



**TRANSFORMATIVE  
EXPERIENCE**  
L. A. PAUL

OXFORD

## CHAPTER 1

# BECOMING A VAMPIRE

Imagine that you have the 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, revelatory 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 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 assuage your fears and explain that modern vampires don't kill humans; they drink the blood of cows and chickens. 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. I understand Reality in a way I just couldn't before. It's amazing. But I can't really explain it to you, a mere human—you have to be a vampire to know what it's like."<sup>1</sup> Suppose that you also know that if you pass up this opportunity up, you'll never have another chance.

### Would You Do It?

To make a choice like this, you'd want to make the best decision you could. And that means you'd want to proceed as rationally as possible.

<sup>1</sup> I thank Jordan S. Ellenberg for suggesting I use vampires to illustrate my thesis. See Ellenberg (2013).

Becoming a vampire would be big—really big. It obviously isn't a decision to be undertaken lightly. You'd want to choose the smartest option, the option that would make your life as good as it could be after you'd made your choice. You wouldn't want to pass up one of the most amazing experiences you could ever have. But you wouldn't want to make a huge mistake either.

The trouble is, in this situation, *how could you possibly make an informed choice?* For, after all, you cannot know what it is like to be a vampire until you are one. And if you can't know what it's like to be a vampire without becoming one, you can't compare the character of the lived experience of what it is like to be you, right now, a mere human, to the character of the lived experience of what it would be like to be a vampire. This means that, if you want to make this choice by considering what you want your lived experience to be like in the future, you can't do it rationally. At least, you can't do it by weighing the competing options concerning what it would be like and choosing on this basis. And it seems awfully suspect to rely solely on the testimony of your vampire friends to make your choice, because, after all, they aren't human any more, so their preferences are the ones vampires have, not the ones humans have.

This book argues that we find ourselves in this sort of situation for some of our most significant life decisions. In other words, in the real world, as we face our own personal series of life choices, some of these choices are, in a very important sense, like the choice to become a vampire. While life choices don't usually involve the possibility of becoming an immortal being, they are fundamentally similar in a different way.

The idea is that, when you find yourself facing a decision involving a new experience that is unlike any other experience you've had before, you can find yourself in a special sort of epistemic situation. In this sort of situation, you know very little about your possible future, in the same way that you are limited when you face a possible future as a vampire. And so, if you want to make the decision by thinking about what your lived experience would be like if you decided to undergo the experience, you have a problem.

In such a situation, you find yourself facing a decision where you lack the information you need to make the decision the way you naturally want to make it—by assessing what the different possibilities would be like and choosing between them. The problem is pressing, because many of life's big personal decisions are like this: they involve the choice to undergo a dramatically new experience that will change your life in important ways, and an essential part of your deliberation concerns what your future life will be like if you decide to undergo the change. But as it turns out, like the choice to become a vampire, many of these big decisions involve choices to have experiences that teach us things we cannot know about from any other source but the experience itself.

When we face a choice like this, we can't know what our lives will be like until we've undergone the new experience, but if we don't undergo the experience, we won't know what we are missing. And, further, many of these new and unknown experiences are life-changing or dramatically personally transformative. So not only must you make the choice without knowing what it will be like if you choose to have the new experience, but the choice is big, and you know it is big. You know that undergoing the experience will change what it is like for you to live your life, and perhaps even change what it is like to *be* you, deeply and fundamentally.

I will not argue that you can't get information from the testimony of others when you make such choices. You can. But I will argue that such guidance only goes so far, for the information such sources can supply is incomplete. So while we should consider the information we can gather from science and from the advice of friends and relatives, in the end, we must decide for ourselves, weighing incomplete evidence in the light of our own personal preferences.

I will not argue that you can't change the way you make the decision. You can. You could change the way you decide, such that the new deliberation does not rely on your expectations about what it would be like for *you* to have the experience, or does not involve weighing advice and testimony from your own personal perspective and deciding how

much of it applies to you personally, but rather, relies solely on impersonal facts about how people, in general, respond to these kinds of experiences. In other words, you can replace your personal approach to decision-making with impersonal decision-making, removing any crucial role for your experience or your individual, personal perspective when you deliberate.

But changing the decision this way gives an unsatisfying answer to the question of how you should make these deeply personal, centrally important, life-changing decisions. For after all, your decision concerns *your* personal future, and so an essential part of your decision is based on what it would be like for *you* to have the experience and to live the life you bring about for yourself. You naturally and intuitively want to make your life choices by thinking about what you care about and what your future experience will be like if you decide to undergo the experience. This is why you are expected to weigh evidence from your own personal perspective and decide how you want to apply it to your own situation. Making the decision into one where what it would be like for you is no longer a consideration, is not a decision you care about in the same way.

