

Expecting the Unexpected

Tom Dougherty, Sophie Horowitz and Paulina Sliwa¹

Forthcoming in Res Philosophica Special Issue on Transformative Experience in 2015. Penultimate draft. Please cite published version when available.

1. Introduction: Is becoming a parent rational?

Like most major life decisions, the decision whether to have children is fraught with uncertainty. How would your child turn out? What would your relationship with her be like? How much would you enjoy parenting? As such, it would seem a paradigm of a decision amenable to philosophers' favorite tool for making decisions under uncertainty—decision theory. Very roughly, according to decision theory, you should gauge how good or bad the possible outcomes of each option are, and weight these “utilities” by how likely you think these outcomes are, in order to calculate how much “expected utility” would result from each option. The decision-theoretic recommendation is that you choose the option with the highest expected utility. Since the purpose of this formal tool is to provide a lamp to guide us through the fogs of the future, we might have hoped that it would help with the decision about whether to have kids.

These would be false hopes, according to L. A. Paul, in an article that has perhaps had more impact outside of academia than any other philosophy essay in recent years (Paul 2015).² Paul argues that parenting decisions are intractable for standard decision theory. This would mean that insofar as we take decision theory to determine what it is rational to choose, we must conclude that it is neither rational to choose to become a parent nor rational to choose not to become a parent. The limits of reason have been reached, and any parenting decision would be a leap of faith.

These are bold and exciting claims. So what could justify them? In this essay, we will focus on a novel argument of Paul's that is based on the claim that becoming a parent is “epistemically transformative”: it gives one new knowledge of what it is like to be a parent and to have experiences related to parenting. These epistemic discoveries are made only upon entering parent-

¹ Listed alphabetically by surname. For helpful comments and discussions, we would like to thank Andy Egan, Brian Hedden, Chris Meacham, Laurie Paul, Josh Schechter, Miriam Schoenfeld, Eric Swanson, the audience of the 2014 Bellingham Summer Philosophy Conference, and anonymous reviewers for *Res Philosophica*. For epistemically transformative experiences during the process of writing the paper, we would like to thank Theodore Swift and Martha Nunenkamp.

² Media discussions include (Burkeman 2013; Gopnik 2013; Lombrozo 2013a, 2013b; Marshall 2013; Rothman 2013).

hood—too late to inform one’s decision to become a parent. On these grounds, Paul argues that a childless person cannot determine how desirable parenting outcomes would be. But without rationally determining the utility of each parenting outcome, this person cannot rationally calculate the expected utility of having a child. Hence, decision theory’s silence.

Although Paul’s primary focus is parenting, this argument is powerful enough to apply more generally to all decisions that determine whether one undergoes an epistemically transformative experience. It would show that we cannot rationally decide to undergo *any* new experience, from tasting Szechuan peppercorns to experiencing one’s first kiss. Could it really be that rational decision theory comes unstuck with all of these decisions?

Our answer will be that decision theory is posed no special problems by epistemically transformative experiences. To see this, we will draw a distinction between knowing *what it is like* to have an experience and rationally estimating *how valuable* that experience is. We agree that one cannot know in advance what it is like to have an epistemically transformative experience. But we disagree that someone cannot rationally estimate how valuable such an experience is. This is because direct experience is only one epistemic route to the value of experiences. Two of the other routes are testimony, and observing others’ behavior. Moreover, in many cases, we have experiences that are in some respects similar to epistemically transformative experiences. These resemblances can yield us partial knowledge of what an epistemically transformative experience is like. This partial knowledge is often enough for us to be able to rationally assign credences about how desirable we would find the experience—our third method of estimating its value. In this way, we will argue that a more nuanced account of the epistemology of value can provide a firmer foundation for decision theory as a theory of practical reason.

Before proceeding, let us clarify what our target is in this article. Since our interest is in Paul’s argument concerning epistemically transformative experiences, our primary focus will be on her first work on epistemic transformation, which is published in the special volume of *Res Philosophica* but has been available online in its finalized typeset form since at least 2013. This argument is novel, but, we argue, fallacious. Having addressed this argument, we turn to subsequent work of Paul’s, which develops a more nuanced overall argument (Paul 2014). This argument is more successful in posing a challenge to decision theory, but only because it appeals to familiar problems for decision theory. The epistemically transformative nature of these experiences does no special work.

2. The challenge of epistemic transformation

Since Paul's argument focuses on discovering what it is like to have a new experience, we will frame our discussion around one of philosophy's most famous characters, who also appears in Paul's essay: Frank Jackson's color scientist Mary. As you will recall, Mary "is confined to a black-and-white room, [and] is educated through black-and-white books and through lectures relayed on black-and-white television" (Jackson 1986, p. 291). We will assume that Mary has survived this social isolation psychologically unscathed by being provided with an ample supply of literature. This has nourished her imagination and allowed her to build up hopes of a future life outside her prison. One of the things she wonders about is whether to become a parent.

From reading glossy black-and-white magazines, Mary has discovered what Paul describes as our culture's "ordinary" way of making a decision whether to have a child: in order to decide whether or not to have a child, someone should consider what the experience of being a parent would be like and consequently "carefully weigh the value of ...[these] future experiences." (Paul 2015, p.2). Paul later characterizes the values of these future experiences as "centering on...the subjective value of *what it is like* to be the person who made the choice" (p.4, emphasis added). To decide whether or not to have a child is thus a choice between different "phenomenal outcomes that involve what it's like for her to have her own child" (p.4). This way of making parenting decisions is Paul's first target.

