

TRANSFORMATIVE CHOICES

Ruth Chang

Abstract: This paper proposes a way to understand transformative choices, choices that change ‘who you are.’ First, it distinguishes two broad models of transformative choice: 1) ‘event-based’ transformative choices in which some event—perhaps an experience—downstream from a choice transforms you, and 2) ‘choice-based’ transformative choices in which the choice itself—and not something downstream from the choice—transforms you. Transformative choices are of interest primarily because they purport to pose a challenge to standard approaches to rational choice. An examination of the event-based transformative choices of L. A. Paul and Edna Ullman-Margalit, however, suggests that event-based transformative choices don’t raise any difficulties for standard approaches to rational choice. An account of choice-based transformative choices—and what it is to be transformed—is then proposed. Transformative choices so understood not only capture paradigmatic cases of transformative choice but also point the way to a different way of thinking about rational choice and agency.

An angel walks into a fractious philosophy department meeting and says to the Chair: “I’ll give you one of three gifts you choose: Wisdom, Truth, or Ten Million Dollars.” The Chair chooses Wisdom. She is transformed! But all she does is sit there, staring down at the table. One of her colleagues whispers to her, “Say something!” The Chair replies, “I should have taken the money.”¹

Some choices are transformative; they *change who we are*. But what is transformative choice? What is transformation and what gets transformed in a transformative choice? How does transformation take place? And if transformative choices are rational, how can they be rational?

¹ Adapted from Carthart and Klein 2007, 79.

In this paper, I moot a view of transformative choices that answers these questions. While there are different phenomena that go under the label, I will assume that the main interest of transformative choices is in how they push the boundaries of what might be called ‘standard approaches’ to rational choice. Do transformative choices require us to abandon these approaches? I explore what I take to be the most natural way of thinking about transformative choices and argue that thinking of them in this way poses no threat to standard approaches. I then propose an alternative view that, I argue, requires us to reject a fundamental assumption of such approaches. Its central idea is that transformation is not only something that can happen to us but is something we can *do*. Transformative choices so understood, I suggest, not only capture paradigmatic cases of transformative choice but also point the way to a different way of thinking about rational choice and agency.

1 Two Models of Transformative Choice

We start by proposing two general models of transformative choice, and in particular, two ways in which you might change ‘who you are.’ We’ll have more to say about what it might be to be ‘transformed,’ that is, to change ‘who you are’ later, but for now, I want to work with the broad, but I think intuitive, understanding of ‘transformation’ and ‘who you are.’

First, you can be transformed by an event or process. One especially salient kind of event is an *experience*. In 1999, Mike May received an operation that partially restored his sight after 43 years of being blind.² We might say that the experience of seeing transformed him—it changed him from being an uncannily talented unsighted person who had broken several downhill skiing records, worked for the CIA, and invented a GPS system for the blind, to a partially-sighted man who has difficulty identifying coke cans at the market and holding a conversation with someone while looking at him or her. May was transformed from a high-functioning unsighted person to a partially-sighted person who struggles—valiantly and, for the most part successfully—to do what sighted people do as a matter of course (Kurson 2007).

Experiences can be, roughly speaking, ‘extended,’ as in May’s case of experiencing life as a sighted person, or ‘discrete,’ as in the experience of seeing or hearing for the first time. Discrete experiences, such as giving birth, climbing Mount Kilimanjaro, or undergoing violent trauma might not only be themselves transformative, but also be the root of an extended experience. A discrete experience, such as being a victim of a violent crime, can cause you to have the extended experience of seeing strangers as threats.

An experience is one kind of event or process, but there are other kinds of events that can transform you. What transforms you needn’t be something

² See, e.g., <http://www.theguardian.com/science/2003/aug/26/genetics.g2>, last accessed May 1, 2015.

subjective in you but something objective in the world. You might be transformed by the event of taking a pill that scrambles your brain and, consequently changes who you are. There need be no appeal to a subjective experience to explain your transformation. More generally, events in the world—including in your own life—can be transformative. When Steve Jobs was fired from Apple, it was arguably the event of being fired that transformed him, not his subjective experience of being fired. Being fired was something that happened in his life, but what ‘changed who he was’ wasn’t the experience of being fired but the event of his being fired. Similarly, the process of military training can transform a cadet; it needn’t be how the cadet subjectively experiences the training that changes her; rather it may be the process of being trained to kill that does the transformative work. And when an assistant professor gets tenure, it can be the achievement of having won tenure that transforms her, not the subjective feel of knowing that she can now be fired only for moral turpitude, however wonderful that feeling might be.

Note that an event—whether experiential—can be transformative independently of whether you *choose* it. The nature and features of certain events—including whether they are transformative—may, of course, *depend* on whether you chose them. Experiencing a violent fistfight may transform you if you didn’t choose it, but might be all in a day’s work if you cage fight for a living. And the arduous discipline of military training might not transform you—or at least not in the same way—if you chose to have it rather than having had it foisted upon you. There are many other factors that will determine whether an event changes who you are—like the social conditions that can determine the meaning of an event (Barnes 2015). For our purposes, we just need to emphasize that in these cases, it is not the choice but rather an event downstream from choice that does the transforming. The first model of transformative choice, then, understands transformative choice as a choice about whether to undergo or bring about an event downstream from choice that transforms you. Call these ‘event-based’ transformative choices.

There is a second way in which you can change who you are. You can be transformed by a *choice itself*. According to a second model of transformative choice, you change who you are by the very making of a choice, not by some experience or event downstream from your choice. Call these ‘choice-based’ transformative choices. Of course, a choice is also an event, but for the purposes of this paper, I will understand ‘events’ as always downstream from a choice in order to mark a distinction between the choice itself transforming an agent and some event or process downstream from choice doing so.

Here are some examples in which the choice itself arguably transforms an agent. Suppose a Hollywood plastic surgeon chooses to give up her 1%-er lifestyle to become a volunteer doctor in a war torn region. Her very choice to give up her material comfort for one of hardship and penury

may change who she is. Or consider Gauguin, whose choice to leave his family to pursue his art in Tahiti transformed him into someone committed to his art at the cost of abandoning prior commitments to his wife and children. And when William Styron's Sophie chose to save her son, Jan, over her daughter, Eva, from certain death in the Nazi death camp, she is transformed by her choice—under duress though it was—into a tragic figure haunted by her role in the death of her daughter. Sometimes the very making of a choice itself can be transformative.

So we might distinguish two models of transformative choice: 1) *event-based transformative choice*, a choice in which an event or process downstream from a choice—perhaps experiential—transforms you, and 2) *choice-based transformative choice*, a choice in which the making of the choice itself transforms you. And there are 'mixed' transformative choices that involve some or other of these two routes to transformation.

My interest in this paper is in understanding whether transformative choices on each model raise any serious challenge to standard approaches rational choice. But I need to be clear about what I mean by 'standard approaches.' Those who worry about transformative choices have tended to assume a rather narrow target—classical expected utility theory, already abandoned by many contemporary philosophers of practical reason—in arguing that transformative choices raise problems for rational choice. By 'standard approaches' to rational choice, I mean something much broader that goes well beyond standard forms of normative expected utility theory and classical rational and social choice theory. My concern is not to show that economists and classical decision theorists have a problematic view of rational choice, though that will be an implication of my argument. I think the challenge that transformative choices pose is much broader. By 'standard approaches' to rational choice, I mean to include *any* reasons- or values- or preference-based normative approach to rational choice that makes the following two fundamental assumptions:

- (1) The rationality of a choice is determined by the value (or utility) of the alternatives or the reasons for and against them,³ and
- (2) That in virtue of which we have reasons or values (or utility) is not a choice itself.

The first assumption is clear; the rationality of a choice is given by the value of, or normative reasons for and against, the alternatives.⁴ So, for example, the rational choice might be the alternative that is at least as good

³ By 'determined' I mean to include both the subjective and objective case. That is, standard approaches hold that the values you assign to an alternative determine what it's subjectively rational to choose and that the values the alternatives in fact have determine what it's objectively rational to choose. As I am understanding 'standard approaches,' the first assumption can thus be undermined in either case, but I will for the most part be concerned with the objective case. So when I talk of 'assigning a value,' I assume that the assignment is correct.

⁴ I'll be using 'values' and 'reasons' interchangeably. Nothing I say here turns on which, if either, is explanatorily more fundamental.

as the others or most strongly supported by reasons. Or it might be what simply ‘ought to be done’ or is ‘good enough.’ So long as the view holds that the rationality of a choice is based on either the value of or the reasons (which might themselves be given by preferences constrained by certain axioms) for the alternatives, it counts as a potential target view to which transformative choices might pose a challenge.

The second assumption is often only implicitly held or can fairly be imputed; choice itself can’t be that in virtue of which you have a reason or you ought to do something. Simply by choosing to jump off a cliff, you can’t make it the case that you have a reason to do so or that it’s a valuable thing to do. To think otherwise would lead to the familiar problem of bootstrapping reasons, viz., that we can simply create reasons for ourselves to x simply by choosing x .⁵ Strictly speaking, most standard views don’t broach the question of that in virtue of which we have reasons. But if they didn’t implicitly assume that choice couldn’t be that in virtue of which we have reasons, the substance of their accounts would be very different. They would have to allow that the mere fact that you choose something could create an additional reason to choose it and thereby be self-justifying.⁶ This consequence is not consistent with views I include among ‘standard approaches.’ And some such views have explicitly denied that choice itself can be that in virtue of which we have reasons (Scanlon 2004). According to standard views, our reasons to x are given by facts that count in favor of x —such as the fact that it’s delicious or that I promised or that I want or prefer it—and those facts are reasons in virtue of facts *other than* the fact that we have chosen a certain way. ‘Standard approaches,’ then, include nearly every view about rational choice that has been a going concern in the last few centuries.⁷

⁵ An especially vivid illustration of the bootstrapping problem is given by Jerry Cohen’s Mafioso objection against Korsgaard’s Kant according to which willing can be a source of reasons. If choosing can be a source of reasons, then the Mafioso can bootstrap his way into having all-things-considered reasons to shoot the kneecaps off his enemy simply by choosing to do so. See Korsgaard 1996.

⁶ Which is not to say, of course, that norms of structural rationality, such as ‘follow through on your choice unless you have a reason not to,’ might not be involved. The interest here is on the ‘rationality of reasons,’ not of the norms governing consistency and coherence among our mental states.

⁷ Revealed preference theory isn’t normative in the sense of interest, and so I exclude it from ‘standard approaches.’ Other theories excluded are a particular kind of neo-Kantianism according to which willing is, strictly speaking, the ground of one’s reasons, such as that I believe championed by Christine Korsgaard (1996). Other neo-Kantians, such as Elizabeth Anderson (1993) and Barbara Herman (1996), explain the ground of practical reasons in terms of some fundamental value, such as the value of humanity, while still other neo-Kantians, such as Thomas Hill (2001), as I read him, take a roughly Humean approach to the ground of practical reasons. I have argued elsewhere that views that may appear to hold that the rationality of a choice is determined by something other than the value of the alternatives or the reasons for them—such as views according to which the rational choice is just ‘the thing to do’ or what meets some test or standard, are in fact views that hold that the rationality of

As we will see, event-based transformative choices purport to challenge the first assumption and choice-based transformative choices the second. But, I'll suggest, event-based transformative choices *don't* raise any difficulties for the first assumption and so, to that extent, can be handled by standard approaches to rational choice. However, as I'll argue, choice-based transformative choices *do* require rejection of the second assumption and thus raise a genuine challenge to standard approaches. So one upshot of the arguments here is that only choice-based, and not event-based, transformative choices, raise difficulties for standard approaches to rational choice.

We start with two accounts of event-based transformative choices, one offered by L. A. Paul, who focuses on a particular kind of experiential event, 'epistemically transformative experiences,' and the other by Edna Ullman-Margalit, who considers transformative events more broadly. These are the only two analytic accounts of transformative choices of which I am aware. Both have as their target normative rational decision theory, and they purport to raise a problem for such theories because they claim that the rationality of a transformative choice cannot be based on the value of the alternatives. If successful, these accounts would also pose a problem for standard approaches, as they are understood, broadly, here. But I have doubts about whether these views, while interesting in their own right, succeed in raising any difficulties for the idea that the rationality of choice is based on the value of the alternatives. They don't seem to raise a genuine challenge to standard ways of thinking about rational choice.

We then turn to the choice-based model of transformative choices. We start by proposing a way to understand transformation—who and what is transformed and how transformation takes place. We then describe how choices might themselves be that in virtue of which something is a reason, that is, 'grounds' for something's being a reason. When we choose in a *thick* sense, that is, by *committing* to an alternative, we create reasons for ourselves to choose it—our commitment is that in virtue of which we have a reason to do something. So by choosing, we can create new reasons for ourselves, thereby transforming 'who we are.' This choice-based view of transformation does not run afoul of the bootstrapping problem that led standard approaches to assume that choice cannot ground reasons because it is part of a more general metanormative view about practical normativity, what I have elsewhere called 'hybrid voluntarism,' according to which our normative power to create reasons through our commitments is suitably constrained. Thus, hybrid voluntarism, while I believe an independently attractive view about the source of normativity, also underwrites an attractive and plausible account of transformative choices. Crucial to choice-based transformative choices is the idea that transformation is something we *do*,

choice is determined by the *comparative* value of or strength of reasons for and against the alternatives. See [Chang Forthcoming](#).

not something that happens to us. I suggest that this account not only captures paradigmatic cases of transformative choice but points to a richer understanding of what it is to be a rational agent.

