

Transformative Experience: Précis and Replies

L.A. Paul¹

University of North Carolina, Chapel Hill

[Forthcoming in PPR. Note: this is not a final draft. Please refer to the final, published version for citation.]

Précis

As we live our lives, we repeatedly make decisions that shape our future circumstances and affect the sort of person we will become. Some of these decisions are major, life-changing decisions: in such cases, we stand at a personal crossroads and must choose our direction. *Transformative Experience* raises questions about how, if we are to make these sorts of life-changing decisions about our futures rationally, we are also to make them authentically.

The philosophical framing for my arguments about making rational, authentic, life-defining choices draws on what we've learned from debates in philosophy of mind about how experience can be necessary for us to have certain imaginative capacities and cognitive abilities. It also draws on debates about color over the intrinsic value of subjective color experience, and on the importance of the first-personal perspective in understanding the self and its possibilities. It ties together the value of experience and its role in prospectively assessing our first personal futures with formal tools drawn from decision theory, causal modeling and cognitive science to address questions concerned with first personal decision making and self-construction in contexts of transformative decision-making.

I explore the tension between rationality and authenticity by considering decision-making from the first-personal perspective in contexts of radical epistemic and personal change. A natural way to make major life choices, such as whether to start a family or to pursue a particular career, is to assess our options by imaginatively projecting ourselves forward into different possible futures. But for choices involving dramatically new, life-changing experiences, we are often confronted by the brute fact that before we undergo the experience, we know very little about what these future outcomes will be like from our own first-personal perspective. Our

¹ I am indebted to Tyler Doggett and Kieran Healy for discussion.

imaginative and epistemic capacities are correspondingly limited, with serious implications for our decision-making. If we are to make life choices in a way we naturally and intuitively want to—by considering what we care about, and imagining the results of our choice for our future selves and future lived experiences—we only learn what we really need to know after we have already committed ourselves. If we try to escape the dilemma by avoiding the new experience, we have still made a choice.

1. *Epistemic and Personal Transformation*

Central to the argument is the notion of a transformative experience. As I develop it, a transformative experience is a kind of experience that is both radically new to the agent and changes her in a deep and fundamental way; experiences such as becoming a parent, discovering a new faith, emigrating to a new country, or fighting in a war. Such experiences are both epistemically and personally transformative.

An *epistemically* transformative experience is an experience that teaches you something you could not have learned without having that kind of experience. Having that experience gives you new abilities to imagine, recognize, and cognitively model possible future experiences of that kind. A *personally* transformative experience changes you in some deep and personally fundamental way, for example, by changing your core personal preferences or by changing the way you understand your desires and the kind of person you take yourself to be. A *transformative experience*, then, is an experience that is both epistemically and personally transformative. Transformative choices and transformative decisions are choices and decisions that centrally involve transformative experiences.

2. *Transformative Experience and Rational Decisions*

The main problem with transformative decisions is that our standard decision models break down when we lack epistemic access to the subjective values for the possible outcomes. Metaphorically, you can't "see" the outcomes in order to knowledgably assess them in the relevant way. As a result, in cases of transformative choice, the rationality of an approach to life where we think of ourselves as authoritatively controlling our choices by imaginatively projecting ourselves forward and considering possible subjective futures is undermined by our cognitive and epistemic limitations. If we attempt to fix the problem by adjusting our decision-theoretic models and eliminating the role for imagination and first personal assessment, the authenticity of our decision-making is undermined.

My target is the ordinary and plausible assumption that, when making big life choices, the ideal rational agent acts authentically by reflecting upon how she wants to realize her future, and perhaps realize herself as a certain kind of person, before she makes her choice. On this approach, you, as the agent, review your options and do a kind of cognitive modeling from the subjective perspective. You imaginatively project different possible futures

for yourself, futures that stem from different possible choices you could make. When you are considering your options, roughly, you evaluate each option by running a mental simulation of what the outcome would be like, should you decide to choose that option. After you run each simulation, you assign it a subjective value, and then you compare all the different values when you make your choice.²

Of course, when you decide, you also take into account any outside testimony and empirical evidence that bears on the question of what to do, but in the end, you evaluate the options by weighing the evidence and considering the expected value of each act from your own perspective. This process of simulation or imaginative acquaintance fits with how normative decision theory is supposed to provide a guide for how agents, if they are making rational decisions about their future, should proceed.

I argue that big life choices often concern transformative experiences, compromising our ability to assign subjective values to the radically new outcomes of these choices. This in turn compromises our ability to use our preferred decision models to make these choices rationally. If you can't "see" the future selves that are the possible outcomes of your choices, you can't model and assess them for their subjective utility. An important issue embedded in this argument concerns the personally transformative nature of the epistemically transformative experience: because you change dramatically, your preferences concerning the new outcomes can also change dramatically. If an experience irreversibly changes who you are, choosing to undergo it might make you care about very different things than you care for now: who you are and what you care about may change when you strike out into the unknown. As a result, having the new experience may change how your post-experience self values the outcomes, including your valuing of your higher-order values, creating a problem for how you are to adjudicate between these different sets of preferences.

A complication thus arises: If, before you make the transformative choice, the dramatic future changes in yourself are phenomenologically inaccessible to you, then from within your first personal perspective, you cannot "foresee" the ways your future self will change or foresee how your high order values will evolve. Thus, you cannot first-personally foresee or understand who you'll become. This creates a deep, existential problem for

² Subjective values, as I understand them, are experientially grounded values attaching to lived experiences. These are the types of values that are involved in transformative decision-making: I describe them as "what it's like" values to emphasize that they necessarily include phenomenal value. (There are other types of values, of course, such as moral and political values, that also come into play when we make big decisions.) Subjective values can be based on more than merely qualitative or sensory phenomenology, as they may also include values arising from nonsensory phenomenological content. They are intended to include contentful features of rich, developed experiences that embed a range of mental states like beliefs, emotions, and desires.

a model of rational choice based on maximizing one's expected utility if the goal is to model a choice where one is, in effect, choosing which future self to realize.

3. *Becoming a Vampire*

I illustrate the situation with vampires. Imagine that you have a one-time-only chance to become a vampire. With one swift, painless bite, you'll be permanently transformed into an elegant and fabulous creature of the night. As a member of the Undead, your life will be completely different. You'll experience a range of intense new sense experiences, you'll gain immortal strength, speed and power, and you'll look fantastic in everything you wear. You'll also need to drink animal blood (but not human blood) and avoid sunlight. Suppose that all of your friends, people whose interests, views and lives were similar to yours, have already decided to become vampires. And all of them tell you that they love it. They describe their new lives with unbridled enthusiasm, and encourage you to become a vampire too. They say things like: "I'd never go back, even if I could. Life has meaning and a sense of purpose now that it never had when I was human. It's amazing. But I can't really explain it to you, a mere human—you have to become a vampire yourself to know what it is like."

