

SYMPOSIUM

Paul's Reconfiguration of Decision-Problems in the Light of Transformative Experiences*

Krister Bykvist ^{(a), (β)}

Ricevuto: 20 settembre 2019; accettato: 24 ottobre 2019

Abstract This paper focuses on cases of *epistemically transformative experiences*, as Paul calls them, cases where we have radically different experiences that teach us something we would not have learned otherwise. Paul raises the new and rather intriguing question of whether epistemic transformative experiences pose a general problem for the very possibility of rational decision-making. It is argued that there is an important grain of truth in Paul's set up and solution when it is applied to a certain narrowly defined set of cases – choices to have a new taste experience in a safe environment, where no important objective values are at stake. But the way she generalizes this approach to large-scale life choices, such as the choice to become a parent, is less convincing. Furthermore, given a proper understanding of revelatory value, there is no need to reconfigure the agent's choice situation in order to enable rational decision-making.

KEYWORDS: Transformative Experience; Rational Decision-making; Revelatory Value; Subjective Value

Riassunto *La riconfigurazione dei problemi decisionali nell'ottica di Transformative Experiences di L.A. Paul* – Questo lavoro si concentra sui casi di *esperienze epistemicamente trasformative*, come le definisce Paul, casi in cui abbiamo esperienze radicalmente differenti che ci insegnano qualcosa che non avremmo appreso diversamente. Paul solleva un problema nuovo e alquanto intrigante, ossia se le esperienze epistemicamente trasformative pongano un problema generale per l'effettiva possibilità della decisione razionale. Si sosterrà come vi sia un importante elemento di verità nella posizione e nella soluzione di Paul, se riferite a un ristretto numero di casi – la scelta di provare una nuova esperienza in un ambiente sicuro, dove non sono in gioco valori oggettivamente importanti. E, tuttavia, il modo in cui Paul generalizza questo approccio investendo un vasto ambito di scelte di vita, quali la scelta di diventare genitore, è meno convincente. Inoltre, data un'adeguata comprensione di valori rivelativi, non c'è bisogno di riconfigurare il contesto di scelta dell'agente per attivare un processo decisionale razionale.

PAROLE CHIAVE: Esperienza trasformativa; Decisione razionale; Valore rivelativo; Valore soggettivo

*This paper is part of Book Symposium on L.A. PAUL, *Transformative Experience*, in: «Rivista internazionale di Filosofia e Psicologia», vol. X, n. 3, 2019, pp. 313-376.

^(a)Department of Philosophy, Stockholm University - SE - 10691 Stockholm (Sweden)

^(β)Institute for Futures Studies, Stockholm - Box 591, SE 101 31 Stockholm (Sweden)

E-mail: krister.bykvist@philosophy.su.se (✉)



IT IS NOT EASY TO be a decision-maker. So many important aspects of our decisions are bound to be hidden from informationally impoverished agents like us. Uncertainty seems therefore inescapable. Orthodox normative decision theory offers a helping hand by providing an account of rational decision-making under uncertainty. But the uncertainty addressed by this theory is severely constrained; it only includes uncertainty about the actual state of nature. For example, the agent is supposed to know that it will either rain or not rain, but she does not know which of these two states will occur. The agent is not supposed to be uncertain about the possible states of nature (e.g., that it will either rain or not rain), the possible consequences of her actions (e.g., that she gets wet if she does not take the umbrella and it rains, that she stays dry if she takes the umbrella and it rains), the set of actions (e.g., taking the umbrella and not taking the umbrella). She is also supposed to have determinate preferences for all possible consequences (e.g., a preference for staying dry over getting soaked), and determinate credences (degrees of belief) about all possible states of nature (e.g., her credence that it will rain is 0.5).