It's worth saying at the outset that since philosophical investigation of transformative choices is relatively new territory, especially for analytic philosophers, I'll be relying on large, abstruse—but ordinary—notions on which we may have little more than an intuitive fix. If you don't share the notions to which I appeal, remember that my aim is to make a case for *one* way of thinking about transformative choice, and in particular, a way that poses a genuine challenge for standard ways of thinking about rational choice.

2 Event-Based Transformative Choices

In an event-based transformative choice, an event downstream from choice transforms you. One particular kind of an event—an experience—might be thought to be an especially good candidate event that can transform an agent.

2.1 Epistemically Transformative Experiences

L. A. Paul has recently proposed that a very particular kind of experience—an “epistemically transformative” experience—can be personally transformative, that is, “change who you are, in the sense of radically changing your point of view” (2014, 10–11). As Paul goes on to say, changing your point of view is a matter of changing your “personal or subjective preferences” (16). And when you radically change your point of view, you “change your post-experience preferences, or change how your post-experience self values outcomes” (48).

An epistemically transformative experience is an experience you can't know what it's like to have without actually having the experience (10). ‘What it's like’ to have an experience goes beyond its raw phenomenal feel and includes attitudes and emotions you might have in response to that feel (27). As Paul explains, “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*. Her knowledge of what something is like, and thus her point of view, changes” (16).

For example, Paul urges that the experience of having a child—the experience of “gestating, producing, and becoming attached to that child”—is epistemically transformative: you can't know antecedent to the experience of having a child, what it's like for you to have a child (77–78; see also Paul 2015). And since the knowledge you gain when you experience having a child “radically chang[es] your point of view,” the experience is also personally transformative.

Since you can't know what an epistemically transformative experience is like before having the experience, you can't assign value to what it's subjectively like antecedent to the experience. This, in turn, Paul thinks, raises a problem for standard forms of normative expected utility theory, which, she notes, presupposes that the rationality of a choice is determined by the values of the alternatives. If your choice is whether to have an experience you can't assign a value to having, then you can't rationally choose it on the basis of its value: normative expected utility theory seems to break down.⁸ Paul also makes the further point that since you are personally transformed, the preferences you have before the experience are different from those you have after the experience and so, she urges, there is the problem of determining which set of preferences should be the ones on the basis of which to evaluate the alternatives. Since this latter aspect of Paul's view is essentially the same as the central point of Edna Ullman-Margalit's account of transformative choices, which we discuss below, we here focus on what is most distinctive about Paul's account—viz., her claim that epistemic transformation entails that you can't know the value of the experience and so can't assign it a value, which you need to do in order rationally to choose whether to undergo it.

Interestingly, although Paul thinks that epistemically transformative experiences lead to a breakdown in rational choice, she doesn't conclude that we should reject normative expected utility theory. Instead, she suggests that we reconceive epistemically transformative choices in a way that no longer poses a challenge to the standard approaches. We should think of the choice about whether to have a child, for example, not in terms of what it's like to have a child but instead as a choice about whether to gain a certain kind of knowledge—viz., the knowledge of what it's like to have a child.⁹ That reframing of the choice, Paul suggests, saves standard approaches from the challenge posed by epistemic experiences because we can assign value to knowing what an experience is like.

I doubt, however, whether epistemically transformative experiences raise any difficulties for any plausible forms of normative expected utility theory or, indeed, for 'standard approaches' to rational choice more broadly understood here. First, it seems that genuinely epistemically transformative experiences as Paul strictly understands them are very rare; we can know what most experiences are like antecedent to having the experience, and

⁸ As Paul writes, "To apply a normative decision-theoretic model for ignorance to a decision about whether to perform an act, you need to know the values of the relevant outcomes, including their relative strengths, and you must be able to compare the values of the outcomes in order to determine the overall structure of the value space. But in the case of a decision involving a[n epistemically] transformative experience, you cannot know what it is like to have that kind of experience until you've had it" (2014, 32). Note too that Paul's worry applies to the subjective reading of the first assumption.

⁹ Paul: "[In transformative choices] we choose between the alternatives of discovering what it is like to have the new preferences and experiences involved, or keeping the status quo" (2014, 122). Thus what it's subjectively like drops away as relevant to the choice.

we can know enough about them in order to assign them at least some rough value. Second, even if we can't know what an experience is like before having it, it doesn't follow that we can't assign it a value. This is because the value relevant to rational choice isn't simply the value of what it's subjectively like. The objective value of experience matters too and typically matters in a way that allows us to assign a rough value to the experience. And third, quite generally, the focus on *experience* is misplaced; transformative choices aren't typically about what *experiences* to have but are choices about whether to undergo certain transformative *events* that go well beyond how things subjectively feel to an agent. Transformation isn't typically about how we subjectively experience things but about how things *in the world* change us. While an interesting phenomenon in its own right, epistemic transformation is not, I think, the right key to understanding transformative choices.

2.1.1 *Epistemic Transformation Is Rare and Atypical of Transformative Choices*

What sort of experiences are most plausibly ones about which we could have no antecedent knowledge of what they're like? I suggest that, as a first cut, we work with the idea that epistemically transformative experiences as the *de novo* exercise of a 'basic' capacity. Trying to give a proper account of what makes something a basic capacity would take us too far afield, but we might give an intuitive gloss of them as follows: a basic capacity is a *sui generis* capacity that belongs to a set of capacities from which all others can be derived for a type of creature; basic capacities are ones that are not the exercise of other capacities but are atomic capacities for that type of creature. And since our interest is in rational choice for humans, we focus our attention on the basic capacities of human rational agents. Seeing, hearing, tasting, touching, smelling—the five traditional senses—are plausibly basic physical capacities. But so too is the normative capacity to recognize and respond to reasons. If you've never exercised those capacities before, it's plausible to think that you can't know in advance of experiencing their exercise what it's subjectively like to experience their exercise.

What about seeing red, tasting Vegemite, or having a child for the first time? Are these cases, which Paul uses to illustrate epistemically transformative experiences, really experiences about which we could have no knowledge antecedent to having the experience?

Paul uses as her touchstone case that of Frank Jackson's Mary, who, after living in a black and white room all her life, emerges and sees red for the first time. Jackson points out that antecedent to the experience of seeing red, Mary can't know what it's like to see red. The point of Jackson's (1982) example is to show that there are some phenomenal properties that require experience in order to know them (and that physicalism was therefore in

doubt).¹⁰ ‘What it’s like’ for Jackson, however, is a simply a matter of phenomenological feel. As we’ve seen, ‘what it’s like’ for Paul goes beyond mere phenomenal feel and includes attitudes and emotional responses to that feel (2014, 12, 27). Paul is right to insist upon this broader notion of ‘what it’s like’ since phenomenal feel mostly isn’t relevant or all that’s relevant in transformative choices. When we are talking about ‘what it’s like’ to have a child, for instance, we aren’t concerned simply with the phenomenological feel of that experience. So when Paul argues that just as Mary can’t know ‘what it’s like’ to see red, the child-free you could not know ‘what it’s like’ to have a child, we must be careful that there is no equivocation. So we might ask, Can Mary know what it’s like to see red in the broader, beyond-mere-phenomenological-feel sense that is relevant to transformative choice, antecedent to the experience of seeing red? As we’ll see, there is good reason to doubt that Paul’s touchstone case is a case of epistemic transformation, and a fortiori, that ‘higher-level,’ more complex experiences of tasting something new and becoming a parent are cases of epistemically transformative experiences. As it turns out, genuinely epistemically transformative experiences are rather hard to come by.

But first, we need to make a clarification. Paul says that an epistemically transformative experience as an experience that you *could not know what it is like* to have without having the experience.¹¹ There is a stronger and a weaker interpretation of this claim. On the stronger interpretation, you can’t have *any knowledge whatsoever* what the experience is like antecedent to having the experience. It would then follow straightforwardly that you can’t assign a value to it on that basis. On the weaker interpretation, although you can have *some* knowledge of what it’s like, that knowledge is not sufficient for you to assign a value to it based on what it’s like. Paul’s text seems implicitly to endorse the stronger reading since she raises objections to claims about how you could have *some* knowledge of what it’s like as a way of bolstering her claim that you can’t know what it’s like and since she doesn’t give an account of what sort of knowledge of what it’s like would be insufficient to assign a relevant value to the experience. But it’s worth considering both interpretations in turn.

Could Mary, who is born and raised in a black and white room, have *no knowledge whatsoever* about what seeing red is like? I think the answer is ‘no’; although Mary may not be able to know the phenomenal feel of seeing red, she might nevertheless have *some* knowledge of what seeing red is like beyond its phenomenal feel. I think this is true for three reasons.

First, seeing red isn’t for Mary simply the de novo exercise of basic capacities. Seeing red is, after all, like seeing black and white in some

¹⁰ See also Nagel 1974 and Lewis 2004.

¹¹ Paul: “[I]n the case of a decision involving a[n epistemically] transformative experience, you cannot know what it is like to have that kind of experience until you’ve had it” (2014, 32). The modality of ‘cannot’ here is significant in evaluating the scope of epistemic transformation but I am unclear as to what she has in mind.

respects. While seeing is a basic capacity, seeing red is not—it involves the exercise of the basic capacity of seeing, which by hypothesis, Mary has exercised before. Since it involves the familiar exercise of a capacity, Mary can have *some* knowledge of what it's like. Thus, one reason to think that seeing red isn't epistemically transformative is that it involves the exercise of capacities that have been exercised before.

Seeing red might also be said to belong to a *type* of experience, where 'types' are individuated by 'what it's like' for humans. All humans have a certain range of subjective responses to the exercise of their capacities. But since there is a wide variation of human subjective responses to experiences, an experience will belong to many different 'types'—different 'what it's like's.' So for example for some humans, seeing red will fall under the 'emotionally neutral' type of experience, while for others it may fall under the 'thrilling' type of experience. The point here is that, given human capacities and variations, a matrix of 'types' of experiences could be constructed so that every possible human experience could be classified as belonging to a range of 'types'—thrilling, emotionally neutral, boring, etc. Since there's no a priori reason to think that such an individuation of experiences wasn't possible, Mary could look up the experience of seeing red in the matrix and see the range of types of experience in which seeing red falls for humans. She would see, for instance, that seeing red falls under many of the same types of experiences that other experiences she's had fall under. If seeing red always falls under a type of experience that includes other experiences she's had, then she would thereby have some knowledge—disjunctive though it may be—of what seeing red would be like for her. It would be like one of those other experiences that fall under the types that seeing red falls under for humans generally. She doesn't need to know *which* type of experience seeing red would fall under *for her*; it's enough that she knows that whatever seeing red will be like for her, it will be like one of the experiences she has already had that fall under the types of experience that seeing red falls under for humans generally. So a second reason to think Mary could have some knowledge of what it's like to see red is that she could know what it's disjunctively like for her to see red.

Finally, Mary can get testimonial evidence about what it's like to see red from those similarly situated—or in the ideal case, from those who share the same physiological and psychological properties that subvene her own experience of seeing red and are otherwise similarly situated. Of course, Mary would need also to gauge the reliability of the testimony and be aware that the testimony given might reflect an ex-poste shift in the way one comes to view an experience after having it. But after confirming reliability, such testimony gives her at the very least knowledge of what an experience is thought to be like after having had it—even if that involves a change in her preferences—by those that share her subvening properties. That is some knowledge about what it's like—it's the kind of experience which, ex ante, those similar to her—or in the best case scenario, her

Doppelganger—experience in such-and-such way, and ex poste, report in such-and-such way. If the testimony is from those merely similar to her but not exactly like her in the relevant respects, it is evidence of what it might be like for her to the extent that she is like those whose testimony she has.¹²

This is not to say that if Mary is to make a rational choice about whether to see red, she must blindly follow the testimony of others or that her choice about whether to see red is no longer first-personal or ‘authentic,’ a worry Paul raises about the appeal to testimonial evidence (2015, 19 and footnote 33; 2014, 105–107). Taking account of evidence about what the experience is like for those similar to ourselves does not entail that when we make a decision about whether to undergo an experience, we don’t act autonomously or authentically or from the first-personal perspective. None of those features of agency require that we act simply from our own subjective preferences, uninformed by external facts, including the testimony of others.¹³ Mary’s Doppelganger might reliably inform Mary that seeing red is just fine. But Mary might nevertheless prefer not to take the risk, just in case the testimony isn’t one-hundred-percent reliable or is ex-poste corrupted. Thus appeal to testimonial evidence doesn’t undermine Mary’s agency. She can have some knowledge of what seeing red is like through testimony.

It’s worth noting that the last two ways in which you can get knowledge about what an experience is like also hold for an agent’s de novo exercise of basic capacities that have already been exercised by other humans. A congenitally blind person can have some knowledge of what seeing is like by knowing, disjunctively, what seeing is like for the human species, and by reliable testimony from those who share her subvening properties. She won’t be able to know its phenomenal feel, but that’s not to say that she can know nothing about what it’s like.¹⁴ On the strong interpretation of epistemic transformation, then, the only genuinely epistemically transformative experiences are those involving *only* the de novo exercise of basic

¹² The modality of ‘can’t’ in Paul’s claim that you can’t have knowledge is important here, but since Paul does not elaborate, I leave open the sort of testimonial and matrix evidence to which any agent might have access. If, for example, everyone could always have reliable testimonial evidence from her Doppelganger, that would make short work of epistemic transformation. Insofar as Paul has a very restricted sense of ‘can’t’ in mind, this reduces the scope—and interest—of her claims accordingly.