So, the question is, would you do it? And the trouble is, how could you possibly make an informed choice? For, after all, you cannot know what it is like to become a vampire until you become one. The experience of becoming a vampire is transformative. That is, it is an experience that is both radically new, such that you have to have it to know what it will be like for you, and moreover, it will change your core personal preferences. You can't possibly know what it would be like before you try it. And you can't possibly know what you'd be missing if you didn't. So you can't rationally choose to do it, but nor can you rationally choose to avoid it, if you want to choose based on what you think it would be like to be a vampire

The vampire case is structurally parallel to a version of Frank Jackson's case of Mary growing up in a black and white room, but where Mary is an ordinary person like you or me, or maybe an ordinary scientist, rather than someone who knows all of complete science at the end of inquiry. The parallel here concerns whether Mary knows what she needs to be able to know if she wants to decide whether to leave her black and white room, if the choice is based on what she thinks seeing color will be like. (She should leave her room if seeing color will be like *this*, but she shouldn't leave it if seeing color will be like *that*.) In this situation, she cannot perform the sort of cognitive modeling that she needs to be able to perform to assign the values of the outcomes for her, and thus to calculate the expected value of leaving her room. She lacks the ability to imaginatively acquaint herself with the future event, what it will be like for her to see color, in a way that can provide a guide for how she should proceed.

We can see, pretty readily, how the puzzle arises in fictional cases like choosing whether to become a vampire or Mary's choice to leave her room. Real life cases of transformative choice involve cases like a congenitally blind adult choosing to have a retinal operation or a person choosing to have her first child. In these cases, you also can't know what it will be like to have the characterizing experience before you have it, and if you choose to have it, it will change you significantly and irreversibly.

4. *Choosing to have a Child*

I elaborate the idea by developing the real-life example of the choice to have one's first child. Having a child often results in the transformative experience of gestating, producing, and becoming attached to your own child. At least in the ordinary case, if you are a woman who has a child, you go through a distinctive and unique experience when growing, carrying and giving birth to the child, and in the process you form a particular, distinctive and unique attachment to the actual newborn you produce. Men can go through a partly similar experience, one without the physical part of gestating and giving birth. For both parents, in the usual case, the attachment is then deepened and developed as they raise their child.

I take the experience of having a child to be unique, because physically producing a child of one's own is unlike any other human experience. As a mother, in a normal pregnancy, you grow the child inside yourself, and produce the baby as part of the birth process. As a father, you contribute your genetic material and watch the child grow inside your partner. When a newborn is produced, both parents experience dramatic hormonal changes and enter other new physiological states, all of which help to create the physical realizer for the intensely emotional phenomenology associated with the birth. These experiences contribute to the forming and strengthening of the attachment relation, and further characteristics of the nature of the attachment manifested between you and your child are determined by the particular properties of the actual child you produce. All of this generates the unique experience associated with having one's first child. Raising a child is then a temporally extended process that extends, deepens, and complicates this relationship.

This unique type of experience often transforms people in the personal sense, and in the process, changes one's preferences. If I am right that the salient details of the nature of the transformative experience of producing and becoming cognitively and emotionally attached to your first child are epistemically inaccessible to you before you undergo the experience, then you cannot, from your first personal perspective, forecast the first-personal nature of the preference changes you may undergo, at least not in a sufficiently detailed way. If so, then the choice to have a child asks you to make a decision where you must choose between earlier and later selves at different times, with different sets of preferences, but where the earlier self lacks crucial information about the preferences of the possible later selves,

and cannot foresee, in the relevant first-personal sense, the self she is making herself into.

Once we see how epistemic and personal transformation work, it becomes apparent that many of life's biggest decisions seem to involve choices to have experiences that teach us things we cannot first-personally know about from any other source but the experience itself. With many big life choices, we only learn what we need to know after we've done it, and we change in the process of doing it. The lesson I draw is that an approach to life that is both rational and authentic requires epistemic humility: life is more about discovery, and coming to terms with who we've made ourselves into via our choices, than about carefully executing a plan for self-realization.

Replies to Pettigrew, Barnes, and Campbell.

I am very grateful to my three symposiasts for their thoughtful, generous, and philosophically rich comments. The questions and challenges they raise bring out further issues and have helped me to develop the idea in several new directions.³

Reply to Richard Pettigrew

Richard Pettigrew proposes a beautifully clear model for how to make transformative decisions under epistemic and personal transformation. The key to Pettigrew's model for epistemically transformative decisions is to replace what we can't grasp epistemically with uncertainty about possible utilities. He then proposes that we model epistemically and personally transformative decisions using sums of weighted local utility functions across time (where each local function can reflect local epistemic uncertainty at that time).

I agree with Pettigrew that decision theory can be reformulated in the way he proposes, but such models will still imply profound epistemic and

³ I am particularly grateful to Elizabeth Barnes, John Campbell, Tyler Doggett, Kieran Healy, and Richard Pettigrew for several insightful discussions and generous written comments that significantly improved these replies. I also thank Herman Cappelen, Enoch Lambert, Casey O'Callaghan, and Alastair Wilson for very helpful comments.

metaphysical alienation for the decision maker, and, as he points out, will still leave us with difficult problems concerning how to weight local utility functions. What Pettigrew has shown us is how to make some versions of these questions precise, and, in cases of transformative choice, he has drawn out just how deep and fundamental the divide between deciding rationally and deciding authentically may be.

1. Replacing the unknowable with the uncertain.

The first main feature of Pettigrew's model is to substitute uncertainty about utilities for the inaccessibility of subjective utilities in cases of epistemically transformative experiences. In effect, we are replacing a utility function whose values for certain outcomes are undefined with a utility function that assigns sets of possible values for those outcomes.

The immediate problem with the model is that it seems to be modeling the wrong decision problem. Mary's not knowing what it will be like for her to see red if she leaves her black and white room is not the same as Mary's being acquainted with a wide range of color experiences but not knowing which one of these experiences will be relevantly similar to her experience of seeing red when she leaves her room. Nor is it the same as Mary's being uncertain about which member of a (potentially enormous) range of possibilities, all of which are epistemically accessible to her, will obtain. Rather, what it's like to see red is simply epistemically inaccessible to Mary until she has the requisite experience.⁴

For this reason, I am not inclined to accept Pettigrew's redescription strategy for epistemically transformative choices. It isn't a redescription of the same decision problem: it replaces that problem with a different one.