From reading less glossy black-and-white scholarly tomes, Mary has also discovered a way to formalize our ordinary decision-making: decision theory. This is Paul's main target, which will be of particular interest to philosophers, and hence the one that we will primarily focus on. She describes it as follows:

To make a choice rationally, we first determine the possible outcomes of each act we might perform. After we have the space of possible outcomes, we determine the value (or utility) of each outcome, and determine the probability of each outcome's occurring given the performance of the act. We then calculate the expected value of each outcome by multiplying the value of the outcome by its probability, and choose to perform the act with the outcome or outcomes with the highest overall expected value. (Paul 2015, p.3)

When this approach is applied to the decision whether to become a parent, one must assign utilities and probabilities to each possible outcome that would result both from becoming a parent and from remaining childless.³

Mary cannot wait to apply these approaches to her parenting decision! But then she picks up the Wall Street Journal, and reads an article reporting the bad news from Paul.⁴ According to Paul, these approaches are useless to Mary because they direct her to focus her decision-making on phenomenal outcomes, and yet Mary is phenomenally impoverished. Paul illustrates this point in terms of the epistemically transformative experience of seeing red for the first time:

For our purposes, Mary's impoverished epistemic situation means, first, that since Mary doesn't know how it'll phenomenally feel to see red before she sees it, she also doesn't know what emotions, beliefs, desires, and dispositions will be caused by what it's like for her to see red. Maybe she'll feel joy and elation. Or maybe she'll feel fear and despair. And so on. Second, because she doesn't know what emotions, beliefs, desires, and dispositions will be caused by her experience of seeing red, she doesn't know what it'll be like to have the set of emotions, beliefs, desires, and dispositions that are caused by her experience of seeing red, simply because she has no guide to which set she'll actually have. And third: she doesn't know what it'll be like to have any of the phenomenal-redness-involving emotions, beliefs, desires, and dispositions that will be caused by her experience of seeing red. Even if she could somehow know that she'll feel joy upon seeing red, she doesn't know what it will be like to feel-joy-while-seeing-redness until she has the experience of seeing red. And these are all ways of saying that, before she leaves her cell, she cannot know the value of what it'll be like for her to see red. (Paul 2015, p. 7).⁵

Similarly, since Mary does not know what it is like to be a parent, Paul would argue, she cannot rationally place a value on becoming a parent. But if she cannot rationally assign utilities to parenting outcomes, then decision theory cannot guide her choice.

³ Moreover, standard decision theory assumes that an agent's preferences are complete: for any two outcomes, she is indifferent between these outcomes, or strictly prefers one to the other. Further, it assumes that an agent's preferences do not form a cycle, e.g. it is not the case that an agent has intransitive preferences by preferring A to B, B to C, and preferring C to A.

⁴ Gopnik, "Is it possible?"

⁵ See also (Paul 2015, pp. 10, 11, 12, 13 and 15).

Although Mary is alone in her room, she is not alone in facing Paul's problem. The same considerations apply to anyone who is deciding whether to undergo an epistemically transformative experience. So in its generalized form, we can summarize Paul's argument as follows:

Premise 1: There is a certain class of life decisions, including parenting decisions, in which an agent is deciding whether to perform an action that has some chance of resulting in an outcome in which she has a phenomenal experience that would be epistemically transformative for her.

Premise 2: If a phenomenal experience would be epistemically transformative for an agent, then she does not antecedently know what the experience would be like.

Premise 3: If an agent does not know what it is like to have an experience, and this experience is constitutive of a "phenomenal outcome," then she cannot rationally judge the subjective value of this outcome for her.⁶

Premise 4: If an agent cannot rationally judge the subjective value of a phenomenal outcome for her, then she cannot rationally choose between options when one of these options would lead to this phenomenal outcome.

Conclusion: Therefore, there is a certain class of life decisions, including parenting decisions, in which an agent cannot rationally decide what to do.

This formulation is broad enough to apply to both of Paul's targets. To target the argument specifically at decision theory, we could specify that in Premises 3 and 4 the talk of judging a value's outcome should be understood in terms of talk of assigning utilities to this outcome.

Is this a sound argument? Premise 2 is true by the definition of "epistemically transformative,"⁷ Premise 4 is highly plausible, and it is trivial to see that the argument is valid. Thus, there are two premises that deserve further investigation—Premises 1 and 3. We will proceed to discuss each, organizing our discussion in terms of increasing importance. We will start with preliminary remarks concerning Premise 1. We will then offer our central criticism of the argument, arguing that we should reject Premise 3. After that, we will

⁶ Thanks to Paul for guidance on how to formulate this premise.

⁷ Though as we note later, we think it is illuminating to draw a distinction between having complete knowledge and partial knowledge of what it is like to have an experience. In light of what we go on to say, Premise 2 is only true when it is read as concerning complete knowledge.

offer a diagnosis of why the conclusion might have seemed plausible, by noting familiar problems that arise for decision theory in these contexts.

3. The broad scope of Premise 1: The pervasiveness of epistemic transformation

We will begin by noting how Paul's argument applies to her main target: decision theory. When applying decision theory to a decision, an agent needs to consider whether she has any credence that an option results in outcomes that involve her having certain phenomenal experiences. If she has some credence that these outcomes will obtain, then she will need to consider how much utility to assign to these outcomes. But if she cannot rationally assign utilities, then she cannot rationally provide herself with the inputs necessary for the decision-theoretic cogs to start grinding.

This feature of Paul's argument means that it has much wider scope than it might at first seem. The argument does not simply concern decisions to become a parent. It also concerns any decision that an agent thinks might lead at some point to her becoming a parent. Suppose Mary has escaped from her colorless prison and gets asked on a date for the first time. If she has some credence, however small, that accepting the invitation will one day lead her into parenthood, then decision theory requires her to assign a utility to the outcomes in which she becomes a parent. This point becomes even more pressing when we consider how many epistemically transformative experience there are: seeing red, tasting a durian fruit, flying in an airplane, falling in love, falling out of love, suffering the ennui of a mid-life crisis, grieving over a loved one's death, climbing a mountain, riding a roller-coaster, fighting in combat, and so on. These are all experiences that are foreign to Mary. Insofar as Mary has some credence that leaving her monochrome prison may result in her undergoing one such experience, Paul's argument would imply that she cannot rely on decision theory to rationally decide whether or not to escape. And this point does not concern just poor Mary. For almost any practical decision we make, we should have some credence that one of our options will bring about an outcome in which we have at least one epistemically transformative experience. Thus, if sound, Paul's argument would show that we cannot appropriately assign a utility to this outcome, and that hence decision theory is stymied. So the argument does not just threaten decision theory's application to parenting decisions. It threatens its application to almost any decision at all.