This limited focus is questioned by a rapidly expanding research field in economics and decision theory. Instead, the aim of this research is to elucidate decision making under a more pervasive kind of uncertainty, "deep uncertainty" or "severe uncertainty", as it is often called.¹ One important kind of uncertainty regarding possible consequences that has recently been discussed is "awareness of unawareness", or "conscious unawareness", i.e., cases where the agent is aware of significant "blind spots" in her current information about possible outcomes.² In these cases, she is aware of the possibility that some action leads to an outcome that, given her currently available information, she does not or cannot fully grasp. She also knows that these outcomes can make a significant difference to the desirability of the action. She might, for example, consider the possibility

of new scientific discoveries or new technological breakthroughs that could be beneficial or harmful given her current aims, without being able to grasp those discoveries and breakthroughs. This inability to grasp can be due to lack of relevant scientific or technological knowledge or, more radically, due to lack of certain crucial concepts used to describe these discoveries or breakthroughs.

Paul's book could be seen as an important contribution to this debate about decision making under "awareness of unawareness", a debate that so far has had very little input from philosophers. More specifically, it could be seen as a philosophical contribution to the debate about decision making under "growing awareness", where the decision-maker knows that possible consequences that are currently unimaginable will be known once she has performed the action.³ In Paul's cases, the possibility that is currently unimaginable is "what it would be like" or "how it would feel" to live a certain life (or a part thereof), and it is unimaginable because the agent has not yet had the experiences that would be part of this life. Once the agent has started to lead the life, she will know how it feels to live it. Paul lists many cases of this kind, some are small scale, others large scale; some are realistic, others unrealistic: eating a durian for the first time, eating vegemite for the first time, becoming a parent, becoming a vampire, becoming a doctor, joining a war, choosing to have a new sensory ability (sight or hearing). These *transformative* choices, as Paul calls them, will not only bring about radically different experiences, they will often also radically change our personality and preferences. Paul claims that such changes pose difficult but different challenges for rational decision-making. One challenge is that choices that result in changes to personality and preferences raise the question of which self to consult when making the choice: the self that exists prior to the transformation or the self that the transformation creates. When deliberating about whether to become a parent, should you consult your current and

career-driven self, or the family-oriented person that you expect you will become after you have your child?

Another problem with this radical change in experience, applied to the choice of becoming a parent, is this. Since one cannot, Paul thinks, know what it is *like* to be a parent before one has actually experienced it, one can neither rationally decide to become a parent, nor decide to stay childfree, by “mentally simulating” the experience of being a parent. But such a simulation of what it would be like for one to experience being a parent is how one should approach this decision, according to the “predominant cultural paradigm”.⁴ According to this paradigm, when making a rational decision about whether to become a parent, one estimates the “subjective value” of experiencing the outcome of the choice to become a parent, and similarly for the experience of living a childfree life, and chooses the alternative that has the highest expectation of subjective value. But since one cannot know what it is like to be a parent before having had the experience of being one, one cannot assign this experience a subjective value. Hence, since a necessary condition for the possibility of making a rational choice in the situation in question is the ability to compare the subjective values of the two outcomes, one cannot rationally solve this decision-problem.⁵

So how should one then make this life changing decision? Paul’s proposal is that when it comes to this and other transformative choices, we should “reformulate” or “reconfigure” the decision-problem so that it is seen as a choice between having and avoiding a *revelation*.⁶ In particular, we should frame the parenthood decision-problem in terms of whether we want to *discover* what parenthood would be like for us. But her solution can be applied to other cases too. When you consider the option of becoming a vampire, you should not frame the decision as involving a choice to realize the outcome described as *what it is like* to be a vampire, but as involving a choice to realize an out-

come described as *discovering what it is like* to be a vampire.⁷ When you wonder whether to join the war, you should not ask yourself what it would be *like* to join the war; you should ask yourself whether you would want to *discover* what it is like to be a soldier. When you wonder whether to become a doctor, you should not ask yourself what it would be *like* to become a doctor; you should ask yourself whether you would want to *discover* what it is like to be a doctor. More specifically, you should ask yourself whether you think the *revelation* of what it is like to live these lives itself has subjective value (regardless of whether the revealed experiences themselves have any subjective value). By invoking this subjective revelatory value, you can now compare the outcome of a transformative experience to the status quo; and you can compare them without invoking any objective values. So, we have not strayed beyond, what Paul calls, “our cultural paradigm”.⁸