Pettigrew recognizes this, but wants to collapse the difference for the purposes of rational decision-making.⁵ He argues that "... you don't need to know the possible phenomenal characters of the experiences that you will have at each outcome in order to know what the possible utility hypotheses are ... You simply need to know the possible values of your utility function—and you do know that, since possible utility values are all real numbers."⁶

But something important has been lost. Consider InvertMary, who is functionally identical to Mary, but phenomenally inverted with respect to her color experience. If InvertMary sees a green apple, she'll have color experiences phenomenally identical to Mary's experiences of a red apple,

⁴ We might even hold that since she lacks the relevant phenomenal concept, she simply cannot represent her experience in the sense that's needed.

⁵ John Collins (2015) explores the possibility of rational neophobia (fear of the unknown) in these sorts of cases, and argues that substituting a synthetic lottery is rationally permissible under certain circumstances.

⁶ p. 3??

and vice-versa. Let's also assume that, for Mary, the numerical value of the utility of what it's like for her to experience red is the same as the numerical value of the utility of what it's like for her to experience green. Since Mary and InvertMary are functional duplicates, their numerical utility values for their experiences of seeing red and their experiences of seeing green are identical. So Mary and InvertMary can expect the same numerical utility if they decide to leave their respective black and white rooms. But it is obvious that what it is like for Mary to see a red fire engine for the first time is different from what it is like for InvertMary to see a red fire engine for the first time, and so something important about the decision-making for each of them differs, even if the numerical values assigned to their subjective utilities are the same.⁷

Suppose, despite this worries, we adopt Pettigrew's redescription strategy. Then we can make transformative decisions rationally—but, as I shall argue, we still face an important kind of philosophical loss, and we cannot escape the problems raised by transformative choice. Agents who use Pettigrew's model for transformative choice will face two distinctive types of alienation from their outcomes. I'll explore each type of alienation in the context of a thought experiment.

2. Alienation under rational choice.

Consider a situation where you desire to have a baby. Your colleague has built a computer, call it "Hal", who can calculate your utilities for you. You consult Hal, and he tells you that you can expect a utility in the range between 2 and 3 if you have a child, and between 6 and 7 if you don't. You attach equal credences to each possible state given that the relevant act is performed. Given this, you will maximize your expected utility by choosing to remain childless, even if you are uncertain about just how much.

You can't understand this result, because although you don't feel like you have a detailed grasp on what the future would be like (everyone tells you life changes dramatically), your own assessment of your utilities for having a child by imaginatively or introspectively prefiguring your future self as a parent assigns a very high utility to having a child, and a very low utility to not having one. In short, you desperately want to have a child.

Given that choosing to act in a way that does not maximize one's utility is not rational, then according to Hal's assessment of your utilities, you can't rationally choose to have a child, even though this conflicts with your assessment of yourself. In this situation, to choose rationally, you must revise your beliefs, allowing the computer's determination of what you are to believe about your utilities to replace your own introspective assessment of your heart's desires.

⁷ This example is *not* intended to suggest that subjective value is based on an internalist notion of experience. See my reply to Campbell REF.

Nevertheless, you believe in Hal, and so you accept his assessment for you, even if it does not comport with what you believe about how you would respond. As a result, you are epistemically alienated from your rational choice by your imaginative incapacities.

But what is Hal doing when he tells you what the range of your future utilities will be? Hal is, in effect, considering you in the actual world, @, at t_1 , and then assessing your utilities at t_2 in different possible worlds W_1 and W_2 . In W_1 at t_2 , you have a baby, and in W_2 at t_2 , you do not have a baby. Hal has to assess your utilities in different possible worlds because he is assessing what the actual world would be like under different possible changes of state. (Before you have a baby, as I discussed above, W_1 at t_2 is epistemically inaccessible to you, but Hal reports back about what he finds.)

Do you exist in W_1 and W_2 ? Yes—or at least your respective counterparts do. Let's call the person who exists in W_1 at t_2 , " C_1 " and the person who exists in W_2 at t_2 , " C_2 ". There is a problem with C_1 , the person who is identical to you (or who represents you) in W_1 at t_2 .

Here is the root of the problem. Normally, with a state change, the agent is kept fixed, in order to assess her utility in the new state. But the state change represented by W_1 does not exist in isolation: because the state change involves an epistemically and personally transformative experience for you, changing the state of the actual world *also* changes your preferences and your psychological capacities. If C_1 is you in W_1 at t_2 , this represents a significant change in your first-personal perspective.⁸

The trouble is that at t_1 , in @, when you consider the choice to have a baby, from your first-personal perspective, C_1 's point of view is psychologically alien to you.⁹ You cannot project your point of view into C_1 's point of view, or grasp her point of view as an extension of your own.

David Velleman's work on personal identity and persistence (REF) brings out the importance of having psychological access to one's future self: "The future 'me' whose existence matters [to me] is picked out precisely by his owning a point of view into which I am attempting to project my representations of the future, just as a past 'me' can be picked out by his having owned the point of view from which I have recovered representations of the past."¹⁰ While C_1 might be, strictly speaking, personally identical to you, from your actual perspective at t_1 , C_1 is not an

⁸ It represents a change the features of the agent whose utility is being assessed, not just the circumstances of the world in which the agent is embedded.

⁹ Or, we might say, C_1 isn't who you, from your @-at- t_1 vantage point, would identify as your psychological counterpart.

¹⁰ p. 76. "Self to Self"

eligible future self, because C₁ is not psychologically accessible to you in any first-personal sense.¹¹

So, in this sense, the utilities that Hal is discovering in W₁ are not the utilities of your future self. They are the utilities of C₁ at t₂, but from your first personal perspective at t₁, C₁ is *not you*.¹² When you consider your decision at t₁, you want to know how you'll respond to the experience at t₂, that is, whether your preferences will be satisfied. Wanting to have *your* preferences satisfied carries with it an implicit, psychological, first-personal constraint: when you make an important personal decision to act in a certain way, you want to know the (range of) utilities that the person who you can first-personally identify as your future self will have. In other words, when you assess your possible acts, you want to have psychological access, in an anticipatory or imaginative way, to each of your possible future selves. For each possible act, you want to grasp the first personal perspective of the self who you think you could make yourself into, and who will live with the result of your choice.¹³

Because, from your first-personal perspective at t₁, C₁ is not you (or, if counterpart theory is preferred, C₁ is the wrong counterpart), using Pettigrew's model entails a type of metaphysical alienation from your possible future selves. If you are facing a possible change, and you are psychologically alienated from the person who will result from this change, then the person who results is not your future self: the metaphysical relation of same-selfhood between who you are now and who you will be after the change does not obtain.