4. Rejecting Premise 3: The epistemology of the value of experiences

Looking more closely at Premise 1 showed that Paul's argument has considerably more breadth than one might first have thought. Before accepting

such a revisionary conclusion, we should examine the argument's crucial step: Premise 3.

Premise 3: If an agent does not know what it is like to have an experience, and this experience is constitutive of a “phenomenal outcome,” then she cannot rationally judge the subjective value of this outcome for her.

In this premise we move from descriptive uncertainty about what a phenomenal outcome is like to evaluative uncertainty about the value of that phenomenal outcome. In support, Paul argues that “the relevant values are determined by what it is like for you to have your child.” Consequently, when deciding whether to have children, “the value of your act...depends largely on the phenomenal character of the mental states that result from it” (Paul 2015, p.5).

But while there is plausibility to the claim that the phenomenal character of an experience is typically relevant to the value of the experience, we should still distinguish the experience's phenomenal character from the experience's value. This is because agents might differ in their attitudes towards the same phenomenal experience. For example, Mary may prefer the taste of sugar to the taste of salt, while her prison guard has the opposite preference. So, the experience of tasting sugar may be more valuable to Mary than her guard. Drawing this distinction allows us to also draw an *epistemological* distinction between awareness of an experience's phenomenal character and awareness of its value.

Once drawn, this epistemological distinction should make us suspicious of Premise 3. From the fact that an experience is epistemically transformative, it only follows that the agent is not antecedently in a position to know what the experience would be like. This is consistent with the agent being able to rationally estimate the experience's value. If you have not given birth, then you do not know what it is like to have the experience of prolonged labor. If you have not experienced a year of isolation in a super-max prison, then you do not know what it is like to be deprived of all human contact for an extended period of time. If so, these experiences would be epistemically transformative. But without having undergone these experiences, you can still judge the intrinsic value of the phenomenal aspect of these experiences.⁸ (Hint: they contain intrinsic disvalue.) The same holds for positive experiences. Given the limited dating opportunities in her prison, Mary does not

⁸ This is consistent with thinking that the experience has *extrinsic* value, e.g. because the labor instrumentally leads to the birth of one's valued child. We focus on the intrinsic value of epistemically transformative experiences, given that it is this type of value that Paul claims one cannot know. Our conception of intrinsic value follows that of Rae Langton. Langton holds that something's intrinsic value is the value that something has “in itself” which we take to be equivalent to the value it has in virtue of its intrinsic properties (Langton 2007).

know what it is like to fall in love with someone who reciprocates her feelings. Nevertheless, her literary window on the world could enable Mary to rationally estimate the intrinsic value of this experience.

So how can Mary rationally estimate the value of epistemically transformative experiences? What kind of evidence could she have? In fact, there is not one single source of relevant evidence. There are at least three. We will illustrate each with examples of epistemically transformative experiences.

4.1 The method of receiving testimony: The Mystery Closet

The first source of evidence concerning the value of an epistemically transformative experience is testimony:

Mystery Closet. From a flyer, Mary learns that the funfair is in town outside her prison walls. She reads that one of its attractions is the “Mystery Closet,” in which customers undergo an experience. The experience is incredibly rare, and so almost certainly customers will not have had the experience before. Out of thousands of customers, every single one has said that they greatly valued the experience. Even better, the organizers have made it free to enter the closet, hoping that it will draw people into attending the funfair.

Easy question: can inexperienced Mary rationally estimate whether the Mystery Closet experience would have intrinsic value for her? The answer is obvious: yes she can. The evaluative testimony of the other customers has given her excellent evidence that the epistemically transformative experience would be a valuable one. Another easy question: if Mary could enter the Mystery Closet, should she choose to do so? The answer is yes, again: it would be rational for her to choose to enter it, given her evidence about the value of the experience.

As well as stylized examples, there are real world examples of uniformly positive testimony. The most obvious examples involve extreme pleasure or pain. We can put our hands on our hearts and say that we do not know what it is like to be high on heroin or crack cocaine. And yet we are still able to rationally assign credences about whether we would find intrinsic value in these experiences. Similarly, we are fortunate enough not to know what extreme torture is like. All the same, we are able to rationally estimate whether we would disvalue this experience. One reason why we are able to do so is that other people have had these experiences, and have testified as to whether these are valuable or not. Our estimates of these experiences’ future value can then rationally guide our actions. If we were given a choice as to whether to undergo torture for a couple of dollars of reward or forgo both torture and reward, it would be rational for us to choose the latter.

These are examples of uniform testimony. But more commonly, testimony will be mixed. Consider:

Durian (Simplified). Mary reads that 50% of people who eat durian say they quite like the taste, but the other 50% say that they find it slightly nauseating.

In this case, Mary's credence as to whether she would enjoy the experience of tasting a durian should be split: she should assign 0.5 credence to the possibility that she would find value in eating durian, and 0.5 credence to the possibility that she would not. Assuming that the intrinsic value of the gustatory experience for someone depends primarily on this person's enjoyment of this experience, this gives Mary split evidence about whether the experience of tasting durian would be valuable for her. Again, this evaluative evidence can guide her actions. It is plausible to think that if Mary is risk averse, then she rationally ought not eat durian, whereas if she is risk loving, then it is rational for her to eat it.