¹³ Paul seems to assume a view of authentic agency whereby an agent simply models possible futures without external, non-self-generated data, and then consults her subjective preferences about those futures. I believe that there are more plausible views of authentic agency that even standard normative expected utility theorists can help themselves to. See Paul 2014, 105–107, 112, 130. We shouldn’t elide authentic, autonomous, and first-personal choice with choice based solely on our preferences about which subjective experience to have.

¹⁴ Later, I’ll be giving another reason to doubt whether the choice of whether to become sighted involves an epistemically transformative experience: what transforms you in such a choice isn’t the experience of being sighted but the goods and bads of the objective fact of being sighted.

human capacities for which no matrix can be constructed and no reliable testimony can be given. It's difficult to imagine such experiences. Maybe evolved new basic capacities unlike any capacity ever exercised by humans before, such as teletransporting oneself by thought alone, would count. But the paradigmatic transformative choices that are of interest to us don't involve such fanciful experiences.

What goes for Mary goes too for the person tasting Vegemite for the first time and, most importantly, for the person having experiences that are typically thought to be personally transformative, such as the experience of having a child. The experience of “gestating, producing and becoming attached to a child” involves the exercise of many capacities you've exercised before. So on the basis of your experience in exercising those capacities, you could have some knowledge of what the experience of having a child will be like for you. Moreover, given our matrix, you can know that the experience of having a child falls under a certain range of types of experiences for humans, and you've very likely had experiences that fall under those types before. Maybe the experience of having a child falls under types that include the experience of being in a family, passing a kidney stone, having a pet, and so on. Since you've had experiences that fall under the same types before, you will know *something* about what it's like to have a child. And, finally, you can get reliable testimonial evidence from parents who share the properties—including the moral and social ones—that subvene what it's like for you to have a child (Harman 2015). That testimonial evidence will give you *some* knowledge of what it's like for you to have a child, especially since gestating, producing, and becoming attached to your child is a multifaceted experience that involves many different ‘sub’ experiences about which you can know something in one of our three delineated ways. None of this is to deny that the experience of having a child will involve many new experiences or that you can know *exactly* what it would be like for you before actually undergoing the experience. But our point is only that it will involve experiences that either you have had before or are like experiences you've had before, or be amenable to reliable testimonial report. And so you can have *some* knowledge of what it is like. If you can have some knowledge of what it is like, we can't simply assume, with Paul, that you *can't* assign some, perhaps rough, value to the experience.

The strong interpretation of Paul's claim, then, doesn't hold of the kind of experiences that seem most relevant to transformative choices. You not only can know *something* of what it's like to see red for the first time, but you can also know *something* of what it's like to try Vegemite or to have your first child. We haven't shown that you can, therefore, assign value to the experience—we'll make a suggestion about that in a moment—but have argued only that since such experiences aren't epistemically transformative in the strong sense, it doesn't follow that we can't assign a value to them.

What about the weak interpretation of epistemic transformation? When Paul claims that we can't know antecedent to an experience what it's like,

she may mean that to be consistent with the claim that we know something of what it's like, or know what it's like to some extent (though, as I've noted, this interpretation does not sit easily with the text). What she may mean by 'knowing what it's like,' in particular, is 'knowing what it's like sufficient to assign it a value on that basis.' That might be consistent with knowing *something* of what it's like.

But what could 'knowing what it's like sufficiently to assign it a value on that basis' amount to? One important question is: If an experience involves the de novo exercise of a basic capacity, is that sufficient to *block* assigning it a value on the basis of what it's like? To answer this question, we need a theory of the ways in which various bits of knowledge of what it's like contribute to the value of what it's like. We want to know, in particular, whether there are 'organic unities' that form between different pieces of knowledge of what an experience is like so that if we lack a bit of knowledge, we lack the knowledge needed to assign a value on the basis of what it's like. Paul doesn't offer a theory and we don't have space to try out a theory here, but we can consider some cases and draw a tentative conclusion.

There are some cases in which the de novo exercise of a basic capacity is clearly such an insignificant contributory factor to what an experience is like that you can nevertheless assign a rough value to what it's like. Suppose for example, that Mike May, while blind, is hit by a car. Pre-sight-restoring-operation, May knows enough about what it would be like to be hit by a car while sighted to assign it a rough value, even though the experience of being hit by a car while sighted would involve the de novo exercise of the basic capacity of seeing. What it's like to be hit by a car is primarily about the hitting, not about the seeing.

There are other cases, however, in which the experience so centrally involves the de novo exercise of a basic capacity that it seems that you can't know what it's like without knowing what it's like to exercise that basic capacity. The experience of listening to Beethoven's Fifth Symphony for a deaf person might seem to be like this. Although that experience involves the exercise of capacities the person has, suppose, exercised before, for instance, the capacity to take sensory input to form an artistic interpretation, what the experience of listening to Beethoven's Fifth is like is so centrally about hearing, that without familiarity with what it's like to hear, it seems that you can't know what it's like to listen to Beethoven's Fifth.

But note that it's yet a further question whether *other* knowledge you could have of what listening to Beethoven's Fifth is like is sufficient for you to assign a rough value to the experience nonetheless. This is, as I've said, a matter of substantive argument, but for my own money, I would bet that knowing the disjunctive range of what it could be like for her to hear Beethoven's Fifth à la our matrix, and having testimonial evidence from similarly situated friends, is sufficient knowledge to assign a rough value to the experience. My own suspicion is that any plausible theory

of what knowledge is sufficient to assign value on the basis of what it's like will, at best, show that *only* experiences involving *solely* the de novo exercise of a basic capacity that hasn't been exercised much by humans before, will preclude the assignment of a value to an experience based on what it's like. Experiences, such as listening to Beethoven's Fifth by a deaf person, which involve not only the de novo exercise of a basic capacity but the familiar exercise of other capacities as well, are experiences about which we can have both 'matrix' and 'testimonial' knowledge, knowledge which is, I believe, sufficient to allow us to assign a rough value to what it's like to have the experience. You can know, for instance, that what it's like to listen to Beethoven's Fifth will be better than what it's like to be skinned alive. The same goes, I believe, for the experience of seeing, hearing, tasting, and so on, for the first time. If this is right, we need to revise our first thoughts about epistemically transformative experiences. They don't include the de novo exercise of any basic capacity, but only those that humans in general have not exercised before. This reduces the interest of epistemically transformative experiences significantly.

Regardless of whether these substantive musings are correct, we can reasonably doubt that paradigmatic transformative choices, such as whether to have a child, involve epistemically transformative experiences since, as we've seen, either, on the strong interpretation, such experiences are *recherché* and don't figure in paradigmatic transformative choices, or, on the weak interpretation, it doesn't follow from the fact that we lack some knowledge of what an experience is like that we cannot assign a value to what it's like. You can know that what it's like to have a child isn't as bad as being boiled in hot oil or slowly dismembered without anesthesia.¹⁵

None of this is to say that there couldn't be a transformative choice that involves having an epistemically transformative experience. Choosing whether solely to exercise a new basic capacity that has never been exercised by humans before plausibly involves such an experience.¹⁶ Transformative choices that are both epistemically transformative and personally transformative in Paul's sense, then, are I think best restricted to just these cases. Paul's case of choosing whether to become a vampire, though purely fictional, is plausibly such a case (though it has the added complication that it might involve abandoning our status as humans, though if Stephanie Meyer of *Twilight* fame is right, vampires and humans can interbreed). Becoming a vampire would presumably involve the de novo exercise of basic capacities never before exercised by humans, and so it's plausible to think that we can't know, as mere humans, what it's like to be a non-human

¹⁵ Another way to put the point is that we can do a bit of cognitive modeling to determine the rough value of an experience. When you simulate what it's like to have a child, you may be unsure what it will be like. But you can be sure that what it will be like is better than what it would be like to be slowly dismembered without anesthesia.

¹⁶ It seems likely that neither matrix nor testimonial knowledge will be of much help in knowing what the exercise of a new-to-humankind basic capacity is like.

vampire. But they don't include choices about whether to become sighted, to taste Vegemite for the first time, or to have a child. Such experiences always seem to involve the exercise of familiar capacities, be representable in a matrix of human responses to such experiences, and be amenable to testimonial evidence about what they're like.

It is at best doubtful that we have many epistemically transformative experiences. And paradigmatic cases of transformative choice—such as whether to have a child, change careers, give most of your wealth away to charity—don't seem to involve epistemic transformation. You can have *some* knowledge of what an experience is like, even if that experience involves the *de novo* exercise of a basic capacity, because such experiences also involve the exercise of capacities you've exercised before, you can know disjunctively what the experience is like, and you can gain suitable testimonial evidence of what it is like. We should be cautious about overgeneralizing from fictional, *recherché* cases, like choosing whether to become a vampire, to the kind of experiences typically involved in a transformative choice.

2.1.2 *Epistemically Transformative Experiences Can Be Evaluated*

If the arguments of the last section are correct, the only epistemically transformative experiences there are involve the *de novo* exercise of a basic capacity that has not been exercised by humans before—not the experiences involved in paradigmatic transformative choices. But now let us set those arguments aside and assume, for the sake of argument, that transformative experiences are typically epistemically transformative as Paul suggests. Does it follow, as Paul argues, that we can't assign a subjective value to such experiences?

Suppose, *arguendo*, that you can't know in advance of suffering a violent trauma what it would be like to suffer it. But you can nevertheless know that what it will be like will be bad. So you can assign a subjective value to the experience even if, *pace* our arguments of the last section, you can't know in advance what the experience will be like. Or suppose that you can't know in advance of rollicking with the angels in heaven what being in heaven is like, but you can know that what it will be like will be good—or at least better than burning in the fires of hell and damnation. So why does Paul think that epistemically transformative experiences can't be assigned a subjective value?

Paul herself acknowledges that “you can know that being eaten by a shark will be horrible” (2014, 27). But she says that such experiences are ones about which “there is no need to deliberate by cognitively modeling in order to assess the subjective value of the relevant outcomes” (27) and so she “will be setting aside decisions like [these]” (27). As she puts it, she wants to focus on epistemically transformative experiences “you're not sure how you'd respond to” (28).

What Paul appears to be claiming is that an epistemically transformative experience can't be assigned a subjective value unless it can be. If it can be, she's not interested in it. And whether an epistemically transformative experience can be assigned a subjective value depends wholly on whether "you're not sure how you'd respond to" it. The idea here seems to be that some epistemically transformative experiences can be assigned a subjective value independently of having to run a simulation of what they're like, that is independent of deliberation by cognitive modeling. Even though you can't know, by hypothesis, what being eaten by a shark is like, you can assign it a subjective value because its subjective value doesn't turn on running a simulation of what it's like. If, on the other hand, you're not sure what an experience is like, you do a bit of cognitive modeling, and if you're still stuck, if "you're not sure how you'd respond to" it, the experience counts as epistemically transformative in the sense Paul is interested in, the kind to which you can't assign a subjective value on the basis of what it's like. So from not being sure how you'd respond to an experience, it seems to follow, as Paul suggests, that you can't assign a subjective value to it on that basis. But not knowing how you'd respond to an experience is one thing, not being able to assign it a subjective value another. What's needed, I think, is an account of why some epistemically transformative experiences can evidently be assigned a subjective value while others putatively can't. That epistemically transformative choices split in this way suggests that being epistemically transformative isn't the feature that is really supposed to be causing problems for expected utility theory in the first place.¹⁷

It might also be wondered why Paul thinks that in order to be able to assign a subjective value to an epistemically transformative experience, we have to *know* what it's like. Standard approaches to rational choice generally recognize that we rarely *know* what will happen in the future. So the value we assign to future outcomes is based in part on the probability that that outcome will come to pass. Why can't Mike May simply have some probabilistic expectation of what seeing is like and assign a subjective value to the experience of seeing on that basis? Seeing might be like many different things, and May could have a probability distribution over the many different ways seeing might be like. He can then assign a value to each way seeing might be like, however varied these values might be. So, a toy illustration: the experience of seeing for the first time might have a 50% chance of being emotionally thrilling to him, a 35% chance of being confusing and depressing to him, a 10% chance of being scary to him, and so on. Having been thrilled, confused and depressed, and scared before, he can assign a rough value respectively to having a thrilling, confusing

¹⁷ There is a large discussion that could be undertaken here about what feature of experience can play the role Paul envisions if, as we've suggested, being epistemically transformative isn't it. But this would take us too far astray. In any case, as we argue below, any event downstream from choice, whatever its features, won't raise a problem for standard approaches to rational choice as they are understood here.

and depressing, and scary experience by coming to see even if the way it is thrilling, etc., might not be exactly the same way some of his past experiences have been thrilling, etc.

May's uncertainty may not only be nonnormative—about what the experience is like—but normative—about what utility to assign to an experience *like that*. So even assuming that he *knows* what seeing will be like, he may be uncertain as to what utility it has. In this case, he can do an expected value calculation on the utilities, though there are more complex and plausible approaches to handling normative uncertainty.¹⁸ Of course, he can't know exactly what the value of the experience will be—maybe he can only assign a rough value—and nor can he know the precise probabilities of the experience being one way rather than another—and that's why he can only assign a probability distribution among the possible outcomes. But having collected all the rough probabilities and rough utilities, he can then tot up the expected utility of seeing by multiplying the rough probability of seeing being like this rather than that and the rough value assigned to its being like this. All of this should sound familiar, because it is essentially a roughed-up version of standard expected utility theory under normative and nonnormative uncertainty.¹⁹ If we can assign probabilities to what an experience will be like for us, rather than *knowing* what an experience will be like for us, and we can assign rough utilities to each way the experience might be—however varied those utilities might be—we are squarely within the framework of expected utility theory. It doesn't matter how varied the possible utilities might be because we can still maximize the expected utility whatever the utilities—however rough—may be. It seems clear that we can assign such probabilities and utilities—*unless* of course the experience is wholly unlike anything we—or any other human—has ever experienced before. That's how we return full circle to the idea that genuinely epistemically transformative experiences are those involving only the *de novo* exercise of a capacity never before exercised by humans.