In fact, the person at t₂ that you (in @, at t₁) want to assess is a person in W₃ at t₂, a world with a state change in which you have a baby but your preferences and perspective remain the same as they were in @ at t₁. That person, call her "C₃," is (perhaps) psychologically accessible to you, but more importantly, she exists in a physically inaccessible world, and so she will not be the self that results from your having a baby.

So the sort of alienation you face doesn't arise from deferring to the dictates of morality or to your social group. We are not in the domain of traditional Existentialism. Rather, rational choice in transformative contexts entails epistemic alienation from the outcomes of one's choice and metaphysical

¹¹ On some metaphysical accounts of personal identity, the *same person* relation merely requires the right sorts of causal or other sorts of continuity. The point here is that *same person* and *same self* are different relations, and the one that matters in these decision contexts is the *same self* relation.

¹² Or, I'd be inclined to say, C₁ is the *wrong* counterpart. "It's the wrong trousers, Gromit, and they've gone wrong!"

¹³ Counterpart theoretically: you want to know the (range of) utilities of a counterpart that is psychologically similar in the relevant first-personal sense to who you are now.

alienation from one's future selves.¹⁴ If we rely on decision models that replace our inability to know or to grasp the utilities with uncertainty over utility values given to us, the epistemic and personal changes that feature in transformative choice will still entail profound epistemic alienation from our possible future outcomes, and profound metaphysical alienation from our possible future selves.¹⁵

3. *Can social science save us?*

Pettigrew suggests that we replace the utilities that are inaccessible via introspection with a range of possible utility values. He suggests that, in many cases, the values of our utilities can be determined using statistical evidence gathered by psychological and social science. If so, the fact that I must dispense with introspection is not a significant loss, for I can (and perhaps, in deference to science, I should!) replace my introspective utilities with [sets of] utilities determined by the statistical data. But this strategy will not evade alienation, for in cases of transformative change, statistical evidence cannot tell an individual what her own, individual-specific utilities are.

One important reason, which I will nevertheless set aside in what follows, is pragmatic. It is not currently possible for the psychological and social sciences to tell us, even allowing for some uncertainty, what our individual utilities are for big life decisions. We have nothing that's even close to good enough data. So, in the immediate future, there's no hope of the science even approximating the job that Hal did in the example above. Moreover, as technology and culture develops, the choices the world offers can be extremely complex and are constantly changing. Entirely new kinds of personal choices arise with major technological and scientific advances. Thus, for real-world big decisions in the immediate future, adequate statistical data is unavailable, so we must immediately face the philosophical losses of transformative decision-making in the real world.

Let us set this issue aside for the purposes of discussion, and pretend that the statistical data is in fact adequate. If we assume that it is adequate, can I use such data to discover my own utilities? No, I cannot.

¹⁴ There is a direct connection to morality, but it isn't through traditional existentialism. Rather, transformative choice suggests that what our best decision models propose conflicts with what is personally, socially, and morally acceptable. In other words, the models for such choice that fit with normative decision theory conflict with traditional kinds of normativity: the normative ideals for personal, social and moral choices. At least, they conflict to the extent that our personal, social and moral choices rely on our ability to imaginatively access our own future perspectives and the perspectives of others. See, for example, Harsanyi REF and Holton and Langton REF.

¹⁵ If we do not employ the replacement strategy, we can frame the epistemic and metaphysical alienation differently. Since your preferences cannot even be formed until you've had the transformative experience, since you cannot conceptually represent the outcomes in the relevant way, it is epistemically and metaphysically indeterminate which counterparts are yours.

The first reason I cannot use such data to discover my own utilities is because what I know from the data is merely general. What the statistical data can tell me is what the *average* effect (or utility value) would be for any member of the population (which, by assumption, we take to be composed of individuals like me). The average effect, however, is perfectly consistent with wide and dramatic variation in the values assigned to utilities (including the range of uncertainty) for any particular individual member who is included in this average. In fact, with real data, we see such variation all the time.

Intuitively, we may wish to use introspection to help us interpret average utility values and ranges of uncertainty with respect to an individual case. In particular, I might wish to consider average utility values for a member of my population with respect to how it comports with my introspective assessments of my own utilities, in order to refine my own personal utility values and my own range of uncertainties. But in the context of transformative decisions, no such introspective method is available.

It may be that discovering such average values are the best we can do, but we must be clear that *to discover the average utility values and its range of uncertainties for a member of my population is not to discover my individual utility values and my own, individual range of uncertainties*. Put another way, to choose the act that you expect to have the highest utility by the lights of the average member of your population is not the same thing as to choose the act that you expect to have the highest utility by your own lights.

The second reason I cannot use such data to discover my own utilities is because the data may *conflate* two distinct types of utilities. Assume that the empirical data tells you the average utility values for members of your population after they undergo a transformative experience. Also assume that you assign these values to the relevant outcomes before you make your own transformative choice.

The trouble is that the data *cannot distinguish between* the future utility for the individual who is actually making the choice, and the future utility for a different individual who merely *replaces* the individual who is making the choice.¹⁶ In the language of the previous section, the data cannot distinguish between the utility for a future self of yours and the utility for a replacement, alien self who merely results from your undergoing the experience. The data conflates two senses of “your” utilities, for the number it reports as the utility value in each case is the same. This is the deep problem with using statistical data to model your preferences about a transformative experience, for with a choice that is both personally and epistemically transformative, it is this distinction that is of the essence.

¹⁶ The replacement individual may well be personally identical to the original individual, but may not be the same *self* as the original individual. See footnote 7 above.

Normally, when faced with a choice that entails personal change, we introspect and reflect in order to be as sure as we can that we truly know our own preferences, but *also* so that we choose in concert with who we really are, and especially, *with whom we want to become*. When undertaking a life-changing decision, we want to knowledgably control who we are making ourselves into, that is, to knowledgably choose our future self, and imaginative introspection is an important guide for doing so. But when the choice concerns an experience that is both epistemically and personally transformative, introspection cannot guide us in this way.

