\Delta$  are making she says,

According to that paradigm, we approach many major life decisions as personal matters where a central feature of what is at stake is what it will be like for us to experience the outcomes of our acts and where the subjective value we assign to an outcome depends upon what we care about, whatever that might be. (TE: 25)

And in order to properly assign these subjective values (thereby properly informing our preferences), we'll need to deliberate in a particular way,

A way to think about how we assess the subjective values of possible experiences is to represent ourselves as engaging in a kind of cognitive modeling. When you are considering your options, you evaluate each possible act and its experiential outcomes by imagining or running a mental simulation of what it would be like, should you act, for each relevant possible outcome of each relevant act. You simulate the relevant possible outcomes for yourself, that is, you simulate what it would be like for you to have each of these experiences. (TE: 26)

So we engage in imaginative forecasting when we imagine the various outcomes and what it would be like to have the experiences they involve, e.g. imagine what it would be like to taste Abracadabra and Zoolander ice cream. The claim in (STAN) is that we won't end up with properly informed preferences over a set of potentially epistemically transformative outcomes when we use imaginative forecasting. And, as I understand it, the reason is that imaginative forecasting cannot help us rationally assign values to the

transformative outcomes in particular and that leaves the preference sets in which those outcomes figure uninformed as well.<sup>5</sup>

Why does imaginative forecasting fail so dramatically when it comes to ETEs/transformative outcomes? As we've seen, the general idea is that given that we don't know what having the relevant ETE is like, we can't imagine what it will be like and so we can't reasonably assign value to that experience by trying to imagine what it will be like. In the next section I'll say more about whether we need to use imaginative forecasting at all in these cases, but in this section I want to press a bit on Paul's thought that the fact that some possible outcome involves having ETEs means that we cannot imagine what that outcome will be like for us. Even if imaginative forecasting won't work as well for ETEs as it does for types of experiences we have had, I don't think that it needs to fail as dramatically as Paul claims it does. My thought is that in most of the cases at issue we have a lot of information — phenomenal and otherwise — about what the outcomes will be like.

Of course Paul is right that when it comes to ETEs we are ignorant in particular ways. Let's call the knowledge or information that we're missing, *exact phenomenal information*. We don't have that sort of information until we've had the relevant type of experience: information about or knowledge of exactly what it's like to be on LSD, taste a new fruit, see a new colour, be blind, have a child, and so on. But how badly off our lack of exact phenomenal information leaves us depends on how much other evidence, information, and knowledge we have. And for most of the examples Paul turns to, typical subjects are in fairly good epistemic positions.

First, we tend to have a fair amount of phenomenal information about the relevant ETEs. We have typically had a range of experiences that are like the relevant ETEs in various ways. People who haven't had children have often spent time with babies or little kids, been dedicated caregivers, been in love and exhausted. And the same is true with respect to new jobs, new cities, new foods, new drugs, and so on. These are often importantly like old jobs, cities, foods, drugs, and so on. We don't know everything of course, but we are far from in the dark phenomenally — we typically have a range of similar experiences to draw on. In this respect, our epistemic blindspots are much less severe than Mary's blindspot was in her black-and-white room.

Moreover, we have plenty of other information about what the relevant outcomes and experiences will be like. We can know a lot about what life is going to be like when we have kids: we know about the moral implica-

---

<sup>5</sup>There are many good questions to ask about imaginative forecasting. Is it the sort of deliberative paradigm Paul claims? Why we should think it better than other methods for assigning values/utilities to outcomes? Why imagination rather than just plain thinking and reasoning? Some of these will be partially addressed in this piece, but detailed discussion will have to be set aside.



tions and financial costs, we've had testimony from friends and relatives, and consumed countless cultural products depicting parenthood in its many guises. On the whole then we have a lot of information (phenomenal and otherwise) about what the central ETEs at issue will be like. Why isn't it enough information? In particular, why can't we use all this data and run a reasonably extensive imaginative forecast, even without the kind of exact phenomenal information that can only come after having had the relevant type of experience? While I agree with Paul that some information about what transformative outcomes will be like for us is missing, the question is whether that specific lack is enough to render our imaginatively informed preferences irrational. And here I'm not yet convinced.

Indeed, there's a case to be made that in at least some of the cases at issue, exact phenomenal information shouldn't carry much deliberative weight at all. Yes, how I'm going to feel once I have a child matters to the value of that outcome for 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 can do enough to make that new person's life great, and so on. From a certain vantage point it becomes quite hard to see why my lack of knowledge about exactly how having a baby is going to feel for me should be so important in the face of all of these other considerations. And it's unclear why in the face of all of these considerations, plus the fact that I have a fair bit of evidence about what the outcome will be like for me, the lack of exact feeling information should prevent me from having sufficiently informed preferences even via imaginative forecasting.