But let's grant Paul's assumption that in order to assign a subjective value or utility to an epistemically transformative experience, you have to *know* what it's like—probabilistic information won't do.²⁰ And let's also grant

¹⁸ For views about how expected utility theory can deal with normative uncertainty, see [Sepielli 2009](#) and [Ross 2008](#).

¹⁹ Things are more complicated than this sentence suggests, but we don't have to attend to those complications for our purposes. The key point is that it is unclear whether assigning subjective value requires *knowing* both what the experience is like and what its utility is if it's like that. Something short of knowing could well suffice for assigning subjective value.

²⁰ Some of Paul's assumptions seem to turn on the very particular form of normative expected utility theory that is her target. The target seems to be a view that requires for rational choice 1) knowledge of what an experience is like through cognitive modeling of what it's like, where 2) experience is all that matters for rational choice, 3) the subjective value of the experience—what the experience is like—crucially matters for the rationality of choice such that if you don't know what it's subjectively like, you can't assign it a value relevant to choice,

that a rough range of subjective value assignments won't suffice as an assignment of subjective value. Let's also grant, for the sake of argument, that there are epistemically transformative experiences beyond merely fictional cases involving the *de novo* exercise of a new-to-human-kind basic capacity, such as the vampyric capacity to conduct one's life without sleep. Because our interest is in transformation in the real world, let's also just grant, for the sake of argument, that *any* *de novo* exercise of a basic capacity is one you can't know what it's like antecedent to having the experience. So seeing for the first time, hearing for the first time, and so on, will be, by hypothesis, epistemically transformative (even though, recall, we've seen good reasons to doubt this). Even granting all this, pace our reasons to think otherwise, we can, as I'll now suggest, nevertheless assign a relevant value to an epistemically transformative experience. So even being as concessive as I think we can be, it seems doubtful whether epistemically transformative experiences raise any problems for the assumption of standard approaches that the rationality of a choice is based on the value of the options. This is because the value of an option that is relevant to rational choice goes well beyond the subjective value of what an experience is like.

Central to Paul's argument about transformative choices is the idea that such a choice "essentially involves your subjective values" (2014, 18) and that in such choices there is no "external reason that trumps or dominates your choice, making subjective deliberation irrelevant or unnecessary" (19).²¹ But she seems to think that so long as the subjective value of an experience is necessarily a part of its value, if you can't assign a subjective value, you can't assign a value to the experience.²² But this doesn't follow.

Return to Mike May who must choose whether to have an operation that will give him the experience of being sighted. Let's suppose, again for the sake of argument, that the experience of being sighted is epistemically transformative for him—he couldn't know what it is like to be sighted

4) for your choice to be 'first-personal,' it must be based primarily on the subjective value of what an experience is like, and 5) the value of an alternative can never be rough—that is, it can only be represented by a standard utility function and not by anything more sophisticated, such as a vector or probability distribution. Moreover, for there to be any reasonable scope for epistemically transformative experiences, the "can't" in the claim that they are experiences the agent 'can't' know antecedently what they're like must be read broadly to include merely contingent factors that may prevent a particular agent from knowing what an experience is like. This further narrows the import of her conclusion. See Paul 2014, *passim*. As we've noted, our understanding of 'standard approaches' is much more wide-ranging.

²¹ It's unclear to me whether by 'external reason' Paul means something like a reason of morality—for example a moral prohibition—that overrides the reasons that would otherwise determine the rational choice or whether she means to include by 'external reason' the objective values of having the experience. I take the conservative interpretation and so argue that the objective values of having an experience are relevant to transformative choices.

²² Recall that we are assuming, for the sake of argument, though we have queried it, that in order assign a subjective value, you have to *know* what the experience is like. So if you don't know what an experience is like, it follows that you can't assign it a subjective value.

antecedent to the experience of being sighted. Let us further suppose that because he can't know what it's like to be sighted, he can't assign a subjective value to the experience of being sighted. Must it follow that he can't assign a value to the experience of being sighted overall?

Note that it would be irrational to think that what it is subjectively like for May to be sighted is *all* that matters in the choice of whether to undergo the operation. The choice about whether to see is not a choice about which subjectively best experiential feel to have. At the very least the objective values of the experience matter too. For example, by having the experience of being sighted, May can gain certain objective goods that improve his well-being, such as, say, greater intimacy with his wife and children, less dependence on others in the execution of quotidian tasks, deeper and more varied social connections with strangers, and so on. Even if we assume, *arguendo*, that May can't assign a subjective value to seeing, he can, we can suppose, know and evaluate the objective goods he will gain by having that experience. This evaluation of the objective value of the experience can be sufficient for assigning a rough overall value to the experience: given the objective value of the experience, he knows that the experience of seeing has greater value overall—taking into account both subjective and objective value—than the experience of having a hot poker in his eye. This is not because the objective values of having the experience ‘trump’ what it's like, in the way that rights might trump or be lexically prior to utility, or because they ‘dominate’ the subjective value of what it's like in the sense that one item might dominate or be pareto superior to another if it is at least as good in all respects and better in at least one. In May's case, the subjective and objective values of the experience weigh against one another in the ordinary way, but because the objective values are such important and significant contributors to the overall value of the experience, a rough overall value can be assigned to the experience even without knowing its subjective value. The subjective value of the experience might then affect the value of the experience within some rough range.²³ So if we can know the objective value of having an experience, even if we don't know its subjective value, we might nevertheless be able to assign it a rough overall value.

What about the case of choosing whether to have a child? Gestating, producing, and becoming attached to a child has a certain subjective feel. But it would be the height of irrationality to think that the choice of whether to have a child is simply a matter of getting the subjectively best experience. The objective values of the experience also matter, and they matter significantly. You can know, for instance, that the experience of

²³ The same goes, I believe, for experiencing life as a non-human vampire or bat or squirrel. You can have knowledge of the objective value of that experience. So, for instance, the experience of life as a squirrel is probably objectively worse, roughly speaking, than the experience of life as a human, at least from where you stand now. It may be less clear in the fictional case of experiencing life as a vampire. The issue of evaluation from different standpoints is one I take up in the discussion of Ullman-Margalit's view below.

having a child will objectively enrich your life in significant ways, be instrumental in bringing a valuable human life into the world, help create in the world a loving family bond between you and your spouse, etc. On the basis of these objective values of the experience, you can know that the experience of having a child will be better than, say, the experience of being skinned alive. So even if you can't know the subjective value of the experience, you can assign a rough value to the experience on the basis of its objective value.

How the objective and subjective values of an experience interact to determine the overall value of an experience is a substantive matter for axiological theorizing and will include investigation of any organic unities among them with respect to the value of the experience overall. But the key point is that we cannot simply assume that if we don't know the subjective value of an experience but know its objective value, that we can't assign it a rough value overall. We might be able to.²⁴

None of this is to deny that there could be some cases in which ignorance of the subjective value of an experience blocks knowing its overall value, even if you know its objective value. The overall value of *some* experiences, such as that involved in having a sumptuous meal, might turn mostly on the subjective value of how we experience them.²⁵ So if you don't know the subjective value of the experience, you can't assign an overall value to it. But the kinds of experiences that most plausibly figure in paradigmatic transformative choice are not like this. When you choose whether to gain sight or hearing, or whether to become a parent, *how it feels to you* is only one small factor in determining the rationality of the choice. What matters significantly more is the way things are, not how you experience them.

Two points help to show why this is so.

Suppose that the experience of having a child will make your life go objectively great—significantly better than it would go if you were to remain child-free. An evil demon, however, plays with your brain and makes the experience of having a child for you a drudge. Paul explicitly says that her interest is in veridical experiences only. But if you experience having a child as drudgery, *and* if that experience is veridical, then there is drudgery in your life—an objective disvalue that is entailed by your veridical experience. One way the subjective badness of an experience might seem to be important to the value of the experience overall, then, is by surreptitiously assuming the objective badness that subjectively bad *veridical* experiences entail. In order to isolate subjective value *per se*, we should assume that your experiences are *nonveridical*, and then ask whether ignorance of the intrinsic badness

²⁴ Just as the question of how some knowledge of what an experience is like might contribute to knowledge of what an experience is like overall is a substantive matter, as we saw in the last section.

²⁵ This can also be disputed; the value of the sumptuous meal may be primarily a matter of how its enjoyment objectively conduces to your well-being. But we don't need to take a stand on that question here.

of your nonveridical experience could block assigning an overall value to the experience on the basis of the objective values of that experience. Once we suppose that an experience does not reflect reality, but is like a dream of hallucination, then, I believe, the import of its subjective value significantly diminishes. If the experience of having a child will, objectively speaking, make your life significantly better, and if what matters to the choice is your well-being, then it seems plausible that you can assign a rough value to the experience despite not knowing the subjective value of your experience.

But what if the evil demon is especially cruel and makes your nonveridical experience of having a child not boring but hellacious torture? Again, the experience isn't veridical, but doesn't the possibility that the experience is torturous show that the subjective value of the experience can make a significant contribution to the overall value of having it? If you can't know whether the experience of having a child will be torturous, even if the experience will be nonveridical, doesn't that block assignment of a rough overall value to the experience on the basis of its objective value? To answer this question, we need once again to isolate the intrinsic subjective badness of the experience. The experience of torture, even if nonveridical, makes a significant contribution to the *objective* badness of your well-being. This is true for every subjectively good or bad experience. Suppose you have traumatic nightmares every night. Even if they don't reflect anything in reality, the unpleasantness of the experience can make your life go objectively worse. Similarly, an experience of hellacious torture, even if fabricated by an evil demon, can take an objective toll on your well-being. As many normative philosophers have pointed out, the objective goodness or badness of something can depend—causally, for instance—on your subjective experiences. So the subjective badness of an experience may amount to nothing more than the objective badness of it, where its objective badness depends in part on how the experience feels. Why shouldn't we think of the subjective badness of hellacious torture in this way? It is bad, not because of any intrinsic subjective badness but because of the way it objectively can harm your well-being.

More concessively, we might imagine a way to isolate the intrinsic subjective badness of an experience. Perhaps after the evil demon makes you experience, nonveridically, hellacious torture, God comes along and immediately expunges the experience from your life so that there is no objective harm from the experience. In this case, we might think that although there is no objective badness of the experience, the subjective badness of the experience remains. Now the question becomes, does the subjective badness of the experience block your assigning a rough overall value to the experience when what matters to the choice is your well-being? It would be strange to think that it does. By hypothesis, having a child will make your life objectively much better than it would otherwise be. What matters to the choice is how well your life goes. Even if what partly contributes to your well-being is the subjective value of your nonveridical

experiences, once we isolate this subjective value from the objective value that such subjective experiences can have, it's hard to see how the intrinsic subjective badness of hellacious torture—understood apart from the way it might objectively harm you—can block your assigning a rough overall value to the experience. Note that this is not simply to assume that well-being cannot be partly a matter of subjective value. It is only to underscore what many other philosophers have persuasively argued that a theory of well-being according to which the *only* thing that makes your life go well is having subjectively good experiences is wrongheaded.²⁶

We can leave our theorizing there because, as I'll be arguing in the next section, in any case, the focus on *experiences*—whether their subjective or objective value—is misplaced to begin with. If that's right, then the issue of whether we can assign a rough value to having an experience based on its objective value alone becomes a side issue.

I've argued that transformative choices aren't about choosing the subjectively best 'what it's like' you can get. They are about choosing what you have most reason to choose, and what you have most reason to choose isn't simply about getting the subjectively best experience. Insofar as the value of the experience matters, it is not only the subjective but more significantly the objective value of the experience that matters. Since the objective value of an experience matters more than its intrinsic subjective value, you can plausibly assign a rough value to an experience even if you don't know its subjective value. And if you can assign an epistemically transformative experience a rough value, you can rationally choose it on that basis.

Nor should we think that the roughness of the value you can assign raises any special difficulties. Normative rational choice theory has the tools to represent rough value (e.g., Pettigrew 2014, Hsieh 2005). And even if two alternatives both have only rough value assignments, there are many ways to understand rational choice on the basis of those rough values. You might rationally treat them as roughly equally good or 'on a par.'²⁷ Since you can assign a value to epistemically transformative experiences,

²⁶ One reading of Paul's target is the view that the rationality of transformative choices turns solely on the subjective value of experiences. I think this is an untenable view of well-being and that even standard forms of expected utility theory need not embrace it. See, e.g., Broome 1991. Another reading is that subjective value is a contributor to the overall value of the experience. This is a more charitable reading of standard expected utility theory and is the reading I suppose in my argument that ignorance of subjective value does not entail ignorance of overall value. Note that this more charitable reading is compatible with preference-satisfaction accounts of well-being since preference-satisfaction accounts do not implausibly presuppose that it's only the subjective feel of preference satisfaction that makes your life go well but the fact of preference-satisfaction.