In this short commentary, I shall focus on cases of *epistemically transformative experiences*, as Paul calls them, cases where we have radically different experiences that teach us something we would not have learned otherwise – what it would be like to lead a certain life. The reasons why I put aside cases of personally transformative experiences, where our personality and preferences change, is, first, that Paul says comparatively little about this problem and its solution and, second, that there already is a quite extensive literature on this problem, which Paul does not address.⁹ In contrast, Paul raises the new and rather intriguing question of whether epistemic transformative experiences pose a general problem for the possibility of rational decision-making. I shall argue that there is an important grain of truth in Paul’s set up and solution when it is applied to a certain narrowly defined set of cases – choices to have a new taste experience in a safe environment, where no important objective values are at stake. But the way she generalizes this approach to large-scale life choices, such as the

choice to become a parent, is less convincing. I should stress, however, that the book has many virtues. One is that, in contrast with the mainstream discussion about growing awareness, the discussion in Paul's book pays close attention to vivid examples of growing awareness and the philosophical issues they generate. Another is the many interesting applications and extensions of her account. For example, she gives a very thoughtful discussion of the controversial question of whether deaf parents who have had deaf child should opt to give their child an implant that will restore the child's hearing.¹⁰

When we ponder whether to try durian for the first time, in a safe environment, knowing that eating it does not put any prudential or moral values (or any other objective values) at risk, it seems sensible to be interested only in how the fruit will taste. Since we cannot know in advance exactly how the fruit will taste – other people's reports will provide limited or even conflicting information – it is sensible to frame the choice as one between *coming to know* how durian tastes and *not coming to know this*. If you are curious, you will have a go and then add this experience to your collection of gustatory experiences. This seems rational, since you prefer knowing how it tastes to not knowing it and nothing else is at stake. The focus on the revelatory value of outcomes seems therefore to be justified in this kind of case. (Whether this value is "subjective" in Paul's sense of the term is a different question, which we will come back to later).

But large-scale life choices are not like gustatory choices in safe environments. First of all, as Paul herself stresses, large-scale choices will often be personally transformative since they bring about a change in your personality and fundamental preferences. But, second, putting this thorny issue aside, choosing to become a parent, join a war, become a doctor, or become a vampire, is not like choosing to taste a durian, for these bigger life-choices involve many important prudential and moral values and they are set in

environments that are far from safe. A lot is at stake, both in terms of prudential and moral values, in joining a war or becoming a vampire. The war can be unjust and sucking blood out of humans is morally problematic to say the least. But even the choice to become a parent can have drastic effects on one's career possibilities. Important moral values can also be at stake. It is hardly morally indifferent to create a child who will live in a happy and loving relationship with you and others. Note also that creating a new life enables a whole new branch of the family tree to be created; again, hardly something that is morally indifferent. Indeed, some have recently argued that having an extra child is the choice that might have the greatest carbon footprint of all the individual choices we make in our life since this child might have a child, who might have a child, and so on for generations.

Paul does, in fact, concede that large-scale life choices often involve important objective values, but she thinks that subjective values are still some of the most «central and important ones and an emphasis on them fits the dominant cultural paradigm». This paradigm says that we should approach many major life decisions «as personal matters where a central feature of what is at stake is what it will be like for us to experience the outcomes of our acts, and where the subjective value we assign to an outcome depends upon what we care about, whatever that might be».¹¹ As I understand her, the cases where we should rely on subjective values are those in which objective values can be put aside (perhaps because the outcomes do not differ much in overall objective value).

This means that her discussion risks having quite a limited applicability since it is not clear that there are many cases left once we have excluded objective values. There do not seem to be many cases of life-changing decisions in which we can sensibly say that we will chose to lead a certain life because we wanted to see what it would be like. "I chose to become a parent because I wanted know

how it felt to be a parent”, “I chose to join the war because I wanted to know how it felt to fight in a war”, and “I chose to become a vampire because I wanted to know how it felt to live as a vampire” are all statements that would sound frivolous and self-absorbed in most cases. In contrast, to say that we chose to try durian because we wanted to know how it tastes seems perfectly acceptable.