(STAN) says that we can't imaginatively forecast our way to informed preferences in the relevant cases. And Paul may be right that imaginative forecasting doesn't get us as far as when we've already had the relevant types of experiences, but the question is whether it's far enough to render our preferences sufficiently informed. I haven't argued that it is, but perhaps one data point to think about is the subject who gathers all the information they can (and does this well), and feeds it into their imaginative forecast (and does this well, too). Then they use that forecast to make a decision. This strikes me as a fairly epistemically responsible subject, one who has done the work to inform their preferences and made a well-thought-through decision. It's not clear to me that we should want to deem this subject's behaviour, preferences, or decision irrational.

## 5 (NON-ST)

At this point one might think: if imaginative forecasting can't properly inform preferences in the cases in  $\Delta$ , then so much the worse for that method. We seem to have plenty of information available to us in many of those cases, so perhaps we can put that information to work in some other way and skip

running the imaginative forecast? And indeed, Paul thinks we can do this. She says,

Are there other ways to make transformative decisions that can meet the normative standard? Yes. One way to make such a decision is to dispense with subjective deliberation and a decision involving subjective values altogether... The idea is that, instead of trying to assess subjective values and use them in your deliberation, you make the decision based solely on the available empirical research, for example, research that compares life satisfaction and well-being for those who have children to those who do not. (TE: 85)

Paul thinks of this as a sort of ‘third-personal’ decision procedure (in contrast with the first-personal imaginative forecasting), and she seems to at least allow that it can be rational to make potentially epistemically transformative decisions this way as well. But she wants to say that we have reason to prefer imaginative forecasting nonetheless. Why? At least in the cases at issue, third-personal decision-making is ‘inauthentic’ and ‘alienating’. It is inauthentic and alienating because it involves largely ignoring our subjective values and deliberating just on the basis of empirical or statistical data. And about this sort of deliberation Paul says,

[I]f we eliminate the first personal perspective from our choice, we give up on authentically owning the decision, because we give up on making the decision for ourselves. We give up our autonomy if we don’t take our own reasons, values, and motives into account when we choose. To be forced to give up the first person perspective in order to be rational would mean that we were forced to engage in a form of self-denial in order to be rational agents. (TE: 130)

I agree with Paul that a decision procedure that forced us to entirely ignore our own reasons and values would be alienating and could render our decisions inauthentic. But I’m not sure that Paul’s third-personal deliberation is a decision procedure like this, nor that there aren’t other feasible procedures that combine first- and third-personal information. If the person who looks to empirical studies on the life satisfaction and well-being measures of people with and without children is going to act on that information, it’s exactly because it connects up in some way with what they already care about. In a Williamsian vernacular, it connects up with their subjective motivational set ( $\mathcal{S}$ ): their desires, projects (egoistic and non), patterns of emotional reaction, loyalties and more.<sup>6</sup> It’s only because of the elements

---

<sup>6</sup>See Williams (1979).

in the deliberator's  $\mathcal{S}$  that the empirical data counts in favour of or against having children for them. The same data can provide a reason to have a child for one decision-maker and provide a reason not to for another. So even the third-personal decision procedure Paul describes is first-personal to some extent and does involve taking one's own reasons, values, and motives into account.

More generally, the point of turning to the empirical data in the sorts of cases at issue is to inform our preferences over the possible outcomes of the decision we're facing. And if we manage to do that then the preferences we end up with are genuinely ours — elements of our subjective motivations sets. But from this perspective they don't need to be inauthentic at all. If I genuinely prefer having a child to not having 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 spending time with my nephew, or I got them based on the empirical data — they are my genuine preferences, and deciding based on them need not be inauthentic or alienating.

In suggesting possible rational decision procedures, Paul turns to what she takes to be a purely impersonal standpoint. Her third-personal method is meant to leave all first-personal subjective information to the side. I've raised doubts about whether her method really is as impersonal as that but there's also a looming question about other hybrid methods that combine third- and first-personal data. Why not use a method like that? Not only would using more first-personal and individualized information help make the relevant decisions more authentically our own, but wouldn't it be more epistemically reasonable or responsible as well? Shouldn't I use all of the evidence at my disposal rather than leaving some aside when deciding?

While Paul might agree that more first-personal data will help make our decisions less alienating, she doesn't think they will help make those decisions more reasonable. In contemplating the sorts of hybrid methods I'm suggesting here, she says, 'The only evidential basis for your assignments is that which you are justified in having, and that does not include phenomenal evidence...' (TE: 89). I take it that her thought here is that phenomenal evidence is necessarily corrosive in the cases at issue, and so attempting to rely on it at all will render our preferences irrational. This is obviously a very strong claim, too strong to my mind. Perhaps the phenomenal evidence you have is not strong enough to justify some preference ordering on its own, but weak phenomenal evidence is still some evidence and should in principle be amenable to supplementation. More generally, why would rationality demand I ignore some of the relevant evidence I have? Isn't that a very paradigm of epistemically bad behaviour? If that's right then not only could more first-personal data render my decision more authentic, but it could also render it more rational.