²⁷ There is a separate set of issues I don't have space to discuss here about how an assignment of rough values to each option or rough relative value to the set of options determines rational choice. So, for example, if we understand rough value as a closed interval range of reals, and two options are represented by overlapping intervals, which is it rational to choose? Economists and decision theorists have offered different answers to this and related questions.

the putative problem they pose for the assumption that rational choice is based on the value of the alternatives disappears.

2.1.3 *Experience Isn't Primarily What Matters in Paradigmatic Transformative Choice*

Epistemically transformative experiences, we've argued, don't pose any threat to standard approaches to rational choice. Strictly, they include only the sole de novo exercise of a new-to-human-kind basic capacity, not experiences typical of paradigmatic transformative choices. But even if we allow, for the sake of argument, that epistemically transformative experiences include de novo exercises of basic capacities that have been exercised before, such as seeing for the first time, and higher-level complex experiences, such as becoming a parent, we find that those experiences—assuming that they are ones we can't know what they're like—are nevertheless experiences to which we can assign a value relevant to rational choice. Since we can assign a value to epistemically transformative experiences, such experiences don't undermine the assumption that rational choice is based on the value of the alternatives.

There is a more serious worry about the focus on epistemically transformative experiences. Is transformative *experience* the right phenomenon on which to focus inquiry into transformative choices? Choices that can change who we are—e.g., a change in careers, divorce, giving a significant portion of our wealth to charity, and so on—aren't primarily about choosing the option that we think will deliver the best experience, whether 'best' is understood objectively or subjectively. As with any choice, a transformative choice is one in which you should do what you have most reason to do, and what you have most reason to do isn't typically a matter of how an *experience* will be for you. The objective and subjective value of an experience may of course be one relevant factor in determining what you have most reason to do, but it is arguably typically of only modest significance. To think otherwise would give experience a distorted importance in understanding choice.

Return to Mike May. At the outset of the paper, we suggested that May's choice was about whether to have an extended transformative experience—to experience life as sighted. But we can now see that there is a better understanding of his case. What mattered in May's choice was not simply the experience of being sighted—the subjective and objective value of that experience—but the goods (and bads) he would thereby gain (and lose) in his life, not just from the experience of being sighted but from *the fact of being sighted*. What transformed him was not the *experience* of being sighted but the fact of being sighted and its many upshots. Of course May could have been the sort of person who made his life a matter of chasing the best subjective experiences. But he would then have misunderstood the nature of his choice about whether to regain his sight.

The same goes for choosing whether to have a child. What matters in such a choice is not getting the objectively and subjectively best experience. Having a good experience is relevant to the choice, but the choice about whether to have a child is primarily one about whether to bring a being into your life and into the world, not about what you are to *experience*. Nozick taught us long ago, experiences aren't what matters in human life. What matters is how things are, not how we experience them. We choose between different ways things are to be, not between different experiences we might have. The focus on experience misses what is important in transformative choice—transformation isn't concerned with how we experience life but with how are lives actually are. The focus on experience makes transformative experience a phenomenon of creatures in Nozick's experience machine or of brains-in-vats. We should broaden our understanding of transformation so that not only experiences, but events in the world can change who we are.

For related reasons we might question Paul's suggested 'solution' to the putative problem raised by epistemic transformation. Recall that, according to Paul, since, by hypothesis, you can't assign value to an epistemically transformative experience, you can't rationally choose whether to have it. If such choices can be rational, then the natural conclusion to draw is that standard approaches to rational choice need revision. But Paul instead suggests that we reconceive the choice of whether to have a child not as one about whether to have an epistemically transformative experience but about whether to have a certain kind of knowledge, in particular, knowledge of what it's like to have the experience of having a child. Since we can assign value to knowing what something is like, such choices would raise no difficulty for standard approaches.

We might doubt Paul's otherwise interesting suggestion because, however odd it might seem to conceive of the choice of whether to have a child as a choice about whether to have a subjective experience, it's odder still to conceive of it as the choice of whether to have knowledge of what something is like. Recasting the choice in these terms seems to misunderstand the nature of the choice. Again, none of this is to deny that some agents might mistakenly think of their choices in these terms. But we should not build a theory of transformative choices on misunderstandings of what such choices involve. At any rate, it isn't a solution to a problem to find a related phenomenon in the neighborhood that doesn't raise the problem.

There are undoubtedly transformative choices about whether to have a transformative experience. You might be poised to ride Full Throttle at Six Flags, an experience that will transform you into someone who is no longer afraid to try activities many would consider terrifying, or be contemplating whether to try Vegemite, an experience that will transform you into a Vegemite fanatic. In these cases, the choice might plausibly be

one about whether to have a transformative experience.²⁸ And for some people, the experience of having a child can be transformative, changing them from a me-first person into someone who can care for another for her own sake. But this isn't to say that their choice of whether to have a child is a choice about whether to have a transformative *experience*, even if it is an experience that transforms them. Transformative choices aren't typically about which subjective feel would be best for you but about ways the world—including you as an agent—are to be. What you are choosing between when you make a transformative choice isn't typically experiences but ways your life might go.

2.2 Event-Based Transformative Choices More Broadly Understood

Transformative choices aren't typically choices between different experiences. Indeed, what matters in a transformative choice isn't simply getting the best subjective feel; what matters is the value of events in the world downstream from the choice, which may include experiences but need not.²⁹ What matters in May's choice about whether to see again are not only the objective and subjective values of the *experience* of seeing but also the objective *goods* (which we can characterize in terms of events) he will have in his life if he is sighted. Indeed, it makes sense to think that it was not the experience of seeing that primarily transformed him but other events, like communing with his wife over a beautiful sunset, responding to visual feedback from his children, and learning new skills that gave him greater opportunities that did the transforming work. So we should allow not only that events, broadly understood, are relevant to assessing the value of an option but that they can do the transformative work in a transformative choice.³⁰

Once we move to events, however, we must abandon Paul's argument that transformative choices raise a problem for standard approaches because her argument crucially turns on the idea that certain kinds of *experiences*—epistemically transformative ones—preclude evaluation and therefore rational choice on the basis of their value. We turn instead to events more broadly, including experiences that aren't epistemically transformative and ask: Can the choice of an alternative that has as a downstream effect a transformative event—experiential or not—pose a problem for standard approaches to rational choice?

²⁸ As I've argued above, if these choices involve one's well-being—that is debatable—then they probably aren't choices *simply* about whether to have an experience but rather choices about whether to make one's life go a certain way. Choices simply about whether to have a certain experience are quite limited indeed.

²⁹ Paul herself sometimes slips into talk of nonexperiential events: e.g., “the process of having a child changes people” (2014, 90 and *passim*).

³⁰ Thanks to Louis Philippe for urging me to clarify the connection between what matters in a transformative choice and what events might do the transforming work.

Edna Ullman-Margalit (2006; 2007) proposes an event-based view of transformative choice.³¹ She argues that a transformative choice is a choice to do something that changes your utility function. In particular, Ullman-Margalit thinks that events downstream from choice change your ‘rationality base,’ the beliefs and desires that form the basis of your reasons, so that your utilities after these downstream events are discontinuous with your utilities before them. You are transformed—you change who you are—by events downstream from choice that change your utility function (2006, 167–168).³² Transformative choices, she thinks, are ones in which “the old ‘rationality base’ is replaced by a new” one (168). Paul shares Ullman-Margalit’s view of transformation; she holds that when an epistemically transformative experience is personally transformative, your utilities before the experience are discontinuous with your utilities after it. So, like Ullman-Margalit, she thinks that personal transformation occurs when your utility function changes.

By way of example, Ullman-Margalit, like Paul, tells the story of a person contemplating whether to have a child. At one point in time he doesn’t “want to become the ‘boring type’ who has children.” But then he decides to have a child. The story continues: With time, “he did adopt the boring characteristics of his parent friends—but he was happy!” Prior to his having a child he “did not approve of the personality he knew he would become if he has children; his preferences were not to have New person’s preferences. . . . As New Person, however, not only did he acquire the predicted new set of preferences, he also seems to have approved of himself having them” (167, footnote 10). Transformative choices, Ullman-Margalit explains, “are choices that straddle two discontinuous personalities” (2007, 60). So a narcissist, for example, might be transformed by having a child because she has never before cared for anyone else for her own sake. Before having a child, the narcissist’s utility function would value self-interested pursuits above all else, but after having a child, her utilities may reflect appreciation of the greater intrinsic value of the well-being of others over some of her pursuits. She is personally transformed by having a child since “what is rational for [her] to do beyond this point [of having a child] is different from the basis for the rationality assessment of [her] actions prior to that point” (2006, 168).

Ullman-Margalit goes on to suggest that such choices raise a problem for standard approaches to rational choice, and in particular, normative expected utility theory, because there is no stable utility function from which to determine the value of the alternatives. Before having a child,

³¹ Strictly speaking, Ullman-Margalit is interested in cases she calls ‘opting,’ cases in which we ‘make a leap of faith’ in choice. But her discussion of ‘opting’ focuses on cases that have as their first feature that they are ‘transformative.’

³² Ullman-Margalit doesn’t distinguish between the choice and its downstream effects, but it’s clear from what she says that she assumes that what changes your utility function are effects downstream from choice.

the alternative of having a child has less value than being child-free, but after having a child, it has more value. Since there is no single, correct value of alternatives but only value relative to a set of preferences, utility function, or ‘personality,’ standard approaches to rational choice break down (167–168).³³ In response to the problem posed by such choices, Ullman-Margalit suggests that all the agent can do is make “a leap of faith,” that is, arbitrarily ‘opt’ for one option over the other (169).

The challenge Ullman-Margalit poses, however, is a one that standard approaches—and normative expected utility theory in particular—can handle. Perhaps the most famous case in the neighborhood comes from Jon Elster’s (1979) discussion of Ulysses and the sirens. Prior to hearing the sirens’ song, Ulysses values his life and that of his men more than hearing the song. After hearing the song, his preferences support the reverse valuation. Elster suggests that Ulysses should exogenously bind his future irrational self to the ship’s mast so that he cannot wreck his ship on the rocks when going mad from the sirens’ song. Rational choice theory, broadly understood, has no difficulty with the case.

Of course Ulysses’s case, since it involves future *irrational* preferences, is not strictly analogous to the challenge that Ullman-Margalit poses. Ullman-Margalit’s choice of whether to have a child involves different sets of putatively *rational* preferences, that is, two perfectly rational but discontinuous and incompatible utility functions before and after the event of having the child. A closer analogy might be provided by Parfit’s (1984) Russian nobleman. As a young man, Parfit tells us, the nobleman is a socialist who wants his old, richer, future self to distribute all of his wealth among the peasantry. But his older, conservative self prefers to keep most of his wealth for himself. It’s not irrational—and we can suppose, not immoral—for the aged nobleman to keep most his wealth, after having discharged whatever duties he might have to be charitable, but nor is it irrational for him to give it all away as his younger self would wish. We just have two perfectly rational but discontinuous and incompatible utility functions before and after life’s intervening events. Parfit suggests a solution akin to Elster’s—the young nobleman should exogenously bind his future self, in this case, Parfit suggests, by creating a contract that only his wife can revoke and getting his wife to agree never to revoke it. This is because, Parfit seems to suggest, the younger self is the ‘real’ self just as the ‘real’ Ulysses is the one not driven to madness by the siren’s song.³⁴

³³ Paul makes the same claim about transformative choices, except that her concern is with epistemically transformative experiences that lead to a change in your utility function while Ullman-Margalit has a broader view of events that may lead to a change in your utility function. Since Paul’s claim is an instance of Ullman-Margalit’s, my discussion of Ullman-Margalit is intended also to be a discussion of Paul’s similar claim.

³⁴ Christine Korsgaard (2009, 202) criticizes Parfit’s solution as requiring the nobleman to put his wife in an “impossible position” because she must wrong either her young or her older husband.

But when you choose to have a child, your child-free self needn't be more 'you' than your parent self. You have two different sets of preferences, both of which reflect different 'you's' at different stages in your life. And both, we can suppose, are rational both at the time that they are had and from the time of choice.³⁵ So how are you rationally to choose whether to have a child when your preferences may radically change after the events involved in having a child from what they are before you have a child?

One solution is to take your present preferences as the basis for determining your rational choice. It needn't be that your present self is the 'real' you; instead it could just be a default rational principle that you must choose according to your present utility base rather than one you expect to have in the future. Note that privileging one's present preference profile would be compatible with those preferences reflecting knowledge of foreseeable future facts, including your future preferences.

Suppose that, without thinking about how having a child might change your preferences, you now prefer to remain child-free, primarily because you are concerned about the impact being a parent might have on your career. You then read Ullman-Margalit and Paul, and you start to think about whether your preferences will reverse themselves if you end up having a child. You talk to your parent friends who didn't want children but ended up having them and take note of their post-child preferences, look up empirical data on the number of parents whose attitudes toward their work changed after having a child, research neurological studies claiming that post-child-birth, you will likely have a hormonal imbalance that contributes to a desire to have yet another child, and so on. While you don't know whether your preferences will change after having a child, you know that it might. You can then take this fact into account in forming your current preferences. The fact that you might be ecstatic about having a child after having one might be grounds now for you to prefer to take the risk of having a child even though the thought, now, of having a child fills you with anxiety and dread. Or that fact might simply change the strength of your current preference to remain child-free, so that you are closer than you were before to preferring to have a child, and that preference could later reverse as new events unfold in your life. Or the fact that you might be miserable with a child, but self-delusional, might strengthen your preference to remain child-free, if, for example, you now strongly prefer never to be self-delusional.