The problem is, neither can the data. For the social-scientific data that provides post-choice utilities can't distinguish between (i) a utility value that represents what your utility will be at t_2 by your own lights, despite your inability to introspect to that result beforehand, and (ii), a utility value that results from replacing yourself with a different, psychologically alien self at t_2 , a value that represents what *her* utility is by *her* own lights.

4. *Forming our future selves.*

How then, in the context of the alienation entailed by a choice that is both epistemically *and* personally transformative, are we to understand the difficult philosophical questions posed by Pettigrew's elegant model of personally transformative decisions?¹⁷

One problem involves how we are to weight changes in utility across time, since, as Pettigrew astutely points out, we cannot assume our higher-order utilities will remain constant. There is no epistemically neutral first-personal perspective that the agent can occupy in order to solve the decision problem in a principled way. Another problem involves how we are to address the fact that such changes are not first-personally foreseeable, and so, if we assign future local utilities based on statistical data, we embed a kind of psychological indeterminacy and subsequent alienation into the model.

To sum up: First, in a transformative context, our future utilities, including our higher order utilities, may change, and we lack a normatively principled decision-theoretical rule for such changes. Second, we are epistemically and metaphysically alienated from these new future local utilities. We lack the ability to epistemically "see" these future utilities from the subjective perspective. Third, such alienation may force us accept a decision model in which the future utilities we assign do not discriminate between utilities for natural psychological extensions of our current selves and utilities for effective psychological replacements of our current selves.

¹⁷ Also see Briggs 2015.

This, then, is the problem. In contexts of transformative choice, how are we to make decisions within the constraints of deliberative decision theory? How are we to determine and follow the relevant diachronic rational norms? And how are we to do so while preserving a role for authentic, informed choice in high stakes cases of great personal importance?

In this way, exploring the questions of transformative experience brings new kinds of philosophical problems to our attention, ones that raise questions about authenticity, alienation, and the human condition framed in formal epistemological and phenomenological terms.

Reply to Elizabeth Barnes

In her fascinating piece, Elizabeth Barnes challenges me to explain why, for a person who has never wanted children and who strongly prefers to remain child-free, choosing *not* to have a child is not rational. In particular, she presses me to explain how, in this kind of case, choosing not to have a child is any different from other choices where it's rational to avoid having a transformative experience.

Now, Barnes understands the issue in a certain way. In particular, she understands it as involving the claim that one can rationally choose to avoid the experience based on the belief that whatever it's like, it's something you don't want *right now*. Even if having the experience would make you very glad that you'd had it, this is completely irrelevant to the rationality of the decision-making process. If, right now, you don't want to have a child, than any changes of mind you'd have afterwards don't matter.

To understand the force of her argument, we need to identify the relevant constraints and consider her examples. First, we are restricting our attention to decisions framed in terms of subjective values and preferences. The decisions of interest are decisions made by an individual and concern her immediate, first-personal future, and are framed in terms of the individual's preferences and values regarding the possible characters of her future lived experiences.

Barnes thinks that the choice to have a child is a member of a class of choices where one can rationally choose to avoid the experience based on the belief that whatever it's like, it's something you don't want right now. To identify this class, she considers two cases where, she argues, the agent can rationally choose to avoid having the transformative experience and its attendant personal changes. In each case, the rational basis for the agent's decision is his belief that whatever it's like to have that experience and to change in that way, and whatever he will think about the choice in the future, it's something he doesn't want right now.

The first example is from Orwell's *1984*: Winston wants to avoid being captured and "reprogrammed" by the Thought Police. Even if he knows that, after the mental reprogramming, he'll be very happy as a devotee of Big Brother, Winston can rationally choose to avoid reprogramming. The second example is from *Star Trek*: Captain Picard wants to avoid assimilation into the Borg, a collective hive mind. Picard rejects assimilation even though he believes that, after the assimilation, he'll be very glad to be Borg.

The examples are chosen to support the claim that, even if I believe that a transformative experience like having a child is likely to result in an outcome where I am very happy to be a parent and take my lived experience to be very valuable, there is a sense in which I can disregard this fact. For, after all, who I'd be as a parent is not who I am now. So why can't I, now, ignore the preferences of that merely possible, vastly different self?

Barnes's argument is that, for the child-free decision-maker in her example, the choice to have a child is like the choice to avoid the Thought Police or to avoid becoming Borg. According to Barnes, in each case, the rational basis for the agent's decision is his belief that whatever it's like to have that experience and change in that way, and whatever he will think about the choice in the future, it's something he doesn't want *right now*, and that's all he needs to consider. Since, for Barnes's child-free decision-maker, having such preference changes are something she doesn't want *right now*, and a rational basis for a person's decision to choose not to have a child is the belief that, whatever it's like to have that experience and change in that way, it's something she doesn't want right now, that's all she needs to consider. Case closed.

Case reopened. I agree that Winston can rationally choose to avoid reprogramming and that Picard can rationally choose to avoid becoming Borg, and that both agents are correct in disregarding the preferences and values of their future selves. But choosing to have a child is not the same sort of choice. The real-life case of choosing to have a child is disanalogous to these cases.

Barnes's fictional cases *are* analogous to real-life cases of mind control, such as being hypnotized, being drugged, or becoming a member of a cult. In these sorts of cases, we are rational in disregarding the values and alien preferences of the selves who would result from the mind-altering manipulation being proposed. In such cases, the transformative experience involves a loss of mental autonomy. Mental autonomy is something that we rational agents currently have experience of, and something to which we assign a very high subjective value—so high that the value of having it can swamp the subjective value of almost any lived experience that lacks it. But there is a further difference. We also think mental autonomy has objective moral and social value, and often, this objective value trumps mere subjective value. Our assignment of a very high subjective value to mental

autonomy reflects this objective underpinning, justifying the irrelevance of the subjective value of future lived experiences without mental autonomy, especially when one's current lived experience with mental autonomy is tolerably pleasant.

Becoming a parent is not like becoming a member of a cult, being drugged, getting hypnotized, or having one's mind assimilated into a collective mind. That is, becoming a parent does not involve a loss of mental autonomy.¹⁸ While becoming a parent does involve significant changes in one's preferences, one's self-definition, and the character of one's lived experiences, external mind control or permanent mental impairment is not part of the outcome. The identity and preference change involved in becoming a parent is deep and far-reaching, but one's mental autonomy and mental capacities are not ordinarily lost or significantly impaired by the change.