Paul is more pessimistic about the possibility of learning from testimony in these cases, but this pessimism is based on considering only the idea that testimony cannot tell us *what an epistemically transformative experience is like*.⁹ This idea is undoubtedly correct; the hallmark of epistemically transformative experiences is that we cannot fully know what they will be like, by testimony or any other means. But all the same, testimony can tell us how valuable an experience is. Paul does indirectly tackle this evaluative testimony when she discusses the evidence provided by survey data about how satisfied parents are (Paul 2015, pp. 17-20). Paul's central response is that this evidence might only provide an agent with "external" evidence about whether parenting would maximize utility for her, but that it is irrational to choose to maximize utility instead of consulting her "subjective...phenomenal preferences."¹⁰ Paul writes:

Imagine Sally, who has always thought that having a child would bring her happiness, deciding not to have a child simply because she knows not having one will maximize her utility. For

⁹ "Perhaps you think that you can know what it's like to have a child, even though you've never had one, because you can read or listen to the testimony of what it was like for others. You are wrong." (Paul 2015, p. 12) In personal communication, Paul agrees that in cases like Mystery Closet and Durian (Simplified) an agent can be rational in accepting evaluative testimony. Nevertheless, she argues that epistemically transformative experiences pose a special problem: in many cases involving epistemically transformative experiences people vary widely as to which value they assign to a particular phenomenal outcome. We agree that when there is such variation, relying on evaluative testimony is more problematic. But we are skeptical that the difficulty here has to do with those experiences being epistemically transformative. We discuss this in more detail in section 5.

¹⁰ Paul also raises the worry that this evidence is not enough to go on. We respond to this when we discuss sparse or messy evidence in Section 5.

her to choose this way, ignoring her subjective preferences and relying solely on external reasons seems bizarre... Now consider Anne, who has always thought that having a child would bring her misery, deciding to have a child simply because she knows it will maximize her utility. Again, the decision procedure seems bizarre from our ordinary perspective. Choosing rationally requires a very different way of thinking about the decision than we ordinarily think it does—to be rational, we have to ignore our phenomenal preferences. (Paul 2015 p. 19)

Unlike Paul, we do not find Anne's behavior bizarre at all. It seems that, like Sally, she has simply revised her earlier beliefs about how good it would be for her to be a parent, in light of new evidence about other parents' happiness. As Paul notes, Anne has received evidence that parenting would maximize utility for her. Anne's utility of course depends on the satisfaction of her preferences. So Anne has received evidence that her preferences will be satisfied by parenting. If we assume, with Paul, that parenting happiness depends on the satisfaction of phenomenal preferences, then Anne has received evidence that her phenomenal preferences will be satisfied. Moreover, insofar as these phenomenal preferences are Anne's own preferences, she has received evidence that her subjective preferences will be satisfied. Therefore, we conclude that external testimony can provide Anne with evidence about how parenting would satisfy her "subjective, phenomenal preferences." The dichotomy between consulting subjective preferences and relying on external reasons is a false one.

It may be helpful in this respect to recall the example of the Mystery Closet, in which customers have novel experiences. This is a paradigm case where prospective customers should care about whether the experience will satisfy their phenomenal preferences. Moreover, since they are making these decisions self-interestedly, they should consult their subjective preferences. Of course, the testimony of previous customers provides them with excellent evidence that they will be glad they went in the closet. In this way, testimony can provide them with external evidence that their subjective phenomenal preferences would be satisfied. Therefore, even if we should make these decisions on the basis of subjective phenomenal preferences, then this consideration is not a good reason for turning our back on evaluative testimony.

It is of course true that, by using testimony, someone is not using first-personal imaginative projection to learn about the satisfaction of her subjective, phenomenal preferences. But our point is that nonetheless the testimony does allow her to learn about the satisfaction of her subjective, phenomenal preferences. First-personal imaginative projection is not the only epistemic route available.

Should we worry that relying on testimony as evidence about the satisfaction of phenomenal preferences would be an “inauthentic” way of making decisions?¹¹ It is hard to say in the abstract, without a developed account of what authenticity of choice involves, but we suggest not. It may be plausible that authenticity requires one to aim at the satisfaction of one’s own preferences (including subjective phenomenal preferences). But we see no intuitive case for thinking that authenticity constrains how one should acquire evidence about how one’s own preferences would be satisfied. After all, it would not be inauthentic for someone to choose to enter the Mystery Closet on the basis of testimony, provided that this testimony had bearing on whether the Mystery Closet would satisfy her own preferences. This seems to us no less true in cases where the stakes are very high, or where the testimonial evidence is messy or inconclusive. As we discuss in more detail later, mixed evidence would make the choice more risky. But as a general point there seems nothing inauthentic about making gambles, when one does so on the basis of how likely, and to what extent, one’s own preferences will be satisfied. We suspect the temptation for thinking that there is a tension between authenticity and testimony-based deliberation comes from running together two ways in which deliberation might be “first-personal.” Deliberation might be first-personal in either of two ways: in the sense that it involves imaginative projection concerning what it is like to have experiences, and in the sense that it aims at the satisfaction of one’s own desires. The first type of first-personal deliberation may pose problems for testimonial evidence, but this type of deliberation has no connection to authenticity. There is plausibly a connection between authenticity and the second type of first-personal deliberation, but this is a type of deliberation that we can conduct on the basis of testimony. Distinguishing these two senses of “first-personal” therefore removes the temptation towards thinking there is a conflict between authenticity and testimony.

4.2 The method of observation: The dog on the beach

Testimony is not our only source of evidence about the value of others’ experiences. Often, this value is revealed in their behavior. This is what makes it possible for us to discover whether speechless animals are having valuable experiences:

Dog on the Beach. Sparky bounds up and down the sand. He dives into the sea to retrieve a tennis ball, before returning to the shore where he vigorously shakes himself dry. He meets a

¹¹ Paul raises considerations of authenticity in her (2014).

new dog, whom he gives a good sniff, and then chases a seagull, with abandon but not success. Throughout, Sparky's eyes are bright, and his ears are perky; he is jumping up and down, his body is wiggling and his tail is wagging.

It does not take a dog-whisperer to realize that Sparky is a happy dog, who is greatly enjoying his experiences on the beach. We know that his behavior indicates that his experiences contain intrinsic value. This is the case even though we do not know what it is like to have these canine experiences—no more than we know what it is like to be a bat (Nagel, 1974).