³⁵ The claim that both sets of preferences can be rational at the time of choice might be doubted. In discussing Parfit's Russian nobleman, for instance, Christine Korsgaard argues that the young nobleman must treat the preferences of his future self as irrational (2009, 202–204). I have argued elsewhere that two incompatible evaluative orderings according to 'given' values or reasons can both be rational if they are 'on a par' (e.g., Chang 2013b). If you don't believe in parity, then the two incompatible orderings can also be understood as different 'sharpenings' of a (nonsemantic) indeterminacy in which one option is better supported by rationality. In any case, the more interesting version of Ullman-Margalit's and Paul's case assumes that both can be rational at the point of choice.

Appealing to your fully-informed present preferences—informed by your knowledge of your future preferences—to determine rational choice is in fact the default view of standard rational choice theory. After all, very few people live and die with continuous utilities over the course of their lives. What Ullman-Margalit and Paul usefully point out is that, although typically a person's utility function changes slowly over time, it can also change quickly and dramatically by an event downstream from a choice. And if you know that your utilities might change dramatically in this way, the question arises, on what basis should you determine what it is rational for you to do? Standard rational choice theory has an answer: choose on the basis of your present preferences—informed by what you reasonably believe your future preferences will be. Of course the kind of normative weight you give to your future preferences when forming your present ones will depend on what your present preferences are. In the choice of whether to become a parent, for example, you might presently be extremely risk averse or you might be a daredevil. Your current attitudes will affect what normative role information about your future preference profile will have for you. But that is par for the course.

There are other ways rational choice theory could deal with choices about whether to undergo events that will change your utility function. Such theories might posit a 'master' utility function—the function your ideal self would have if it knew all the facts, present and future, relevant to any choice you might make in your life. This master utility function could then order the two sets of preferences with respect to any given choice. If, for example, the master utility function favors the ordering, (have a child, remain child-free), then, even if your present preferences favor remaining child-free, since the preferences you would have after you have a child are better—presumably with respect to your well-being—than your present preferences, the rational thing to do is to follow not your present preferences, but your future preferences. This is the rational thing to do because your future preferences reflect the preferences of your master utility function.

Or, rational choice theory could posit a principle of rationality according to which when you know that one of the options for choice will change your utility function, you should just 'wait and see.' Such a principle might counsel that you put off the choice if possible, or take an incremental approach to the choice by breaking it down into smaller sub-choices that can be made over time, which could have the effect of turning transformative choices into ordinary choices by which one is transformed over time.³⁶

The important point for our purposes is that standard approaches to rational choice have ways of dealing with choices with downstream effects that lead to a change in your utility function, personality, or point of view

³⁶ See also [Ullmann-Margalit 2006](#) who suggests that one strategy for dealing with cases of transformative choice is to break them down into smaller choices.

without compromising either of its two fundamental assumptions. While your personality may change, the rationality of your choice is based on the value of the alternatives. And while the value of the alternatives may be relative to different personalities, the rationality of the choice is based on the value of the alternatives nonetheless. Until these approaches are shown to be untenable, the problem supposedly posed by transformative events downstream from choice in evaluating the value of the alternatives is one that standard approaches to rational choice can solve.

3 Choice-Based Transformative Choices

As we've seen, the first assumption of standard approaches to rational choice, namely, that rational choice is determined by the value of the alternatives, can be kept intact even in the face of choices about whether to undergo events that will change our utility functions, personalities, or point of view. While it's plausible that many transformative choices will be event-based, such choices don't seem to raise any genuine challenge to standard approaches. This is not to say that there are no such choices, but only to suggest that perhaps they aren't the most interesting kind of transformative choice around.

In the rest of this paper, I want to begin to lay the groundwork for an account of choice-based transformative choices, choices in which *the choice itself* does the transforming work. As we'll see, the account of choice-based transformative choices challenges the second assumption of standard approaches, namely, that choice cannot be *that in virtue of which* we have reasons or it's rational for us to choose something.

But first we need to try to clarify what we mean by 'transformation.' Both Ullman-Margalit and Paul suggest that transformation is a change in your utility function, personality, or point of view. I agree, but I want to tweak and expand their expected-utility-theory-based idea of transformation, in part so that it is not wedded to any particular substantive view about how normativity is to be modeled. I'll use the term 'reasons' to indicate considerations that count in favor of an alternative, whatever those considerations might be—preferences, evaluative facts, duties, excellences, and so on. (The points I want to make can also be put in terms of 'values' and even pro tanto 'duties' (with appropriate bells and whistles) but I will stick with 'reasons' since, at least to my ears, that is the most neutral-sounding normative term.) As I'll be suggesting, you are transformed—change who you are—whenever your reasons change in a way that alters your normative character. And in choice-based transformative choice, you alter your normative character in a distinctive way: you change 'who you are'—change the reasons of your normative character—*in virtue of* your choice.

3.1 Transformation

A transformative choice changes ‘who you are.’ But what might this involve?

First, we ask, Who is transformed? A *person* gets transformed in some way, but not into a nonperson—not into a bat, or squirrel, or vampire (if vampires aren’t people, teen fiction notwithstanding). So a person must remain before and after a transformative choice. Furthermore, *you*, whatever you metaphysically are, must remain before and after a transformative choice. That’s how we can intelligibly say that *you* have undergone a transformation. Transformative choices don’t alter your numerical identity or your personhood; they merely change the way *you*, the person, are from how *you*, the person, were before the choice. A transformation that meets both these conditions is what we might call a *personal* transformation. We can understand the transformation of choice-based transformative choices as personal in this sense.

This way of understanding transformation ensures that transformative choices include paradigmatic cases of transformative choices such as whether to regain one’s vision, have a child, change careers, and so on, that are part of ordinary human life. Excluded are ‘radical transformative choices,’ choices that either transform you from a *person* into a non-person or *you* into something that isn’t you—such as choices of whether to turn into a tree or cockroach or ghostly spirit. Such choices no doubt raise interesting questions for the rationality of choice but aren’t relevant here.³⁷ If being human is inessential both to being you and to being a person, transformative choices will include those about whether to undergo some significant enhancement, such as one that would enable you to live for 500 years or would allow you to jump to the moon, that might transform you from being human to being ‘superhuman.’

Second, we ask, What is the feature of *you*, *the person* that gets transformed in a transformative choice? As we’ve seen, both Paul and Ullman-Margalit understand personal transformation as a change in your utility function—your preferences over alternatives is different before and after a transformative event. Both of them plausibly understand transformation as a *normative* phenomenon. Although May, upon regaining sight, changes dramatically in nonnormative ways, these nonnormative changes aren’t transformative in the sense of interest unless they subvene or ground a *normative* change. As I will put it, you change ‘who you are’ when you change your *reasons*, which includes a change in the strength of your reasons as well as coming to have new reasons you didn’t have before. After

³⁷ See Kemp 2015 for a defense of the idea that what I am here calling ‘radical’ transformation is not something you can undertake or will but is something that happens to you. The transformations of interest here are not radical in this way since there is a you—whatever that entails—that is a person, to whom we can attribute and who can, in principle, undertake the transformation.

his transformation, May's reasons to read in Braille are weaker than they were before, and he has normative reasons to go to museums that he didn't have before. The transformation of transformative choices, I suggest, is essentially a matter of changing your normative reasons.

The transformation of interest is also objective; the reasons you come to have are reasons you *in fact* come to have, not reasons you merely believe you have. Sometimes we can be transformed and realize it much later. You might choose to become a parent and only after a few months of changing diapers realize that you are no longer the self-centered, me-first person you were before. Similarly, you can falsely believe that you are transformed—say, by attendance at an awesome rock concert—but in fact you are much as you always were. This implies that in a transformative choice, you may not know that an alternative you choose will transform you.

Now, for transformation to take place, the change in your reasons can't be a change in any old reasons you happen to have. The choice needs to change 'who you are.' I suggest that 'who you are' is understood normatively, as your normative character or normative identity. 'Who you are' is who you are *qua* normative agent, and the reasons that belong to 'who you are' are the reasons that make up your *normative character*, roughly your normative personality, or 'the sort of person' you are, normatively speaking. Transformation, then, is a matter of changing your normative character, by having reasons that make the normative you different from how you were before.

We'll have more to say about normative character shortly. But one possible misunderstanding should be put aside here. It might be thought that the normative you is your 'deep self' or, as it is sometimes put, 'who you really are, deep down inside.'³⁸ But I think identifying your normative character with your deep self, assuming that this idea is even coherent, would be a mistake. This is because transformation may be a much shallower phenomenon than changing who you 'really are, deep down inside.'

Suppose, for instance, that your deep self is given by the reasons that play an organizing, structuring, executive, or some higher-order role vis-à-vis your other reasons. They might be general Bratmanian 'self-governing' policies such as 'be consistent' or 'always do what's best for me,' or substantively thick Aristotelian 'master ends' such as happiness or flourishing, that structure your other reasons.

But your higher order policies aren't plausibly the reasons that determine your normative character. This is for two reasons. First, we need a story as to why your actual policies determine your normative character (or for that matter, your deep self) as opposed merely to being the policies under which you happen—perhaps because of some brainwashing—to be

³⁸ I assume for the sake of argument that the idea of 'who you really are, deep down inside' is coherent.

operating. The Bratmanian story, according to which one's self-governing policies have special status because they ensure one's metaphysical identity over time as an agent, won't help us here because we have put aside 'radical' transformation, the transformation that alters your metaphysical identity. In any case, it isn't plausible to think that whatever policies determine you to be a single metaphysical agent acting over time are the same policies that make up your normative character, since your normative character is rather more specific.

Second, since transformative choices change your normative character, the reasons that determine that character should be reasons that are changed in a transformative choice. But paradigmatic transformative choices don't require changes in your general policies, master ends, or higher-order normative principles. You might have a general policy to do whatever you believe is morally right, to love thy neighbor, or even to look out only for yourself and your loved ones. When you choose to have a child, live in the countryside or change careers, your general policies and master ends may well remain intact. It's not that, having moved to the countryside, you become morally reprobate. And yet a move to a rural life can change your normative character. You are no longer the always-overcommitted multi-tasker who never stops to smell the roses. You're now the sort of normative character who sits on a porch swing, doing nothing while enjoying the sunset. In short, can change who you are, normatively speaking, by becoming a parent, or a country-dweller, or a graveyard-shifter, by changing your reasons without changing the reasons of your deep self.

There are other possible suggestions about the reasons that determine your deep self, but I think all of them will fail as reasons that determine your normative character. This is because your normative character is shallower than 'who you really are, deep down inside.' Some people seem much the same, deep down, throughout their lives. Nevertheless, they can have different normative characters at different point in their lives. Parfit's Russian nobleman might be such a person. Nelson Mandela and Mother Teresa might be others.

When you make a transformative choice, you can make yourself into a person who would make you cringe at some other point in your life. Such transformations might involve changes in the deep self. But they need not.

So far we've said that transformative choices involve changing your reasons, and in particular, the reasons that determine your normative character. But there's a third, crucial question we should ask about transformation. How does a transformative choice change your reasons? Since the most interesting change in your reasons is coming to have new reasons you didn't have before, I'll focus on that case.

Standard approaches to rational choice explain how you come to have reasons by appealing to something essentially independent of your choice—typically some relation between an alternative and your mental states, a normative fact about the goodness of alternative, or the normative fact

that whatever is the reason is a reason. To keep things simple, I'll focus on explanations of how you come to have reasons that appeal to normative facts, though the same points can be made in terms of preferences and other mental states.

An example will help. Suppose you could have a burger or salad for dinner. You choose to have the burger. As a consequence of choosing to have it, you order a burger. Your choice of the burger, and the subsequent action of ordering it, changes the facts going forward. You've now ordered the burger, you'll be having red meat for dinner, and so on. The fact that you've ordered a burger (along with pre-existing facts about your tastes and so on) gives you a reason you didn't have before, a reason, say, order a cold Heineken (and not a hot eggnog) as a dinner beverage.

Why does the fact that you've ordered a burger give you a reason? That is, *in virtue of what* does the fact that you've ordered a burger give you a reason to order a beer? According to standard approaches, the answer is given by a normative fact, viz., *If you order a burger (in such-and-such circumstances), you have a reason to order a cold beer*. A normative fact that connects downstream effects of your choice with reasons is that in virtue of which you have those reasons. We might call such facts 'grounding' normative facts. When you satisfy the antecedent of such facts, you thereby 'trigger' the consequent reasons of such facts.

There are many other examples. There is a normative that, *If you punch someone in the nose (in such-and-such circumstances), then you have a reason to make amends*. So if you fulfill the antecedent condition of this normative fact, that is, if you punch someone in the nose, you have a reason to make amends. You have that reason in virtue of the grounding normative fact according to which punching someone in the nose 'triggers' a reason to make amends. If you chose to have a child, as a downstream effect of so choosing you may have a child. If you have a child, then you have a reason to nurture and care for her that you didn't have before. You have this reason in virtue of the normative fact that, *If you have a child (in such-and-such circumstances), then you have a reason to nurture and care for her*. Again, your having a child 'triggers' a reason to nurture and care for her. According to standard approaches, you come to have reasons not in virtue of your choosing anything, but in virtue of a grounding normative fact that connects downstream effects of your choice with reasons.