The difference stems from whether making the choice should involve cognitively modeling yourself forward into the shoes of your possible future selves, the selves who would result from the decision. Some decisions should involve this sort of self-projection and some should not, for some sorts of decisions turn on the subjective value of what it's like to be that future self, and some do not. In particular, decisions involving the loss of mental autonomy do not. When an act results in a loss of mental autonomy that degrades the status of the future self's testimony and lived experience, it can be rational to disregard what it is like to be that future self when choosing how to act.

The problems raised by transformative choice, such a choosing to have a child, concern a different type of decision. Such decisions include cases where we must consider the possibility of becoming a self that is epistemically alien to us. But in addition, there cannot be an objective value that trumps considerations based on the subjective value of what it's like to be that alien future self. Transformative choices occur in cases where, in effect, our way of framing the decision presents us with an open field of possibilities. They are decision situations where, after the relevant objective moral, social and other nonsubjective values are taken into account (whether or not they are reflected in our subjective preferences), there are still multiple particular courses of action available to the rational agent. In such situations, we ordinarily want to make the choice in question by cognitively modeling ourselves forward into the perspectives of our possible future selves, so we can choose who we are making ourselves into in an informed, authentic manner. Choosing whether to have a child is just this sort of choice.

It is a distinctively modern situation to be in: one where we can choose different ways to realize our future selves without having our path laid

¹⁸ OK, apart from 3am feedings. There isn't a lot of mental autonomy there.

down for us by the authorities. In this situation, we are permitted to make our own way through the field of possibilities, and we do so, in part, by assessing our subjective preferences for how we'd like our lives to go. Philosophers have devoted a lot of attention to decisions where objective values are the values we are concerned about. But many of our big life decisions involve "open field" cases where the choice to hand is focused on maximizing our expected subjective value, because purely objective values, empirical facts, and other constraints have already been taken into account.

Thus far, I have rejected Barnes's comparison, arguing that the cases of reprogramming by the Thought Police and assimilation by the Borg are disanalogous to the transformative choices of interest, such as the choice to have a child. Since becoming a parent is not analogous to being drugged or mentally controlled, there is no justification *of this sort* for the child-free person to dismiss what her future lived experience would be like and what her future preferences would be when she makes her decision.¹⁹ However, as I argue in *Transformative Experience*, one *can* rationally refuse to discover what it's like to become a parent. How? By reframing the argument. If you are like Barnes's child-free decision-maker, you can rationally choose to forego the discovery of what it's like to become a parent by choosing to keep your current preferences rather than to discover new ones.

But there is a deeper way to interpret Barnes's view, for her point is more subtle than the first reading I've given it. Part of what Barnes is suggesting is that Winston and Picard should disregard the values and desires of their possible, mind-altered future selves because those selves are damaged in some way. They are cognitively impaired agents, and thus their wishes, once they exist, should not be allowed to affect the rational decision-making process.

In other words, I interpret Barnes as raising a deeply interesting question: when, exactly, should we regard major cognitive changes in ourselves as destructive of our mental autonomy? When would making a radical epistemic and personal change in myself count—from my point of view before the change—as my making myself into a cognitively impaired agent? Obviously, if I were to choose to undergo a lobotomy I'd be making myself into a cognitively impaired agent. But when does radical change in myself, simply as mere radical change, amount to a loss of control over how *I* think?

When I undergo a transformative change, I change my epistemic capacities and my core personal preferences, and as a result I change the character of the way I think and the way I first-personally experience the world. My response to a life-changing transformative experience will define and infuse the character of the ways I experience and value the world and myself,

¹⁹Dear reader, please note: all this simply suggests that the choice to not have a child based on expected subjective value is no more rational than the choice to have one. We are all in the same boat.

perhaps for the rest of my life. And essentially, in contexts of transformative change, I must decide whether to undergo such a change without being able to first personally forecast or model how the change will go, and thus without being able to grasp the nature of the cognitive change from my first personal perspective. This is precisely why the combination of radical epistemic change along with radical personal change is so threatening.

In such contexts, when does choosing to undergo such a change amount to giving up one's cognitive capacities in the pejorative sense? When is it the case that, before I make a decision to become a different sort of self, I can rightly regard my future self as cognitively impaired, relative to my current self? Should I, from the perspective of my current self with her current preferences, regard *any* dramatic change of my preferences, especially transformative changes to my core personal preferences, as a kind of cognitive impairment? Where is the line between revising one's preferences in response to experience such that one autonomously *learns from* the experience, versus being *controlled by* the experience?

If any transformative change counts as cognitive impairment, then Barnes's thesis endorses an unhappy conservatism: don't ever leave your small town, don't ever get a college education, and don't ever change your current political perspective, because, by your current lights, the self that results from such experience will be cognitively impaired. Barnes, of course, is not arguing for this sort of conservatism. But where do we draw the line in dismissing the epistemically inaccessible subjective perspectives of our possible future selves? What we've found is a connection to the point made by Richard Pettigrew in his comment [REF], who gives a formal presentation of a related question—how are we to weight our local utility functions over time when framing and contemplating the possibility of transformative change, especially when those future changes are first-personally inaccessible to us?

This, of course, is just the sort of question I intended to raise when writing *Transformative Experience*. I regret to say that I do not know the answer.

Reply to John Campbell

What must imagination be in order for it to play a suitable role in authentic decision-making? In his thoughtful and insightful comment, John Campbell argues that in some cases, for such a decision to be authentic, it must involve "...imagining how the external, mind-independent environment is, as well as the mental life that is located in that environment, and affective, in that your exercise of imagination directly engages your emotions and actions, without any need for further reflection." [REF]. I agree, although my conception of imagination may not be quite the same as his.

1. *Imaginative knowledge.*

To make an authentic decision in contexts of personal change, one that reflects an informed, first-personal grasp on who you are and what you care about, you often want to know how you'll respond to the effects of your acts, including whom you'll become. Who you take yourself to be now and whom you are making yourself into is informed by your ability to imaginatively evolve your first-personal perspective into your different possible futures.

When you make a decision in this way, you use your imagination to project yourself mentally forward into the first-personal perspectives of your possible future selves. On my view, for many big, life-changing decisions, you want to authentically assess your options by assessing the subjective value of your possible future lived experiences. Ideally, the assessment involves a determination of the subjective value of each possible outcome of your decision, that is, each possible lived experience, by imaginatively grasping what it would be like for you to live in that future. That is, you want to assess what it would be like for you to first-personally occupy the self who lives that experience in that outcome, and so you *imaginatively empathize* with your possible future selves.²⁰ In this way, imaginative empathy can play a central role in authentic future-self-creation, or authentic self-realization.