The same is true of our fellow human animals. We can observe people's facial expressions, their body language and other forms of their bodily behavior. On this basis, we can discover whether their experiences have intrinsic value. Moreover, we can do so even when we ourselves have not had these experiences. Suppose Mary watches footage of a drunk person who is smiling, laughing and uncharacteristically telling her friends how much she loves them. As a lifelong teetotaler, Mary does not know what this person's inebriated experiences are like, but she can tell that the drunk is having a pleasurable experience. Alternatively, suppose Mary observes someone suffering from clinical depression, who is eating less, sleeping less, and is removing herself from social engagements. Even if Mary does not know what the experience of severe depression is like, she can still infer that this experience does not contain intrinsic value. By using her knowledge of the value of others' experience, Mary can make inferences about how much value that experience would have for her. In this way, observation provides Mary with evidence with which to rationally estimate the value for her of these epistemically transformative experiences.

4.3 The method of inference from similar experiences: Vegemite

So far, we have argued that we can rationally estimate the value of an epistemically transformative experience by considering how much value this type of experience has for others. But we often also have specific evidence bearing on what our own personal preferences are likely to be. Experiences fall into broader kinds. If someone has had some experiences that are members of a kind, then she can inductively come to know something about what the other members of this kind are like. Thus, our third source of evidence regarding the value of an epistemically transformative experience is to consider its resemblance to other experiences that we have had.

To illustrate this point, let us consider an example that Paul takes from David Lewis. According to Lewis, you cannot come to know what it is like to taste Vegemite without actually having tasted it:

If you want to know what some new and different experience is like, you can learn it by going out and really having that experience. You can't learn it by being told about the experience, however thorough your lessons might be. . . . You may have tasted Vegemite, that famous Australian substance; and I never have. So you may know what it's like to taste Vegemite. I don't, and unless I taste Vegemite (what, and spoil a good example!) I never will. (Lewis 1990, p. 292)

Quoting this passage, Paul endorses Lewis's claim that tasting Vegemite for the first time is epistemically transformative. Since it is transformative, she argues, we cannot rationally assign a value to tasting Vegemite.

But this overlooks the fact that even if we cannot have complete knowledge of the phenomenal feel of an epistemically transformative experience in advance, we can still have partial knowledge of this. This partial knowledge can be a basis on which to rationally estimate the value of the Vegemite-tasting experience. For example, Mary can read that the experience of tasting Vegemite is an experience of tasting something intensely salty and savory. This testimony is enough for Mary to know that tasting Vegemite has some similarity to the experience of tasting soy sauce, parmesan, or anchovies.¹² If Mary has been revolted every time that she ate intensely salty and savory foods, then tasting Vegemite is unlikely to be an intrinsically valuable experience for her. More generally, awareness of these resemblances and of one's preferences can provide a guide to whether a new experience would be valuable.

As with testimony, Paul does consider resemblances:

Being around other people's children isn't enough to learn about what it will be like in your own case. The resemblance simply isn't close enough in the relevant respects. (Paul 2015, p. 13)

And this is plausible, so far as it goes. Arguably, one cannot fully appreciate what it is like to be a parent by being around other people's children. But even so, we can have partial knowledge of what this is like. In turn, this partial knowledge can provide a rational guide for our estimates concerning the value that parenting would have for us. Suppose a childless kindergarten teacher takes great pleasure in being around children, caring for them, and

¹² Since testimony of qualitative resemblances is different from evaluative testimony, the third epistemic method of making inferences from similar experiences is distinct from our first epistemic method of receiving evaluative testimony. Receiving qualitative testimony that Vegemite is intensely salty does not by itself allow one to estimate the value of eating Vegemite. By contrast, receiving evaluative testimony that torture is intensely disvaluable does allow one to estimate torture's value.

seeing them develop and flourish, and does not particularly mind the associated unpleasant tasks. This person clearly has some grounds on which to assign credence to how much he would value the experience of parenthood.

Indeed, if we filled in the details of the Mary case in the right way, we might even imagine that Mary is able to make predictions along these lines about her experience of seeing red. Because this would be an epistemically transformative experience, Paul argues that Mary cannot know whether she would value it (Paul 2015, p. 14). But we can imagine the case in such a way that it is plausible that Mary can justifiably have high credence that she would value it. Suppose that Mary's aesthetic sensibility is heavily biased towards finding sights beautiful; she finds value even in sights that are not conventionally beautiful. Further, Mary burns with a deep yearning to understand all aspects of the human experience—she wants to feel what others feel, as she values the insight this brings her of their lives. Moreover, Mary's curiosity knows no bounds; she is an adventurous sort who loves novelty for its own sake, and is never ruffled by the exotic. Now, consider the fact that seeing red for the first time is a member of the kinds, “visual experience,” “experience that has been had by many other humans,” and of course, “epistemically transformative experience.” In light of this fact, if Mary is aware of her aesthetic sensitivity, her interest in other humans and her yen for the new, then she is in a position to rationally estimate the value of seeing red for the first time.

4.4 Summary: Why we should reject Premise 3

In light of these considerations, we conclude that there are counterexamples to Premise 3:

Premise 3: If an agent does not know what it is like to have an experience, and this experience is constitutive of a “phenomenal outcome,” then she cannot rationally judge the subjective value of this outcome for her.

The plausibility of Premise 3 relies on a restricted view about what counts as the admissible evidence concerning the value of experiences: the premise is true only if, as Paul suggests, the only admissible evidence is complete knowledge of the phenomenal character of the experience. Our arguments in this section aimed to show that this restricted view is false. We can use testimony, behavioral observation and inference from similar experiences to rationally estimate the value of new experiences.