Note that according to this standard explanation of how you come to have reasons, choice *per se* may be normatively irrelevant to your having new reasons. Even if you have a child by accident and not by choice, the fact that you have one triggers reasons to nurture and care for her in virtue of the same (or similar) normative fact.³⁹ If transformative choices change your reasons in this way, then choice may be irrelevant in explaining both

³⁹ Whether it's the same grounding normative fact turns on a matter of substance, viz., whether how you come to have a child—by choice or not—affects which reasons having a child triggers. In the same way, it's a substantive matter whether the fact of choice is a relevant fact that

what transforms you and how you are transformed. All the work is done by effects downstream from choice—the events, broadly construed—that trigger reasons in virtue of a grounding normative fact connecting those events to those reasons. Event-based transformative choices, then, can give you new reasons in the standard way. If we understand transformative choices on the model of event-based transformation, we take the ‘choice’ out of ‘transformative choice.’

Choice-based transformative choices, by contrast, put the ‘choice’ back into transformative choices. Choice is not only what transforms you but how are you transformed; that is, choice is *that in virtue of which* you come to have the reasons of transformation. You change your normative character through choice itself.

3.2 An Account of Choice-Based Transformative Choices

The idea that we come to have new reasons in virtue of choice itself may seem unpromising. After all, standard approaches assume that choice itself can’t ground your having reasons for obvious reasons: if choice can ground your having reasons, you can choose your way into any reasons you want. It would seem, then, that rejecting this assumption of standard approaches leads to intolerable bootstrapping.

But choice-based views of transformative choices can help themselves to a general theory of the grounds of normativity that avoids this result. Elsewhere, I’ve proposed that reasons can have one of two different grounds; for any consideration that is a reason there are two different sorts of consideration that can make it a reason, one a normative fact and the other an act of will. Correspondingly, because a given consideration can have two different sources, it can be two different reasons.⁴⁰ The view is ‘hybrid voluntarism,’ so called because it understands the grounds of practical normativity as a hybrid of two kinds of considerations that can make something a reason. Some reasons are grounded in something other than our wills—they are ‘given’ reasons because they are given to us—while other reasons are grounded in our wills—they are ‘will-based’ or ‘voluntarist’ because we create them through an act of willing.⁴¹

triggers reasons. In some cases, surely it is. But choice in such cases is not a ground but a triggering fact.

⁴⁰ Alternatively, we might say that the overall normative strength of a single—and single kind of—reason may have two different sources. To draw the starkest contrast between the two sources of normativity, I assume that reasons are individuated by both ‘content’ and source and so we have two ‘kinds’ of reasons, given and will-based. The substance of the view, however, can be put equivalently in terms of the normative strength of a single kind of reason having two different sources.

⁴¹ See [Chang 2009, 2013a,b](#). Talk in terms of new will-based reasons helps to underscore the fact that the source of the normativity of the reason is in the will, but the view can also be equivalently stated in terms of the will being that in virtue of which a given reason has greater (or lesser) strength than it had before.

Suppose, for instance, that your child wants a new toy. Why is this fact a reason to buy her one—in virtue of what is it a reason, assuming that it is? One answer is that satisfying her desires makes her feel safe and happy, which is good or valuable in some way.⁴² The fact that satisfying her desire is good is that in virtue of which the fact that she wants a new toy is a reason for you to buy her one. Notice that this explanation makes no appeal to your will. What makes something a reason is something other than your will—its being a reason is ‘given’ to you, not willed by you. It’s given by the goodness of buying her a new toy. So the fact that she wants a new toy is a *given* reason for you to buy her one.

A reason can also be a reason because of an act of your will. Suppose that you *will* that your child’s needs and interests are normative for you, or, as I will put it, you *commit* to her needs and interests by putting your will—your very self as a normative agent—behind those needs and interests. By committing to her in this way, your commitment can be *that in virtue of which* her desire for a new toy is a reason for you to give her one. Satisfying her desire for a new toy will serve her needs and interests by making her feel safe and happy. Your commitment to her needs and interests is thus what makes her desire for a new toy a reason for you. By committing to something, you can make something a reason for you to perform some action.

Notice that the same fact—that your child wants a new toy—can be the ‘content’ of two different reasons, distinguished by what makes the fact a reason. A fact can be a reason to do something in virtue of a normative fact, such that doing that thing is good or that the fact is a reason to do it, or it can be a reason to do something in virtue of your act of will—your commitment to something. You can have a reason to give her a new toy in virtue of the fact that doing so would be good, but you can also have a reason to give her a new toy in virtue of your commitment to her needs and interests.

Now if we left things there, will-based reasons would indeed lead to intolerable bootstrapping. Perhaps buying your child a new toy would be bad for her, but you could nevertheless bootstrap your way into having reasons to buy her a cornucopia of toys by committing to satisfying her every desire.

Hybrid voluntarism maintains that there is a hierarchy among your given and will-based reasons. It holds that your given reasons, so long as they don’t ‘run out,’ always determine what you have most reason to do. So if you have most given reason not to buy your child a new toy—if it would be bad for her, for instance—you shouldn’t buy her one, all things considered. When your given reasons ‘run out,’ however, your will-based reasons can determine what you should do, all things considered.

⁴² Or that buying her a new toy will satisfy your desire that she feel safe and happy. Kate Manne offers an interesting twist on this answer: buying her a new toy will satisfy *the child’s* desire to feel safe and happy. See [Manne Forthcoming](#).

Sometimes your given reasons will ‘run out.’ They can ‘run out’ in one of two ways, but for simplicity’s sake we can focus on just one of them. Your given reasons ‘run out’ if the options are comparable but neither option is better than the other and nor are they equally good. That is, they are *on a par* with respect to what matters in the choice. Or, put equivalently in terms of reasons: the strengths of the reasons are comparable, and they don’t favor one option over the other and nor are they of equal strength. Your reasons are *on a par*. Two items are on a par if they are comparable and yet neither is better than the other and nor are they equally good.⁴³

If your given reasons are on a par, they don’t tell you what you should do. You now have the normative power to commit to one of the options, thereby creating a will-based reason in favor of it.⁴⁴ The silence of your given reasons is one sense in which when your given reasons ‘run out,’ you can create will-based reasons in favor of an alternative. Your will-based reasons can, in turn, make it the case that you have most all-things-considered reasons to choose one option over the other.

Return to the choice of whether to have a child. Suppose, for simplicity, that what matters in the choice is simply your well-being (perhaps because you are a single parent, an unborn child has no well-being, you know that your child will be carbon neutral, etc.). Suppose that, with respect to what would make your life go best, having a child is better in some ways, remaining child-free is better in some other ways, and having a child is neither better nor worse than remaining child-free, overall. There are just different tradeoffs to be made whichever course you take. The options are on a par with respect to your given reasons. Typically options that are on a par will bear very different values while nevertheless being in the same neighborhood of overall relevant value.

According to hybrid voluntarism, if your given reasons are on a par, you can *commit* to some feature of being a parent that counts in favor of being a parent and thereby create for yourself a new will-based reason to have a child. Your new will-based reason, then, may make it true that you have most reason, all things considered—considering both given and will-based reasons—to have a child.

This hierarchy between given and will-based reasons eliminates the worry that you can bootstrap your way into any reasons you like. You can create will-based reasons only when your given reasons have run out. If you have most given reasons to remain child-free, say, because having a child will make you suicidal, then you can’t will yourself reasons that make it the case that, all things considered, you should have a child. Hybrid voluntarism holds that your given reasons have ‘first dibs’ in determining what you should do.

⁴³ See Chang 2002. I don’t have space here to discuss the possibility that the options are incomparable. I address that possibility in Chang 2012, *Forthcoming*.

⁴⁴ You can also create a will-based reason against choosing an option by committing to not having some feature of one of the alternatives in your life.

When you commit to some feature of having a child, such as forming a parent-child bond, you thereby choose to become a parent in a *thick* sense of ‘choose’: you select an alternative by putting your very agency behind it. By committing to forming a parent-child bond, you now ‘stand for’ forming such a bond—you, your very self as an agent, are *for* forming such a bond.⁴⁵ You can also choose to become a parent in the *thin* sense of ‘choose’; you choose in a *thin* sense whenever you select an alternative or merely intend to go for it without throwing your very agency behind the alternative. The thick sense of choice involves an act of will, a commitment of your normative self; the thin sense doesn’t. Crucially, choice-based transformative choices involve choice in the thick sense: when you choose an option, you are committing to that option by putting your agency behind its features.⁴⁶

We can now see how in choice-based transformative choices your choice can be both what transforms you and that in virtue of which you are transformed. In deciding whether to have a child, by hypothesis, the given reasons are on a par. You have the normative power to commit to one of the options or one of its features. You might commit to forming a parent-child attachment. That commitment just is choosing to have a child in the thick sense. That commitment then creates new will-based reasons for you to have a child, that is, your commitment is that in virtue of which you now have a new will-based reason to have a child. Your new will-based reason then interacts with your other, given, reasons and guides your choice in the thin sense. You may now have most all things considered reasons to choose to have a child. Your new will-based reason transforms you because it is a reason that determines your normative character. You are now the sort of person who has most all things considered reasons to have a child. Before the choice you were the sort of person for whom the reasons for having a child and remaining child were on a par. By choosing, you change the reasons that determine your normative character.

Your normative character then, that is, the normative you, is given by all the will-based reasons you have before the transformative choice. Those will-based reasons are the reasons that make you ‘who you are,’ normatively speaking. When you create a new will-based reason for yourself, you transform yourself into someone who has reasons she didn’t have before. These new reasons are reasons you have created for yourself through your commitment to something. In this way, through choice, you transform yourself from one normative character into another.

⁴⁵ For further discussion of this form of internal commitment that can be the source of reasons, see [Chang 2013a](#).

⁴⁶ Commitment needn’t be a conscious, deliberate steeling of the will. We can be unaware of and even deny commitments we have in fact made. Think of the man who has a self-conception as a swinging bachelor but who would sacrifice his life to protect the needs and interests of the person he’s made a life with for the last several decades.

The way you come to have new will-based reasons on this account, is incompatible with the second assumption of standard approaches to rational choice. You come to have reasons in virtue of choice itself, not in virtue of something essentially independent of choice. And to think that the choice itself could play the role of an event downstream from choice—that is, as something that fulfills the antecedent condition of a normative fact in virtue of which when you make the choice, you have a reason, would lead again to intolerable bootstrapping; there are no normative facts of the form, *If you choose x, then you have a reason to pursue x*. Thus, choice-based transformative choices, as I've proposed we understand them, require fundamental revision of standard approaches to rational choice.

Now, when faced with a transformative choice in which the alternatives are on a par, you might *not* commit to one alternative or the other. You might choose in the thin sense without choosing in the thick sense. If you choose without committing, you either 'drift' into an alternative or 'plump' for it. Drifting encompasses many different ways of selecting an alternative. You could drift by selecting by omission (you never seriously consider whether to have a child), of by taking the path of least resistance (all your friends are becoming parents, so it might be easiest for you simply to follow the crowd), or by allowing some emotion, such as fear, determine what you will in fact do (you are terrified that when you are on your deathbed that you will profoundly regret never having a child). When you drift into an option, you don't put your agency behind it or its features but let reasons of the world cause you to take one path rather than another.

You can also, instead of committing or drifting, 'plump' for an alternative, that is, arbitrarily select one option over another for no reason at all.⁴⁷ Choosing in the thin sense by drifting or plumping can, of course, lead to downstream events that are transformative. Importantly, however, both drifting and plumping are ways transformation can *happen to you* rather than be *by you*. When you choose in the thick sense, by committing to an option, you transform yourself by creating for yourself new reasons you didn't have before. Your choice does the transforming work.

It's worth noting that according to Ullman-Margalit and Paul's event-based views of transformative choice, at the point of choice, there is no comparative fact about the merits of the alternatives. Paul suggests that the alternatives are what I have called 'noncomparable' (2014, 102, footnote 55). Two items, such as fried eggs and the number nine, are noncomparable with respect to tastiness when there is a formal precondition for the possibility of comparability that is not met. In the case of fried eggs and the number nine, the formal condition is that the 'covering value,' that

⁴⁷ Plumping is not the same as 'picking' between two equally good alternatives; note that having a child and remaining child-free are, by hypothesis, not equally good or equally well-supported by reasons. When you pick between equally good alternatives, you choose for the reason that the options are equally good and it's better to select one than to remain poised between them like Buridan's ass.

is, the respect in terms of which the items are being compared, namely ‘tastiness,’ fails to cover both items: the number nine is not the kind of thing that can be tasty. Paul thinks that alternatives in a transformative choice are noncomparable because one of them involves an epistemically transformative experience (that is also personally transformative), which precludes assigning a value to it. Since there can be no value assigned to any alternative that is epistemically transformative, a formal precondition for the possibility of comparing them, namely that each alternative has some value, isn’t met. We’ve seen some reasons to doubt, however, whether epistemically transformative experiences preclude assigning a value to an alternative.

Ullman-Margalit, by contrast, seems to think that there is no comparative fact about the merits of the alternatives at the point of choice because their value is relative to the utility function that exists at the relevant time. Each alternative has value, but the value is different depending on when the evaluation takes place. Before you have a child, it’s better for you to remain child-free since your utility function favors that option. After you have a child, it’s rational for you to have the child—your utilities after the event of having a child favor your having the child. But it’s puzzling why we should relativize the rationality of our actions to the utility function that exists at the time. As a general principle, holding that the rationality of an action is relative to the utility function held at the time leads to untoward consequences. It would make, for example, a course of action where you undermine all your other aims and projects rational so long as your utility function at the time favored doing so. And if we tried to constrain your utility function, such as by restricting the ways in which you can come to have it, we will end up denying as rational certain intuitively rational choices to undergo certain transformations.