Intuitively, the subjective value of a lived experience is not merely a matter of the phenomenal character of the internal characteristics of one's inner life. It's a richer value, a value that includes what it's like to live '*in this*', as Campbell puts it. That is, it encompasses the value of what it's like to live in a particular set of circumstances, where those circumstances may include the external environment.

So the character of one's inner life plays an important role in determining the subjective value of lived experience. But we need not understand this in a purely internalist sense. Often, what we care about is what the experience is subjectively like for a person, *given that that person is in the circumstances that she is in.*

Hence, subjective values need not be *merely* phenomenological or *merely* experiential. One way to put this is that, by assumption, an agent making choices about her futures assigns subjective values to outcomes concerning possible lived experiences, where the value of the lived experience can include what it would be like for her to "live" that outcome in the environment in which she is in.

²⁰ It's worth noting the parallel here with Harsanyi [REF], who argues that "the basic intellectual operation ... is imaginative empathy. We imagine ourselves to be in the shoes of another person, and ask ourselves the question, 'If I were now really in *his* position, and had *his* taste, *his* education, *his* social background, *his* cultural values, and *his* psychological makeup, then what now would now be *my* preferences between various alternatives."

Campbell is correct, then, that my approach to authentic decision-making must make room for an approach that extends past valuing experience understood as merely valuing one's purely internal, sensory phenomenal character, and thus should extend past an internalist conception of the imaginative task involved in grasping such value.

Subjective values are values of lived experience, and such experience often includes one's experience of the environment, and one's experience of her environment often includes her beliefs about her environment. So subjective values need not be internalist, at least not in the sense that Campbell urges us to reject. They do not need to be understood as merely concerned with the character of one's internal mental life, and I never intended them to be so understood. They are, instead, concerned with the character of one's lived experience, which can include her experience of her environment.²¹ Since authentic decision-making can involve knowledgeably imagining the subjective value of lived experiences, I would not want to be committed to an internalist conception of imagination that excludes this.

Campbell, however, worries that I am implicitly endorsing a purely internalist conception of imaginative understanding, because my argument for authenticity and imagination exploits thought experiments in the philosophy of mind that assume internalist conceptions of qualia. Those thought experiments demonstrate the power and importance of experience in generating our imaginative capacities. However, they do so in a context where, at least arguably, an internalist conception of qualia is assumed.

I appreciate the chance to set the record straight. When using these thought experiments to frame my arguments, I am not doing so in an internalist context. That is, I do not assume that subjective value is determined merely by the inner, purely qualitative state of the self who is transformed by the new experience, nor do I assume a restrictive, internalist view of the imaginative act required for her to assess the subjective value of lived experience. Imagining the subjective value of your future lived experience is not necessarily just a matter of imagining what your mental state will be like. Moreover, I do not assume that imagination concerns only the affectless character of one's inner life. In many cases, the subjective value of an experience, as well as the imaginative act needed to authentically grasp this value, will depend partly on the environment.

Fortunately, my arguments do not depend on an internalist, affectless approach to imagination.²² The thought examples will do their job as long

²¹ In the book, I attempt to capture this by noting that they are concerned with the character of one's *veridical* lived experience, which builds in an externalist safeguard.

²² Does my view require us to factor imagination into an "internal experience" part and an "external experience" part? As I've indicated, I don't think it does. I'm not a fan of internalism *or* of externalism about qualia: I find the distinction misguided. However, to defend this here would take me too far afield. I thank Campbell for pressing me on the point (in conversation). I think there are large and complicated issues about how to understand

as the information gained by the qualitative discovery (even a merely qualitative discovery), functions as a necessary element in the epistemic and personal transformation of the agent. It is true that imagining the subjective value of your future lived experience may not merely involve imagining what your qualia will be like as you respond to the transformation. But you must still grasp what your qualia will be like in order for you to be able to cognitively evolve yourself forward under the transformative change, and in order to grasp what it will be like for you to live in your possible future circumstances. In cases of epistemic transformation, experience, even experience understood in purely internalist terms, is still needed to teach you what your future could be like, since qualitative experience of the relevant kind is needed to give you the imaginative capacity to first-personally represent or model your possible future selves in those possible future circumstances.²³

That is, what it will be like for you to live in those circumstances is informed by and infused with the qualitative information you gain when you undergo the epistemic transformation and will have a significant effect on how your core personal preferences are formed and developed. So experience is needed in order for the agent to assess the subjective value of her possible future lived experiences, whether or not we endorse internalism.

Before Mary in her black and white room decides whether to exit and discover what it's like to see color *out there*, she cannot imagine what it would be like to see red, where this includes her inability to imagine and assess the nature of what her lived experience in a colorful world would be like. She cannot imagine what it is like to see red, nor can she imagine any of the effects of that discovery on the character of her lived experience, and so she cannot imagine what it is like to see red *out there*, nor can she imagine any of the effects of that discovery on the character of her lived experience *out there*.

If a congenitally blind adult were to decide to have a retinal operation that allowed him to see for the first time, he would discover what it was like to live in the world as a sighted adult. Before he has such an operation, he doesn't know what it would be like to be sighted, and thus he cannot imagine what it would be like for him to live in the world as a sighted adult, with all the gains and losses that entails.

And, of course, knowledge of the character of one's life experiences after the retinal operation, or of one's life experiences out in the world after growing up in a black and white room, along with one's preferences given

qualia: for example, see Mark Johnston's criticisms of what he describes as the "Wallpaper View" in his "The Function of Sensory Awareness."

²³ I make no assumptions about whether the discovery is, in fact, purely qualitative. My point is that the argument will still go through under these conservative assumptions.

those experiences, is precisely what is wanted for authentic decisions in these particular cases.

Campbell's point about authentic decision-making, then, is that in some cases, in order for your first-personal decision to be authentic, your imaginative assessment of your future must be knowledgeable. In such cases, the authenticity of the decision depends on a grasp on the nature of the world as well as a grasp on the nature of your self. As Campbell puts it: "Unless you factor in that the possibility being offered to you is one in which you gain *knowledge* of some aspects of your environment, you've missed a principle factor that you ought to be taking account of in your assessment."

The point here is that while there is an essential role for the first-personal qualitative element to authentic decision-making, authenticity can require more. Authentic decision-making can require imaginative knowledge of what my future circumstances will be, where such imaginative knowledge carries with it a direct affective, emotional engagement that allows me to cognitively and emotionally empathize with my possible future selves.