5. Familiar epistemic problems for would-be parents

We have argued that Paul's argument fails: from the fact that an experience is epistemically transformative, it does not follow that one cannot make a

rational decision about whether to undergo it. In more recent work, Paul offers a more restricted version of the argument. Paul has narrowed her interest to high-stakes cases.¹³ In high-stakes cases, the transformative experiences that purportedly create trouble for decision theory involve not just phenomenal ignorance, but also conflicting and inconclusive testimony about what it is like to undergo them, as well as changes in the agent's core preferences. In this section, we agree that in these more restricted cases it may well be tricky to employ decision-theoretic reasoning to guide one's decision. This, however, can be traced back to some familiar, and more general, challenges for epistemology and decision theory. The fact that these experiences are epistemically transformative is irrelevant. At the same time we offer an alternative explanation of why Paul's original argument may have seemed compelling. We will start by discussing problems that arise from the kind of evidence that we have available when making life-changing decisions. We will then discuss problems raised by preferences in life-changing decisions.

At several points in her discussion, Paul emphasizes how hard it is to know what one's future experience is like. She characterizes this problem as one of qualitative ignorance:

Qualia-Ignorance: Of one specific experience, not knowing what it is like to have this experience.

After all, this is why epistemically transformative experiences are meant to pose a special problem for decision-making.

But at key points, Paul also appeals to another type of ignorance.¹⁴ Recall her discussion of Mary seeing red, quoted here in full in section 2.¹⁵ The intuition elicited by this discussion is that Mary's ignorance leaves her unable to assign a value to her experience. But why? We suggest that the main part of the explanation is that Mary is unsure whether her experience would be a frightening experience, a stressful experience, a satisfying experience, and so on.¹⁶ This is simply an instance of a more general type of ignorance:

¹³ Paul clarifies this in her comments on this essay at the 2014 Bellingham Summer Philosophy Conference. Similarly, in (Paul 2014, p. 18), she focuses on "decisions about whether to undergo an experience that will change your life in a significant new way." We take this to be a refinement of her earlier argument in (Paul 2015), the scope of which more broadly included low-stakes decisions to see red for the first time or to taste vegemite for the first time.

¹⁴ See (Paul 2015 pp.7, 9 & 14).

¹⁵ (Paul 2015, p. 7). See also (Paul 2015, pp. 11, 12, 13 & 15).

¹⁶ We pass over a more minor point in the quoted passage where Paul notes that Mary "doesn't know what it will be like to feel-joy-while-seeing-redness until she has the experience of seeing red." We find this consideration to have no intuitive appeal: it should be clear to Mary that feeling-joy-while-seeing-redness will have positive value for her.

Which-Ignorance: Of many specific experiences, not knowing which of these experiences one will undergo.

Which-ignorance is independent of qualia-ignorance.¹⁷

In her more recent work, it also looks as if Paul appeals to which-ignorance as posing difficulties for decision-theoretic reasoning. In discussing the transformative choice of becoming a vampire, Paul argues:

What if it turns out, given your delicate sensibilities, that once you've transformed, you can't stand chicken blood—all you'll want to drink is human blood, in particular, the blood of male virgins. (One of your vampire friends confides that he is actually quite finicky now that his palate has been educated about platelet terroir.) But contemporary vampire society frowns on drinking human blood, since it isn't good for public relations. And so, if you become a vampire, for the foreseeable future, you'd have to eat food that absolutely disgusts you, and you'd have to constantly confront and overcome your repulsive urge to attack innocent little boys. [...] The problem here is that you can't predict how your preferences will change. Something that seems disgusting now might seem preferable to the finest of wines once you've been vampirically rewired. (Paul 2014, p. 45)

Again, the problem Paul points to is that you cannot know which preferences you will acquire once you turn into a vampire. And so you cannot know which experiences you will be having: one of relishing chicken blood or one of being disgusted by it?

It is important to note that which-ignorance by itself poses no problem at all for decision theory. In fact, it is exactly this kind of ignorance that gives decision theory its purpose. Decision theory is a formal tool for acting when one is unsure about the causal consequences of various options; it guides these choices in light of one's credences in these causal consequences obtaining. Decision theory does not guide our actions by assuming we have knowledge of the actual utility of the outcomes that will in fact obtain as the result of our actions. Instead, it guides us to perform the actions that have the highest expected utility, which is based on how likely we consider various outcomes to obtain. All we need in such situations is the ability to assign rational credences to various outcomes' obtaining, and to assign utilities to

¹⁷ We can see this by noting two points. First, there can be which-ignorance without qualia-ignorance: when the sky is gray, one can be unsure whether one will undergo the familiar experience of walking home in the rain or another familiar experience of walking home dry. Second, there can be qualia-ignorance without which-ignorance. Suppose that there is a single qualitative experience corresponding to what it is like to be a bat using echolocation to find an insect (Nagel 1984). Since it is a single experience, we do not have which-ignorance about it, but we do have qualia-ignorance about it.

those outcomes. And, as we have argued, this is something we can do when making decisions regarding epistemically transformative experiences.

But situations involving which-ignorance may prove tricky for decision theory in other ways. To apply decision theory we need to be in a position to assign rational credences to various outcomes of our action. But our world is often extremely epistemically uncooperative. For one, it is often ungenerous with the evidence that it provides us. Paul brings this out when discussing the possibility of making use of survey data about other parents' happiness in order to inform our decisions about whether to have children. One of her objections is that "[t]here just isn't enough evidence available to support this sort of reasoning"; so, we should "hold off on deciding, due to lack of conclusive evidence" (Paul 2015, p. 19). Similarly, Paul argues that if "we assign values and credences based on insufficient evidence, and calculate the expected value of our acts using such assignments, our decision does not meet the normative standard for rationality" (2014, p. 23). In addition, the world sometimes provides us with different pieces of evidence that are so messy that it is unclear what the evidence supports. Even if we have plenty of survey data and detailed testimony from many friends who are parents, how should we evaluate this evidence to form our overall credences? As Paul (2014, p. 28) points out, this is particularly a problem when agents vary widely as to which value they assign to the outcome in question.

In light of these challenges, we might say that there is a problem of sparse or messy evidence: either the evidence is too sparse to support any rational assessment at all, or the evidence is too messy to support the type of reasoning required for the precision of decision theory.