The paradigm Paul alludes to sounds familiar if one thinks about how typical agents care about their own future pleasure or pain. We want to know what a future experience feels like so we can decide whether it will be pleasant or painful. But Paul has in mind something that goes beyond pleasure and pain. Subjective values, according to her, are not merely values of pleasure and pain. Instead, «they can be grounded by more than merely qualitative or sensory characteristics, as they may also arise from nonsensory phenomenological features of experiences, especially rich, developed experiences that embed a range of mental states, including beliefs, emotions, and desires». ¹² One problem with this characterization of subjective values is that they are both supposed to “depend on” what we care about and be “grounded by” sensory and non-sensory phenomenological features. But it is not clear how we are supposed to understand this double nature of subjective values.

One option is that it is the agent’s responses that call the shots: an outcome (or a part thereof) has subjective value for an agent if and only if the agent cares about it in virtue of some of its sensory or non-sensory phenomenal features. The agent is not required to care about any particular phenomenal features. On this view, subjective value depends on what we care about, since if we do not care about an outcome it lacks value for us. Subjective value is also grounded in phenomenal features in the sense that if we care about an outcome *in virtue of* such features, it has subjective value. The more the agent cares about it, the more subjective value it has, which can be, but need not be, a

matter of the agent caring more when the intensity of some phenomenal feature is greater. Since the phenomenal feature need not be sensory, the agent may care about the non-sensory phenomenal features of finally *discovering* how intense the taste of durian may be, assuming that these discoveries themselves have phenomenal features. So, discoveries can have subjective value for the agent.

This response-dependent notion of subjective values does not square well with other parts of Paul’s account, however. First, as explained above, Paul’s own solution to the problem of transformative choice is to assign subjective value to revelations, which, on the response-dependent model, means that the agent is simply assumed to care about the non-sensory phenomenal features of the *discovery* of having a certain experience, not taking into account the phenomenological features of the experience itself, which are supposed to be unknown to her. But many normal agents do not care about such things (I myself being one). We must distinguish the phenomenological features of the objects of knowledge – the taste experience of eating a durian – from the phenomenological features of the knowledge state itself (without its object). What is common is that people want to know how a certain culinary item tastes (a durian, for instance); but they do not care about how the *knowledge* of this experience itself feels.

Of course, Paul could claim that even though we often do not in fact care about the phenomenological features of these epistemic states we are nevertheless *rationally required* to do so. But this is hardly part of our cultural paradigm of rational decision-making. After all, we do not teach our children to care about the phenomenological features of discoveries as such. What we might do is to teach them to be more daring and explore various new taste experiences in safe environments. Finally, there might not be anything here for us to care about, since it is questionable whether these epistemic states

themselves (without taking into account their objects) have any distinct phenomenological features at all.

Second, Paul's claim that in order to know the subjective value of a future experience you *must* have already experienced it seems false, if the agent's responses call the shots. In order to know its subjective value, you only need to know that you will respond to this experience on the basis of some of its phenomenal features. You need not now know what these features are like. For example, in order to know that the experience of being a vampire will have subjective value for you, you do not need to know now what it is like to be a vampire; you only need to know that you will like the experience of being a vampire in virtue of some of its phenomenal features.

It is true, as Paul repeatedly points out, that you may not know how or how strongly you will respond to your future experiences. You may not know whether you will like or dislike your vampire experiences, for example. But this problem seems not to be of the same severity as the original one, and definitely of a different kind. Rational choice as Paul herself defines it, is about maximizing *expected* subjective value on the basis of evidentially supported credences. In order to do this, we need to have an evidentially supported credence distribution over alternative hypotheses about whether and how much your future self will like or dislike its experiences. The fact that the agent does not know which of these hypotheses will be true is not relevant for the possibility of rational choice.