2. *The lived experience of others.*

There is another dimension to the subjective value of lived experience that Campbell emphasizes, and that I agree can play a role in our authentic decision-making. This is the subjective value of the lived experiences of other people affected by our decision. For some decisions, when determining the subjective value of an act, we desire to assess the subjective value of the lived experiences of others as well as our own.

This is especially important for decisions made for those dependent upon us, such as decisions made for our children, and decisions that are selfless to some extent, that is, decisions that are to some extent made for the sake of others. If I make a medical decision for my child, my decision is heavily influenced by my assessment of the quality of her future lived experience and the future lived experiences of the rest of the family.²⁴ If I decide to devote my life to teaching disadvantaged children, my decision is heavily influenced by my assessment of how the childrens' futures would be improved if I took the job.

Part of authentically engaging with the world around you and with your possible future selves can include imaginatively knowing how you understand yourself in relation to others. Such knowledge can be needed for you to first-personally understand how you'll respond to different

²⁴ I discuss problems with informed consent and subjective decision-making for others in chapter three and in the Afterword of *Transformative Experience*.

possible scenarios involving other people, and in this way to grasp a dimension of who you are.

I might authentically decide to devote my life to others, and decide to privilege the positive first-order subjective values of others over my own first-order subjective values, perhaps because I think my selfless act may result in an outcome with a high objective value. I can make this choice authentically based on my imaginative knowledge of what I care about, given my understanding of myself in relation to others. If the choice is made authentically, my affective engagement with bringing about the objectively valuable outcome can generate a higher-order subjective value that I'd describe as a kind of "meaningfulness".

3. Expert testimony

Campbell raises another problem for an overly internalist picture of authentic decision-making. Such a picture threatens to make it legitimate for Sally to privilege her personal desire to have a baby over the cautions of the experts, simply because the experts couldn't possibly understand what it would be like for her to have a child. If experience is necessary to understand major life changes, is there a tension between authentic decision-making and decision-making guided by scientific expertise?

There is, but authenticity had better not provide a cover for one's desire to follow the dictates of astrology or new age healing. Let me clarify what I intended to argue by examining the puzzle that Campbell poses for me.

Campbell points out that my example of Sally, who privileges her personal desire to have a baby over the cautions of the experts, might seem to be parallel to the case of Billy, who privileges his personal desire to take dangerous drugs over the cautions of the experts. In the example, Billy is clearly wrong to ignore the empirical evidence. But isn't Sally just as wrong?

Yes and no.

In the book, I use the example of Sally to emphasize the idea that unreflectively replacing one's introspective assessment of subjective value with an expert's assessments of subjective value is inauthentic. If a person unreflectively defers to experts when making a transformative choice, she is abdicating her responsibility for her actions in a way that makes her act inauthentic.²⁵

However, Campbell is correct to point out that, in a context where the empirical evidence is perfectly clear-cut and clearly specifies the values (and credences) for Sally, it could be inauthentic for Sally to privilege her

²⁵ I'm particularly indebted to Tyler Doggett here.

introspective assessments over what she knows from the scientists.²⁶ Authenticity in this sort of case requires her to recognize evidence provided by the experts as evidence she should use in her decision. This does not license unreflective replacement of one's introspective assessment with an expert's assessment, but it might license *reflective* replacement of one's introspective assessment with an expert's assessment.

However, when we restrict the case to one involving transformative experiences where our empirical evidence for the outcomes is incomplete or insufficiently fine-grained, the situation changes, because we can find ourselves unsure about how the empirical results apply. In transformative contexts with incomplete evidence, the problem with replacing our introspective deliberations with an expert's assessment of subjective values has another dimension.²⁷ On a natural reading of the example, Sally's case exemplifies this additional dimension.

If the evidence against choosing to have a child were as clear-cut as the evidence in the example against taking dangerous drugs (I will assume the evidence is clear-cut in Billy's case, for after all, the drugs are labeled as "dangerous"), then it would be bizarre for Sally to ignore that evidence. But if the situation is anything like it is in the real world, the empirical evidence in Sally's case would be nothing like as complete as we are assuming it to be in Campbell's case of Billy.

Social science gives us excellent information for decisions involving populations, such as those concerning public policy or institutional guidelines. But when the empirical results can vary for different people, that is, when the evidence allows that more than one outcome is possible for the population of which you are a member, the data may be too general to guide you perfectly, as an individual, to the value for your very own outcome.²⁸ In fictional cases, we need not worry about the complications of empirical methodology. But in real-world cases of transformative

²⁶ I say "could be", because there are in-principle problems with how Sally might be expected to interpret the evidence from the experts that I am eliding for the purposes of discussion. See footnote 26, below, and my reply to Pettigrew.

²⁷ In my book, I raise two kinds of problems involving incompleteness for our interpretation of evidence in cases of transformative experience. The first concerns the incompleteness of real-world evidence in terms of observational and external validity: the data might simply be too messy or incomplete in some practical sense, and the epistemic inaccessibility of transformative experience hampers our ability to use introspection to close the gap. For related issues, see Cartwright REF. The second concerns a distinctive, in-principle problem for the individual who is asked to interpret the utility values given to her by an expert. Before she undergoes the change, we can assume she knows the value that the individual who results from the change will assign to the outcome. But she doesn't know if the self who results from the change is *her* self, or whether that self has been transformed into an epistemically alien self, a self with very different preferences from her current ones. This matters for how she is to interpret the utility value. See my reply to Pettigrew for further discussion.

²⁸ See my reply to Pettigrew [ref]

experience, with real-world evidence, we must recognize real-world limitations.

The usual solution to this problem is for the individual to rely on introspection to attempt to close the gap between the empirical results for a population and the results she can expect as an individual. That is, she uses the empirical data in conjunction with introspection on the sort of person she is in order to assess how she is likely to respond to the experience. In such a context, for someone to make a big life decision without considering her introspective evidence *is* bizarre—and inauthentic. (In my book, I also raise the worry that in cases of transformative experience our introspective deliberations are ineffective, but this is a different point.)

So Campbell is correct to point out that truly authentic decision-making may require knowledge of the way the individual relates to the world, and thus may require knowledge of how she will respond to various events and interventions she may undergo. If the individual has perfect or near-perfect evidence of how she will respond, whether that evidence comes via testimony or via reflection, to decide authentically, she should reflectively recognize and employ that knowledge in her decision-making. However, when the evidence is incomplete, simply relying on the incomplete evidence *without* introspectively trying to close the gap seems inauthentic, precisely because, in such a case, she lacks relevant knowledge of how she will respond.