How to use sparse or messy evidence to form credences is a challenging problem for epistemology. To illustrate, consider the following case of Adam Elga's:

Stranger. A stranger approaches you on the street and starts pulling out objects from a bag. The first three objects he pulls out are a regular-sized tube of toothpaste, a live jellyfish, and a travel-sized tube of toothpaste. To what degree should you believe that the next object he pulls out will be another tube of toothpaste? (Elga 2010, p.1)

This case nicely illustrates the difficulty of assigning credences when our evidence is sparse and messy: you have not got much to go on, and it is unclear how to put together the scant pieces of evidence that you have. In these respects, Elga's case is similar to the kinds of evidential situations that we often find ourselves in when considering epistemically transformative experiences like becoming a parent. We have observed our friends and others be-

coming parents (or not). We have heard or read all kinds of testimony. But how should we weigh all of this evidence together?

Indeed, this problem is particularly likely to arise with the epistemic methods that we discussed earlier. Take testimony. Our previous durian example was artificially simplistic. A more realistic variant would be:

Durian (Complex). Mary reads in the Lonely Planet Guide to Asian Fruits that many people consider durian a delicacy, while a minority find the taste disgusting. Her internet pen pal says that he considers it the “king of fruits.” Her prison guards say that it is not such a big deal either way. Knowing this, Mary is deciding whether to eat a durian for the first time on her release.

In this case, it is much harder for Mary to estimate how much value she would get from eating durian. One problem is that it is hard to tell how much value is derived from tasting “a delicacy” or the “king of fruits.” But more pressingly, it is hard for her to estimate how likely it is that her experience will be like that of the majority or that of the minority. How many people were consulted by the Lonely Planet before it judged what the majority and minority preferences were? And just how major is the majority: 90%? 70% 50.01%? While this is in doubt, it is hard for Mary to use this evidence to estimate the value she would get from eating a durian. In all these types of case we might think that the evidence is simply too sparse or too messy to license precise reasoning.

One might conclude that in cases of messy or unclear evidence we are not licensed to form any kind of doxastic attitude.¹⁸ Taking this line would mean throwing out much of epistemology as well as decision theory. One might think that this goes too far: after all, we do have some information in situations like Elga’s. It is just not clear exactly how it adds up. In light of this observation, some people—though not Elga—have taken this type of case to call for a partial revision of standard approaches in formal epistemology and decision theory. They argue that messy cases show that, sometimes, epistemic rationality does not require us to assign *precise* credences. Rather, in some cases we are rationally required to assign “mushy” credences, which are understood either as a range of precise credences, or sometimes as some other kind of coarse-grained doxastic state. But these problems are not unique to decisions about epistemically transformative experiences; they arise across

¹⁸ Just as one might think that in cases of extreme ignorance, where we have no evidence bearing on a proposition, one should not form any doxastic attitudes at all towards this proposition.

the board.¹⁹ For most of us, the possibility that Elga's stranger has another tube of toothpaste in his bag does not involve any epistemically transformative experiences—just familiar ones that are hard to assess under the particular circumstances. Whether a particular sparse or messy body of evidence concerns epistemically transformative experiences is doing no special work here.

These points bear on parenting decisions. Alas, our epistemically uncooperative world has furnished us with evidence that is less helpful than we should like. There are two key issues in this regard. First, there is a plurality of possible parenting outcomes that might obtain: postpartum depression, the parental pride that floods social media with baby photographs, and so on. Someone can have sparse or messy evidence about whether each outcome would obtain. (This leads to the aforementioned which-ignorance of the outcomes of parenting decisions.) Second, someone can have sparse or messy evidence about the value that she would get from a particular parenting experience. For example, if Mary's only testimony about a particular experience is limited to some rather abstruse poetry, then it will be hard for her to estimate how valuable the experience would be for her. But although these issues surface with parenting decisions, no special work is done by the fact that parenting experiences are epistemically transformative. So in these respects, parenting is simply an interesting new example of a familiar epistemic problem.

In addition to epistemic problems, there are also problems concerning preferences". We will end by noting two of these. The first is that decisions such as whether or not to have children may involve incommensurable preferences. To see this, suppose for simplicity that you have good evidence that whichever choice you make, you would be happy and fulfilled. But you would be happy and fulfilled in very different ways: you are deciding between the freedom to pursue your own projects and the joy of watching your child grow and develop. As such, your preferences may be incommensurable, and there may be no way of assigning precise utilities to each experience in a way that adequately captures your attitudes. Since decision theory requires precise utilities,²⁰ decisions involving incommensurability present a challenge to

¹⁹ For criticism of mushy credences, see White [2009]. For a defense, see Schoenfield [2012]. Also see Sturgeon [2008] for further discussion of when different types of evidential situation might warrant different types of doxastic attitude. See Carr [ms] for an argument that we can accommodate intuitions supporting mushy credences without abandoning the standard Bayesian framework.

²⁰ Formally, the problem is that incommensurable preferences are likely to be negatively intransitive—we strictly prefer A+ to A, we do not weakly prefer A+ to B, and we do not weakly prefer B to A—and incomplete: it is neither the case that we are indifferent between A and B, strictly prefer A to B nor strictly prefer B to A. As we mentioned earlier, an assumption of standard decision theory is that rational agents have complete and acyclic preferences over all outcomes

standard decision theory—a challenge that is the subject of ongoing debate.²¹ But this challenge is orthogonal to the issue of epistemic transformation—we can have incommensurable preferences about things we have already experienced, and we can also have commensurable preferences about epistemically transformative experiences.

The second difficulty is that life-changing decisions often involve a shift in one's preferences or desires regarding the outcomes in question. For example, the experience of becoming a parent may change one's preferences about being a parent, or the experience of becoming a vampire (we might suppose) may involve changing one's preferences about whether to be vampire or human.²² This raises the difficult question about how practical rationality requires you to act when your current preferences diverge from the future ones. (This question is of course the close cousin of the familiar problem about whether future desires give one present reasons for action.²³) Since standard decision theory tells you only which actions are rational in light of your current preferences and credences, it is indeed silent about what rational significance your future preferences have for you. And so if future preferences are rationally significant for present choices, then this means that one would have to either concede that decision theory is not fully comprehensive as a theory of practical reason or to find a way to extend decision theory so it provides guidance about how to act in light of preference-shift.²⁴ While there are genuine philosophical problems here, these challenges are again independent of epistemically transformative experiences. Though life-changing choices may involve both epistemic transformation *and* preference shift, it seems to us that the challenges these choices pose to decision theory are just the familiar ones; epistemic transformation does not pose an additional challenge.