The other interpretation of subjective value denies that the agent's responses call the shots; phenomenological features of outcomes can in their own right contribute to subjective value. At times, Paul says things that suggest this response-independent interpretation. She claims, for instance, that the subjective value of the outcome of eating a durian is partially a matter of the phenomenal intensity of what it is like for you to taste a durian, and «so the magnitude of the posi-

tive or negative value is not just determined by the fact that the durian tastes good (or bad) to you, but by how intense your taste experience is».¹³ Similarly, when discussing the option of becoming a vampire, Paul claims that it is possible that the *phenomenal intensity* of the experience of becoming a vampire will swamp the phenomenal intensity of the experiences in other outcomes.¹⁴

This is a very controversial view of value, to say the least, and is hardly part of our cultural paradigm about how to assess possible lives. It is not even part of the cultural paradigm about how to assess wines. Not even wine connoisseurs, who think that there are wines that we *ought* to like even if we do not find them at all pleasing, think that this value is improved whenever some aspect of the tasting experience becomes more intense. The best wine is not the one with the most intense acidity, sweetness, and tannins. What matters most is the balance of the wine – how the various taste elements come together.

Paul could reply by saying that the wine example I gave is close to what she had in mind, except that we should consider non-sensory phenomenal features as well, and that it was therefore a mistake to focus on the sheer intensity of experiences. The idea would then be that an outcome (or a part thereof) has subjective value for an agent if and only if the agent cares about it in virtue of a certain combination (possibly very complex) of sensory and non-sensory phenomenal features. These features go beyond pain and pleasure, likes and dislikes, and it is not just a matter of the intensity of the features. It is not up to the agent to decide which features those are and what the right combination is.

Again, it is very doubtful that this is part of our cultural paradigm about how to assess outcomes, however. It is true, as Paul points out, that we often use “mental simulation” in deliberations, whereby we project ourselves into our possible future outcomes and try to assess “what they would be like”. But, typical-

ly, we do this in order to predict how we would act, react, and how much pleasure and pain we would feel. It would be reading far too much into this method to understand it as a way of identifying some complex combination of phenomenal properties, which are supposed to go beyond pain and pleasure and our likes and dislikes, that in part makes the outcome valuable. At least, this use of the method is not typical. It would only be useful for *texture fetishists*, who think that the experiential texture itself, independently of whether it is painful or pleasant, liked or disliked, contributes to the value of outcomes. They deserve the name “fetishist” since even if such a person were to consult a crystal ball that informed him exactly how happy and wealthy he will be if he becomes a parent, how it will affect his career, how pleasant or unpleasant it will be, how happy the child will be, and so on, he would not be able to assess the outcome of being a parent, since he still lacks information about the *exact* texture of becoming a parent. I think it is safe to say that such persons are rare, and also that there is not a general rational requirement to become one.¹⁵

Even if I am wrong about this, this interpretation of subjective values seems to make Paul’s own solution problematic. Note that even on this interpretation, subjective value depends in part on what we care about, but that means that bringing in the subjective value of revelation is no solution if the agent happens *not* to care about how revelations feel. So, we are back to the problem of having to insist that we are rationally required to care about how revelations feel. Furthermore, since subjective value is in part defined by some combination of phenomenal properties and it is not up to the agent to decide what this combination will be, we need to know the phenomenal features that in part make revelations valuable. But what are they and how can we know them?

Finally, if we go for this account of subjective value, there does not seem to be any need to bring in the subjective value of reve-

lations. After all, when we are purchasing wine so that we can enjoy a good wine but don’t have the ability to taste it in advance, it seems very sensible to rely on the advice of wine connoisseurs whose judgements we trust and who know what kinds of wine we tend to enjoy. They might of course disagree among themselves, but we can still consider various hypotheses regarding who is likely to be right, and give more credence to the ones we trust more. The same approach could be used for life choices, assuming with Paul that they are similar to wine choices.

The main upshots of this are that (a) on both interpretations, Paul has to assume that we are rationally required to care about the phenomenal feel of revelations of experiences, which seems very implausible and definitely not part of our cultural paradigm about rational decision-making, and that (b) on neither interpretation of subjective value will the agent necessarily be stuck in her decision-making for she can entertain different hypotheses about subjective value without having to have experienced those values herself.