In any case, there is a more worrying puzzle. Both Ullman-Margalit and Paul seem to think that at the point of choice, both alternatives are *rational*. If they didn’t think this, then transformative choices would amount to an instance of Elster’s Ulysses case: at the point of choice, one alternative is irrational, and so one should bind oneself in a way so as not to choose it. Standard approaches have no difficulty accounting for the rationality of such choices. If, instead, the alternatives are both rational at the point of choice, then how could there be no comparative fact about how the alternatives relate? Paul explicitly denies that the alternatives are *incomparable*, and I think she is right to do so.⁴⁸ Ullman-Margalit doesn’t say much about how the alternatives in a transformative choice relate, but she seems, like Paul, to assume that they aren’t equally good. If the alternatives are neither incomparable, nor equally good, and one is not better than the other but both are rational, then what relation holds between them? They are, I suggest, on a par.

⁴⁸ See Chang 2012, Forthcoming.

The options involved in transformative choices can evaluatively relate in one of three ways.⁴⁹ First, one option might be better (or good enough), in which case it is rational for you to choose it. You might, for instance, choose not to wander into a dangerous neighborhood where it is likely that you will be set upon by a gang of thugs. You know that being beaten up will be traumatic and transformative—and *worse* than not being beaten up. So you rationally choose not to have the transformative event in your life because having it is worse than the alternative. Or, to take the reverse case, you might face a choice of whether to claim your lottery winnings. Winning the lottery is a mixed bag, but overall, it's better to have the money than to forgo it. If you choose to claim your winnings, you'll undergo transformative events downstream from that choice. You rationally choose to have a transformative event in your life because it's better than the alternative.

Second, your options might be equally good with respect to what matters in the choice. This will, however, be rare in cases of transformative choice. If one option isn't better than the other, then it's unlikely that they will be equally good. It might be that for you, for example, having a child is neither better nor worse than remaining child free with respect to you well-being. Does it follow that they are equally good? We can run a test to determine whether they are. If we improve (or detract) from one of the options a bit—perhaps we throw in a part-time nanny if you have a child—does it necessarily follow that having a child is now better than remaining child-free. We can imagine a scenario where it doesn't necessarily follow. If this is right, then neither option is better than the other and nor are they equally good. For if they're equally good, an improvement in one, even if small, must make the improved option better (Chang 2002). But this doesn't plausibly hold for transformative choices between options, neither of which is better than the other. Moreover, if the options are equally good, it would be intrinsically rational for you to flip a coin between them. It seems odd to think that it's intrinsically rational for you to flip a coin in deciding whether to become a parent (Chang 2012).

Third, the options might be on a par. Typically, items are on a par when, with respect to some things that matter in the choice, one option is better, with respect to other things that matter, the other option is better, and yet neither is better than the other overall. For many people, remaining child-free and becoming a parent will be on a par. I've had more than one philosopher-friend say to me: 'Having a child might never allow me to achieve what I want to in my work, on the other hand, it might allow me to have a life enriched in ways that I can't now forsee. I know that whichever option I choose, I will be transformed into a different sort of person than how I am now.' If the alternatives are on a par, it would be a mistake to

⁴⁹ There are also the options that they are incomparable or indeterminately comparable due to vagueness in language. I argue against this options elsewhere. See [Chang 2002, 2012, Forthcoming](#).

continue search for reasons in the world that might make it the case that one is better; if they're on a par, there are no such reasons. Instead, since your given reasons are on a par, you have the normative power to choose one of the options, in the thick sense, by committing to one of its distinctive features. When you commit to an option, you create a will-based reason for you to choose it. You've now changed your reasons—you used to be the sort of person for whom having a child and remaining child-free were on a par. After committing to remaining child-free, you thereby become the sort of person for whom there is most reason to remain child-free. Your commitment changes your normative character.

Choice-based transformative choices capture paradigmatic cases of transformative choice. We've already looked the case of choosing whether to have a child. What about choosing whether to have a cochlear earplant or an operation that will restore your vision? Return to Mike May. Being unsighted involves having certain goods that are precluded if you are sighted, and vice versa, and perhaps those goods are, overall, in the same neighborhood of value although very different in value.⁵⁰ Perhaps May faced a choice in which his options were on a par. On the choice-based view of transformative choice favored here, May could transform himself by committing to the goods of being sighted, and thereby give himself most reason to undergo the operation that would restore his sight. He could transform himself by his choice by creating for himself new reasons he didn't have before in virtue of that choice, and perhaps thereby making it true that he is someone who now has most reasons to become a sighted person.

Thus when the Hollywood plastic surgeon gives up her luxurious lifestyle to volunteer in a war torn region, she commits to features of being a volunteer which then creates new will-based reasons to live differently than she has before. These new reasons change her normative character. When Gauguin chooses to abandon his family for his art, he creates a will-based reason to pursue his art and changed his normative character decisively into what it was. When Sophie, faced with the choice about which of her children to send to the gas chambers, creates a will-based reason to save Jan rather than Eva. It's this commitment that arguably tortures her for the rest of her life. And when the philosophy chairperson of the story with which we began this paper chooses wisdom over truth and money, the goods are on a par. But by committing to wisdom, she now is the sort of person for whom wisdom matters more than truth or money. What makes

⁵⁰ It might be difficult for sighted people to understand the depth of the goods gained by being congenitally blind. One grows up with different capacities that sighted people cannot share. Of course, once you are sighted, it seems much worse to be blind, especially if you cannot then achieve the special capacities of the congenitally blind. May was born sighted, became blinded at a very young age, and had the opportunity to become sighted at a late age, so his case is correspondingly more difficult.

the story droll is the suggestion that for such a person, money should be chosen instead of wisdom.

Again, this is not to say that there aren't transformative choices of the event-based kind. You can be transformed from a fun-loving, happy-go-lucky person into someone withdrawn and fearful by a traumatic event downstream from a choice to, say, pursue a dangerous career. Some transformative choices aren't between options that are on a par, but between options one of which is better than the other. In such choices, there is no problem for standard approaches since the rational thing to do is to choose the better option. Indeed, I suspect that many paradigmatic transformative choices are 'mixed'—involving both choice-based and event-based transformation. My suggestion is only that in cases in which it might appear that an event does the transforming work, the choice itself may also transform the agent. Focusing only on event-based transformative choices leads us to overlook choice-based transformation. And only the latter raises problems for standard approaches to rational choice.

I'll end the paper by pointing out two upshots that the account of choice-based transformative choices favored here have for our thinking about rational choice. First, the account of transformative choices I've offered is neutral as to the nature of options for choice. What you choose between, in a transformative choice, may not itself be transformative. It is often assumed that transformative choices must be 'big' choices, choices between options that, if pursued, will transform your life. But on the choice-based account, transformative choices can be made between options that don't themselves transform you. You can change your normative character in small, mundane ways. A beach vacation and a mountain retreat might be on a par. Neither option will transform you. But now suppose you commit to some feature of the beach vacation, thereby making it true that you have most reason to go on the beach vacation. You change 'who you are' in this small way by creating for yourself a new reason you didn't have before. You are now, to that small extent, a beach person rather than a mountain person. The same goes for choosing between desserts that are on a par. If you create for yourself a will-based reason to go for the chocolate mousse over the fruit cup, you change your normative character to the extent that before the choice, those desserts were on a par, and after your choice, the chocolate one is better. In short, you always have the opportunity to transform yourself whenever you face options that are on a par. That's because when options are on a par, we have the normative power to create reasons for ourselves that may then change us from people for whom two options are on a par to people for whom one of them is better. We can change who we are in both big and small ways.

The most important and far-reaching upshot of the account proposed here is that transformative choices, so understood, require us to reject standard approaches to rational choice. Transformative choices, then, suggest that rational agency is not simply a matter of *recognizing* and then

responding to reasons given to us by the world. Instead, part of what it is to be rational is to *create* reasons for yourself, that is, put your agency behind a consideration by committing to it. By creating reasons for yourself, you change who you are—you transform yourself—through an activity of your own will. This is transformation in the deepest sense—not something that happens to us but something we do ourselves.

Ruth Chang

E-mail : ruthechang@gmail.com

References:

- Anderson, Elizabeth. 1993. *Value in Ethics and Economics*. Cambridge, MA: Harvard University Press.
- Barnes, Elizabeth. 2015. "Social Identities and Transformative Experience." *Res Philosophica* 92 (2): 171–187. <http://dx.doi.org/10.11612/resphil.2015.92.2.3>.
- Broome, John. 1991. *Weighing Goods*. Oxford: Blackwell.
- Carthcart, Thomas and Daniel Klein. 2007. *Plato and a Platypus Walk into a Bar . . . : Understanding Philosophy Through Jokes*. New York, NY: Abrams.
- Chang, Ruth. 2002. "The Possibility of Parity." *Ethics* 112: 659–688. <http://dx.doi.org/10.1086/339673>.
- Chang, Ruth. 2009. "Voluntarist Reasons and the Sources of Normativity." In *Reasons for Action*, edited by David Sobel and Steven Wall, 243–271. New York, NY: Cambridge University Press.
- Chang, Ruth. 2012. "Are Hard Choices Cases of Incomparability?" *Philosophical Issues* 22 (1): 106–126. <http://dx.doi.org/10.1111/j.1533-6077.2012.00239.x>.
- Chang, Ruth. 2013a. "Commitments, Reasons, and the Will." *Oxford Studies in Metaethics* 8: 74–113. <http://dx.doi.org/10.1093/acprof:oso/9780199678044.003.0004>.
- Chang, Ruth. 2013b. "Grounding Practical Normativity: Going Hybrid." *Philosophical Studies* 164 (1): 163–187. <http://dx.doi.org/10.1007/s11098-013-0092-z>.
- Chang, Ruth. Forthcoming. "Comparativism: The Grounds of Rational Choice." In *Weighing Reasons*, edited by Errol Lord and Barry McGuire. Oxford: Oxford University Press.
- Elster, Jon. 1979. *Ulysses and the Sirens*. Cambridge: Cambridge University Press.
- Harman, Elizabeth. 2015. "Transformative Experience and Reliance on Moral Testimony." *Res Philosophica* 92 (2): 323–339. <http://dx.doi.org/10.11612/resphil.2015.92.2.8>.
- Herman, Barbara. 1996. *The Practical of Moral Judgment*. Cambridge, MA: Harvard University Press.
- Hill, Thomas. 2001. *Human Welfare and Moral Worth: Kantian Perspectives*. Oxford: Oxford University Press.
- Hsieh, Nien-he. 2005. "Equality, Clumpiness, and Incomparability." *Utilitas* 17: 180–204. <http://dx.doi.org/10.1017/S0953820805001512>.
- Jackson, Frank. 1982. "Epiphenomenal Qualia." *The Philosophical Quarterly* 32: 127–136. <http://dx.doi.org/10.2307/2960077>.
- Kemp, Ryan. 2015. "The Self-Transformation Puzzle: On the Possibility of Radical Self-Transformation." *Res Philosophica* 92 (2): 389–417. <http://dx.doi.org/10.11612/resphil.2015.92.2.11>.

Acknowledgements This paper was presented at the *Transformative Experience* conference at St. Louis University, September 19–20, 2014. Thanks to Hille Paakkunainen for helpful commentary and Jon Jacobs for his invitation and editorship. Thanks also to various audiences at the University of Texas and Princeton for discussion, especially Johann Frick and Alec Walen, and to two anonymous referees for their comments.

- Korsgaard, Christine. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Korsgaard, Christine. 2009. *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.
- Kurson, Robert. 2007. *Crashing Through: A True Story of Risk, Adventure, and the Man Who Dared to See*. New York, NY: Random House.
- Lewis, David. 2004. "What Experience Teaches." In *There's Something About Mary: Essays on Phenomenal Consciousness and Frank Jackson's Knowledge Argument*, edited by Peter Ludlow, Yujin Nagasawa, and Daniel Stoljar, 77–103. Cambridge, MA: MIT Press.
- Manne, Kate. Forthcoming. "Democratizing Humeanism." In *Weighing Reasons*, edited by Errol Lord and Barry McGuire. Oxford: Oxford University Press.
- Nagel, Thomas. 1974. "What Is It Like to Be a Bat?" *Philosophical Review* 83 (4): 435–450. <http://dx.doi.org/10.2307/2183914>.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Paul, L. A. 2014. *Transformative Experience*. Oxford: Oxford University Press.
- Paul, L. A. 2015. "What You Can't Expect When You're Expecting." *Res Philosophica* 92 (2): 149–170. <http://dx.doi.org/10.11612/resphil.2015.92.2.1>.
- Pettigrew, Richard. 2014. "L. A. Paul on Transformative Experience and Decision Theory I." <http://m-phi.blogspot.de/2014/08/l-paul-on-transformative-experience-and.html>.
- Ross, Jake. 2008. *Acceptance and Practical Reason*. Ph.D. thesis, Rutgers University.
- Scanlon, Thomas. 2004. "Reasons: A Puzzling Duality." In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, edited by R. Jay Wallace, Phillip Pettit, Samuel Scheffler, and Michael Smith, 231–246. Oxford: Oxford University Press.
- Sepielli, Andrew. 2009. "What to Do When You Don't Know What to Do." *Oxford Studies in Metaethics* 4: 5–28.
- Ullman-Margalit, Edna. 2007. "Difficult Choices: To Agonize or Not to Agonize?" *Social Research* 71 (1): 51–78.
- Ullmann-Margalit, Edna. 2006. "Big Decisions: Opting, Converting, Drifting." *Royal Institute of Philosophy Supplement* 58: 157–172. <http://dx.doi.org/10.1017/S1358246106058085>.