²¹ See Hare (2010) for a defense of prospectivism. See Bales A, Cohen D, and Handfield, T. [forthcoming] for criticism and an alternative proposal.

²² “Your effort to evaluate testimony is complicated by the fact that even people who seemed quite anti-vampire beforehand can change their minds after being bitten, suggesting that some sort of deep preference change is indeed occurring. Although your friends, as vampires, report that they are happy with their new existence, it isn't clear that their pre-vampire selves would have been happy with the change. For example, your once-vegetarian neighbor who practiced Buddhism and an esoteric variety of hot yoga now says that since being bitten (as it happens, against her will), she too loves being a vampire...Which preferences matter more? Your current, human preferences, or the preferences you'd have if you were bitten? How can you rationally choose to ignore your current preferences when making your choice? If you choose to become a vampire simply because you think that the fact of becoming a vampire will make you into a being who will be happy with the choice you've made, you are not choosing by considering your own (current) preferences.” Paul (2014, pp. 46-47).

²³ See, e.g., (Nagel 1970), (Parfit 1984), (Harman 2009), (Brink 2010).

²⁴ Discussions of preference-shift and decision theory include (Weirich 1981), (Ullmann-Margalit 2006), (Arntzenius 2008), (Briggs 2010).

7. Conclusion

Life-changing decisions, such as the decision of whether to become a parent, are indeed difficult. They pose serious challenges for decision theory. And they often involve epistemically transformative experiences, too. But we have argued that, contrary to Paul, the challenges these choices pose for decision theory do not arise *because* they involve epistemically transformative experiences. Rather, life-changing experiences present us with a tangle of well-known difficulties for decision theory: the fact that our evidence about the value of future experiences is often sparse or messy, that our preferences may be incommensurable, and that these preferences may change in the future. Thus, when it comes to life-changing decisions, there are many factors that make it hard – or perhaps even impossible – to rationally decide what to do. But the fact that these decisions involve epistemically transformative experiences is not one of them.

References

Arntzenius F. (2008). “No Regrets, or: Edith Piaf Revamps Decision Theory,” *Erkenntnis* 68, pp. 277–297.

Bales A, Cohen D, and Handfield T. (Forthcoming). Decision theory for agents with incomplete preferences. *Australasian Journal for Philosophy*.

Briggs, R. (2010). Decision-theory paradoxes as voting paradoxes. *The Philosophical Review* 119 (1), 1-30.

Brink D. (2010). Prospects for temporal neutrality. in Callender C. (ed.) *The Oxford handbook of philosophy of time*, Oxford: Oxford University Press.

Burkeman, O. (2013). This column will change your life. *The Guardian*.

Carr, J. (ms). Imprecise evidence without imprecise credence.

Chang, R. (ed.) (1997). *Incommensurability, incomparability and practical reason*. Cambridge: Harvard University Press.

- (2002). The possibility of parity. *Ethics* 112(4), 659-688.

Elga, A. (2010). Subjective probabilities should be sharp. *Philosophers' Imprint*, 10(5), 1 - 11.

Elster, J, and Roemer, J. (eds.) (1991). *Interpersonal comparisons of well-being*. Cambridge: Cambridge University Press.

Gopnik, A. (2013). Is it possible to reason about having a child? *The Wall Street Journal*.

Hare, C. (2010). Take the sugar. *Analysis* 70(2), 237 - 247.

Harman, E. (2009). 'I'll Be Glad I Did It' Reasoning and the Significance of Future Desires," *Philosophical Perspectives* 23 *Ethics*, 177-199.

Jackson, F. (1986). What Mary didn't know. *The Journal of Philosophy*, 83(5), 291-295.

Langton, R. (2007). Objective and unconditioned value. *Philosophical Review*, 116(2), 157-185.

Lewis, D. (1990). What experience teaches. In Lycan W. (ed.) *Mind and cognition: A reader*,
Oxford: Blackwell, pp. 499 - 519.

Lombrozo T. (2013a). Is having a child a rational decision? *NPR*

- (2013b). Is it rational to have a child? Can psychology tell us? *Psychology Today*.

Nagel, T. (1970). *The possibility of altruism*. Princeton NJ: Princeton University Press.

- (1974). What is it like to be a bat? *Philosophical Review*, 83(4), 435-450.

Nozick, R. (1974). *Anarchy, state and utopia*. New York: Basic Books.

Paul, L. A. (2015 a). What you can't expect when you're expecting. *Res Philosophica*, 92(2), 1-23. Accessed at <http://dx.doi.org/10.11612/resphil.2015.92.2.1>

- (2015 b). *Transformative Experience*. Oxford: Oxford University Press.

Parfit D. (1984). *Reasons and persons*. Oxford: Oxford University Press.

Rothman J. (2013). The impossible decision. *The New Yorker*.

Schoenfield, M. (2012). Chilling out on epistemic rationality. *Philosophical Studies* 158(2),197-219.

Sturgeon, S. (2008). Reasons and the grain of belief. *Noûs* 42(1), 139-165.

Ullmann-Margalit, E. (2006). Big decisions: Opting, converting, drifting. *Royal Institute of Philosophy Supplement*, 58, 157-172.

Van Fraassen, B. (1984). Belief and the will. *Journal of Philosophy*, 81(5), 235-256

Weirich, P. (1981). A bias in rationality. *Australasian Journal of Philosophy*, 59 (1), 31-37.

White, R. (2009). Evidential symmetry and mushy credence. In Szabo Gendler T. & Hawthorne J. (eds.), *Oxford studies in epistemology*. Oxford: Oxford University Press, pp. 161-186.