Paul would resist (b), partly because she thinks that the agent will not be able to assign epistemically reasonable credences to these hypotheses about subjective value. Since the agent cannot rely on first-hand knowledge of the experience, she has to rely on *third-personal* information. Paul points out, however, that relying on third-personal information is problematic because there might not be any robust information to go by.¹⁶ Paul is surely right to remind us about this problem, but some pieces of information can be better than none. Take the example of parenthood. While she might be right in claiming that one cannot know for sure how parenthood will affect one’s desires and beliefs,¹⁷ one can at least form informed expectations by talking to those who have already gone through the transformation and by reading the relevant empirical literature.¹⁸ The same can be said about the impact of parenthood on career prospects and happiness (or subjective well-being). Paul is, of

course, correct in pointing out that these statistics won't tell anyone for sure how happy they will be as a parent (or what that happiness or misery will "feel like"), or how it will affect their career. But by complimenting data about the average effects on people sufficiently similar to oneself with information gained by asking people one trusts how parenthood affected them – and, perhaps more importantly, *why* it affected them the way it did – one can certainly form informed and reasonable expectations about the effects parenthood will have on one's career and subjective well-being. These expectations can then help one decide whether one's desires are best served by becoming a parent or by remaining childfree. In any case, much more needs to be said to show that this problem makes it *impossible* to have epistemically reasonable credences about your future subjective values.

Paul has another argument against relying on third-personal information, namely, that it would threaten our *autonomy* or *authenticity* as agents. She gives the example of Sally, who has always believed that having a child will make her happy and fulfilled, but decides not to have a child just because empirical evidence suggests that remaining childless would maximize her expected subjective well-being.¹⁹ Paul claims that her choice would be "bizarre" and that she would be «giving up autonomy for the sake of rationality». ²⁰ It is not clear why Paul thinks this. If Sally wants to have a child because she believes that it would make her happy and maximize her expected subjective well-being and is informed that having a child would not maximize her expected subjective well-being, why would it be bizarre and a sacrifice of autonomy for her to choose to remain childless? She is just making use of relevant empirical information to make a choice on the basis of her preference for happiness and future subjective values. As Paul herself points out, "rational authenticity" is about making the best decisions one can in order to fulfil one's dreams and aspirations²¹ and that is exactly

what Sally is doing by making use of relevant empirical information. Paul must have in mind a very different Sally, one that has no prior preference about whether to have a child but decides to remain childless just because some social scientist told her to, where the social scientist thinks she should not have a child because it would not maximize her expected subjective well-being. That would be a rather bizarre behaviour and also a sacrifice of autonomy, but this would of course go way beyond making use of third-personal empirical evidence – it would be to give someone else control over your future.

Another problem for response-involving subjective values, whether the responses call the shots or not, is that it seems that we can have conflicts between responses. As Paul points out, my response now towards a future experience, which I already know first-hand (or the response I would now have if I knew the experience first-hand), need not agree with my future response towards the same future experience.²² I am not sure this is a major problem, however, if we put aside cases where our fundamental preferences change, for this means that we are not talking about conflicts of *ideals* here, only conflicts between responses towards the "feel and flavour" of a certain life in cases where we can safely put aside important objective values such as moral and prudential values. But normal agents do not seem to have *unconditional* attitudes towards the feel and flavour of their lives when no important values are at stake. For example, now being a fermented cabbage lover, I favour my future experience of eating fermented cabbage *on the condition* that I will later still favour it. To use Parfit's term, these favourings are *conditional on their own persistence*. In contrast, when my favouring expresses an ideal, I favour being honest, healthy, and prudent in the future even if I will later lack any concern for these things. Something like this distinction is implicitly assumed in the old joke "It is a good thing I do not like fermented cabbage because if I did I would eat something I hate".

So far, I have questioned Paul's use of subjective values in rational decision-making. But let us now assume that Paul is right about the following things: (a) rational choice, in many life choice cases, has to do with maximizing expected subjective value, (b) we need to know first-hand the phenomenal features of an experience in order to know its subjective value, (c) revelations have their own phenomenology. Even if these controversial claims are granted, it is still not clear that her own solution works. Remember that the idea is to invoke the subjective value of revelations. But in order to do this the agent needs to already know this subjective value. Otherwise the reframing of the decision situation in terms of revelations would be of no use. But if I have not yet had a certain experience, how can I know how it feels to have the *revelation* of that experience? If I do not know how it feels to have the revelation, I do not know its subjective value and we are back where we started.