(non-st) says that other methods of informing our preferences in the

cases in  $\Delta$  are inauthentic and will leave us alienated from the decisions we make. This is true for Paul since she holds that the only method other than imaginative forecasting that can properly inform our preferences in the cases in  $\Delta$  is a maximally impersonal method — a method that asks us to ignore our own personal aims, goals, and desires. I’ve raised some doubts about how impersonal Paul’s impersonal method is in the end, but I also think other, hybrid methods for informing our preferences are in better shape than Paul allows.<sup>7</sup>

## 6 WHERE DOES THIS LEAVE US?

Paul thinks there is one way that we all want to (and need to if we want to be authentic) make potentially epistemically transformative decisions. But the fact that those decisions are potentially epistemically transformative means that we cannot rationally make them in this desired way. Since most of us have made many decisions of the relevant sort, Paul seems to be claiming that many of us have failed to make those decisions rationally. Or if we did manage to make them rationally, we were inauthentic or alienated from those decisions and the outcomes they brought about. This is a weighty conclusion and we are in a quite bad situation when it comes to these decisions (according to Paul).

Paul sometimes claims that it is decision theory that’s to blame for leaving us in this predicament. I don’t think that’s right. Decision theory 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, wherever that preference came from, decision theory says do it (all else equal of course).

That said, decision theory is very permissive in this way. And perhaps this fails to capture how we think we really ought to make many decisions. Decision theory typically makes no demand that our preferences be informed and places no constraints on which decision procedures we should use, effectively flattening a great deal of the complexity of the sorts of decisions Paul is interested in here. I take it that Paul is worried about this flattening and I’m sympathetic. Decision theory is highly idealized and as such is plausibly

---

<sup>7</sup>There is an additional wrinkle here as well in that Paul does think that in some cases in  $\Delta$  we can choose rationally and authentically but without going through any imaginative procedure. For instance, I can rationally and authentically decide not to get mauled by a shark without doing any imaginative forecasting, despite never having been mauled by a shark (or any animal) before. See TE pp.27-8, 127. Although Paul says that wants to leave this sort of case aside (she says we can know what it would be like to be mauled by a shark without doing the imaginative forecast), cases like this nonetheless open up some interesting space in which potentially epistemically transformative decisions can be made rationally and authentically without imaginative forecasting.

stripping away all sorts of important aspects of real-life decision-making.

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  $\Delta$ . 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.

Where are we left in the end? I think it depends in part on how bad one thinks the result that we can only choose rationally at the expense of choosing authentically 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. I've tried to press a little bit on some of Paul's key claims, pointing to some exit ramps.<sup>8</sup> In the end Paul and I disagree about how badly off ETEs leave us. But we don't disagree about how fruitful reflecting on their role in rational decision-making is. I hope this discussion has brought out some of the ways in which potentially epistemically transformative decisions make some of the deepest questions about rational decision-making strikingly salient: What does it take to have truly rational preferences and truly authentic ones? How should we go about assigning values to the outcomes we're deliberating about? How much should what some outcomes will feel like for us weigh in our decisions to bring those outcomes about? What should we do in the — entirely pervasive — situation in which we don't know everything we can about the consequences of our actions?\*

---

<sup>8</sup>Paul offers her own escape route, ultimately suggesting that we can restore rational and authentic decision-making in the cases in  $\Delta$  if we make decisions based on our preferences for having new experiences. So I can rationally and authentically decide to have a child based on my (imaginatively informed) preference for trying new things over maintaining the status quo. I don't find this solution all that appealing myself. Evidence gathering and rational inquiry can't make my decision authentic and rational, but my desire to have a new type of experience can?

\*Thanks to Laurie Paul, Harvey Lederman, and Amia Srinivasan for very helpful discussion.

## REFERENCES

- Harsanyi, John C. (1996). Utilities, preferences, and substantive goods. *Social Choice and Welfare*, 14:129–145. [3]
- Hume, David (2007). *A Treatise of Human Nature: A Critical Edition*. Oxford University Press. [5]
- Paul, L. A. (2014). *Transformative Experience*. Oxford University Press. [1]
- Savage, Leonard J. (1954). *The Foundations of Statistics*. Wiley. [5]
- Williams, Bernard (1979). Internal and external reasons. In Harrison, Ross, editor, *Rational Action*, pages 101–113. Cambridge University Press. [10]