There are two main responses Paul could make. She could say that the phenomenology of revelation is to a sufficiently large extent *invariant* with the object of the revelation. However, it is difficult to see why this must be true. For example, at least to me who has tried both, coming to know how vegemite tastes seems phenomenologically quite different from coming to know what it is like to be a parent. Indeed, it is difficult for Paul to deny this possible variance, since she adopts an expansive account of phenomenology, according to which the phenomenology of an experience may depend on its content. So two epistemic experiences with very different contents – i.e., coming to know very different experiences – cannot be assumed to have the same phenomenology.²³ But if this cannot be assumed, nothing is gained in terms of helping the agent decide, for instance, whether to have a child or not, by reformulating the choice as one between having or foregoing the experience of the revelation in question. For just as he cannot determine the subjective value of being a parent, so he cannot determine the

subjective value of experiencing the revelation that parenthood brings with it.²⁴

The other response is to say that the phenomenology is *different* but argue that there are crucial phenomenological similarities between the revelations of very different experiences. Even if the subjective value of revelations is somehow grounded in their intrinsic phenomenal features, they need not be grounded in all such features. So, in order to know the relevant aspects of the experience of a certain revelation, you need not have had this experience in all its phenomenal detail and richness. You only need to know what this experience is like in certain respects, which you might already have encountered in the revelatory experiences you have already had. With this knowledge, you can extrapolate from your past experience to the new future experience. But exactly the same move could be made for the experiences that are the *objects* of the revelation. Not all phenomenal features of becoming a parent are relevant for its subjective value; only some are. Perhaps in order to (at least roughly) assess the subjective value of parenthood it is enough to have had experiences of sleep deprivation, strong attachment, love, and devotion towards a vulnerable person. Indeed, Paul herself suggests this kind of move in her Afterword when she talks about the choice to become a vampire. One can then wonder why we need to bring in the subjective value of revelations if it is conceded that we can roughly assess the subjective value of the experiences that would be revealed by our transformative choice. We need an argument that shows that it is easier to identify the subjective value of revelations than to do same for the experiences that are the objects of these revelations.

Of course, these problems would all disappear if we simply denied that the *experience* of the revelation must have value. Instead we could say that whether a certain revelation has value depends on whether the agent wants to *know* what it feels like to live a certain life – i.e., to decide whether or not

she wants to be in a particular *epistemic state*. This way of understanding revelatory value seems to mesh much better with how we speak. When we are curious about how a durian tastes, we say that we want to *know* how it tastes, but we do not say that we want to know this because we care about how it *feels* to know how it tastes. There are many different reasons for why we want to know how it tastes. Many are instrumental: we want to know how it tastes because we think it is likely we will like it, or we want to know this so we can tell others to try it if it tastes good (and we assume others have the same tastes), or we want to know how it tastes, even if it tastes awful, because we want to be seen as gustatory daredevils. But we could also just be interested in collecting pieces of knowledge of taste experiences, independently of whether it tastes good or bad. This account of revelatory value would be a radical break from Paul's framework, however.

First, this option goes against the "cultural paradigm" Paul adopts, namely that we should make decisions by performing mental simulations about what the future would be like experientially in all details, going beyond future pleasures and pains, and future likes and dislikes. To ask the agent to decide whether they prefer to be in a certain epistemic state or not is not to ask her to imagine how things would be like in all their experiential details.

Second, the agent Paul had in mind cares only about subjective values and thus gets stuck in cases of transformative choices. But the agent we now imagine cares (intrinsically or instrumentally) about *knowing* what something is like, even though she denies that this revelation itself has any subjective value. She might even deny that the object of the revelation can be subjectively good: "I know the taste will be pretty awful, but I want to show others I can take it".

Third, since, unlike the texture fetishist, the agent we imagine is not stuck with values that can only be known by first-hand experience, she does not face the problem Paul

starts off with: the problem of not knowing the feel and flavor of the alternative outcomes and therefore not knowing how to assess them. She faces instead the more familiar problem of trading the revelatory value of new experiences ("Do I want to *know* how it tastes?", "Do I want to *know* how it feels to live like that?") against objective values, such as morality and prudence ("Is it *OK* to eat it?", "Is it *OK* to live like that?"). To do this sensibly she needs to first figure out whether she cares about knowing a new experience for its own sake or for the sake of other things she cares about. If it is the latter, she needs to assess the likelihoods that the revelation will bring about the things she cares about for their own sake. Again, this is a familiar problem for all kinds of values the agent considers. The important point is that, for this agent, who I think is not unlike many of us, there is no need to reconfigure the choice situation. The revelation of new experiences is already one of the intrinsic or instrumental values she considers.

Acknowledgements

I would like to thank Tim Campbell, Julia Mosquera, and Orri Stefansson for very incisive and useful comments on an earlier draft of this paper.

Notes

¹ For a survey of this literature, see C. HELGESON, *Structuring Decisions Under Deep Uncertainty*, in: «Topoi», online first 2018, August 14th – doi: 0.1007/s11245-018-9584-y.

² For a recent contribution, see, for instance, E. KARMO, M.-L. VIERO, *Awareness of Unawareness: A Theory of Decision Making in the Face of Ignorance*, in: «Journal of Economic Theory», vol. CLXVIII, 2017, pp. 301-328.

³ For a recent contribution to this debate, see M.-L. VIERO, *An Intertemporal Model of Growing Awareness*, Working Paper 1388, 2017, Economics Department, Queen's University – online available at URL: https://www.econ.queensu.ca/sites/econ.queensu.ca/files/wpaper/qed_wp_1388.pdf.

⁴ L.A. PAUL, *Transformative Experience*, Oxford

University Press, Oxford 2014, p. 25.

⁵ *Ivi*, p. 31.

⁶ *Ivi*, pp. 113-115.

⁷ *Ivi*, p. 114.

⁸ *Ivi*, p. 38 and 113.

⁹ See e.g. P. BRICKER, *Prudence*, in: «The Journal of Philosophy», vol. LXXVII, n. 7, 1980, pp. 381-401; D. PARFIT, *Reasons and Persons*, Oxford University Press, Oxford 1986; A. GIBBARD, *Interpersonal Comparisons: Preference, Good and the Intrinsic Reward of Life*, in: J. ELSTER, A. HYLLAND (eds.), *Foundations of Social Choice Theory*, Cambridge University Press, Cambridge 1986, pp. 165-193; J. ELSTER, J. ROEMER (eds.), *Interpersonal Comparisons of Well-Being*, Cambridge University Press, Cambridge 1991; K. BYKVIST, *Prudence for Changing Selves*, in: «Utilitas», vol. XVIII, n. 3, 2006, pp. 264-283; D. MCKERLIE, *Justice Between the Young and the Old*, Oxford University Press, Oxford 2013.

¹⁰ L.A. PAUL, *Transformative Experience*, cit., p. 7.

¹¹ *Ivi*, p. 25.

¹² *Ivi*, p. 12.

¹³ *Ivi*, p. 35.

¹⁴ *Ivi*, p. 43.

¹⁵ For more on texture fetishists, see K. BYKVIST, O. STEFANSSON, *Epistemic Transformation and Rational Choice*, in: «Economics and Philosophy», vol. XXXIII, n. 1, 2017, pp. 125-138.

¹⁶ L.A. PAUL, *Transformative Experience*, cit., pp. 88-89.

¹⁷ *Ivi*, p. 81 and 91.

¹⁸ Paul may also be correct in pointing out that one cannot know, before becoming a parent, *what it will be like* to have the beliefs and desires that one will come to have after becoming a parent. But the point is that those who are interested in what new attitudes one will come to have, rather than what it will be like to have those attitudes, have access to vast evidence that can inform their decision.

¹⁹ L.A. PAUL, *Transformative Experience*, cit., p. 87.

²⁰ *Ivi*, p. 88.

²¹ *Ivi*, p. 105.

²² *Ivi*, p. 16 and 81.

²³ L.A. PAUL, *Transformative Experience*, cit., p. 12, footnote 16.

²⁴ In fact Paul seems in places to admit as much, see L.A. PAUL, *Transformative Experience*, cit., pp. 93-94.