So, in many ways, large and small, as we live our lives, we find ourselves confronted with a brute fact about how little we can know about our futures, just when it is most important to us that we do know. For many big life choices, we only learn what we need to know after we've done it, and we change ourselves in the process of doing it. I'll argue that, in the end, the best response to this situation is to choose based on whether we want to discover who we'll become.<sup>2</sup>

<sup>2</sup> In the Afterword, I'll explore connections to empirical and theoretical issues in economics, statistics, and psychology, and will consider ways to develop new models for rational decision-making in contexts of transformative choice.

# L.A. Paul: "The Transformative Experience" | Edge.org

The worries or puzzles that I've been thinking about and exploring come from drawing together a number of strands in philosophy that haven't been drawn together before. The first strand involves a relatively new area of inquiry in philosophy that goes by the name of formal epistemology. It's an interesting and engaging new development, and formal epistemologists are interested in the way that individuals make decisions. They're interested in looking at formal decision theory, but they're interested in doing this within the context of epistemic questions. The thought is to explore how we can make rational decisions by taking agents to have psychologically real utilities or desires, by thinking in terms of particular degrees of beliefs or credences and also psychologically real preferences, and thinking about how we want, in an epistemic context, to think about individuals making decisions so that these individuals can know how they should act. The "should" is important; we're exploring these questions from a normative perspective.

I'm interested in normative decision theory, as opposed to behavioral decision theory. I'm interested in what the epistemic gold standard is that we as individuals should be aspiring to reach when we make decisions. In particular, what I'll talk about a little bit more concerns the normative gold standard for when we make important decisions.

The formal epistemologist usually thinks of the individual in a third-personal sense. Namely, it's as though we're observing individuals, and thinking about their epistemic states and how they're making their decisions. But there's another perspective that's also important and draws in another strand from philosophy, a strand of work that's been important over the last 30 or 40 years in philosophy. People like Dave Chalmers have made important

contributions to philosophy involving the notion of consciousness and trying to understand what consciousness is.

What I want to look at closely is what philosophers have learned about the value of experience—how we've learned about what experience teaches us. A lot of times this discussion occurs in the context of worrying about the mind-body problem, or questions about physicalism. That's not my focus. I want to get a better understanding and think about how important it is, in some contexts, that we have certain experiences in order to know or understand certain information. There are disputes in the philosophical community about whether experience is required to know certain facts, or what exactly it is that experience teaches. I'm not worried about that dispute. I just want us to be able to see that sometimes experience is important. It's necessary, at the very least, for us to have conceptual or imaginative abilities in order to grasp certain kinds of imaginative content.

If we draw the strands together—formal epistemology, normative decision theory, consciousness with a focus on what experience teaches—then we get a different perspective on how we need to think about how we make big decisions. What I'm going to say is going to connect a little bit to what Molly talked about earlier today and some of the things that Josh Knobe talked about the last time we had this session. When each of you thinks about how you make a big decision, you need to consider how you—your current self—wants to perform some act or decide what to do in order to maximize the utility for your future self. The choices I'm especially interested in are ones that are life-changing decisions. As I said, they're high-stakes—the things that we care about very much.

What we ordinarily do is imagine ourselves into different possible scenarios: "Maybe I could do A, maybe I could do B, maybe I could do C. What should I do? How do I want to live my life? What kind of person do I want to be?" You

can think of this philosophically as what kind of future self do I want to become? What kind of future do I want to occupy? I care about what it's going to be like to be me after I undergo this central experience that's part of this big decision. That's the question I'm interested in.

As a philosopher, it's kind of funny to tell a bunch of scientists about a fictional example. The reason why it's important to look at these fictional examples, or at least the one about vampires I'm giving you, is because the structure is present in a number of real-life cases. It's important to get the structure out there so we can understand it. We're not worried about questions about morality here. Obviously those questions are important, but as a metaphysician I don't think about morality, I lack the relevant expertise.

We're going to pretend that modern-day vampires don't drink the blood of humans; they're vegetarian vampires, which means they only drink the blood of humanely farmed animals. You have a one-time-only chance to become a modern-day vampire. You think, "This is a pretty amazing opportunity, do I want to gain immortality, amazing speed, strength, and power? But do I want to become undead, become an immortal monster and have to drink blood? It's a tough call." Then you go around asking people for their advice and you discover that all of your friends and family members have already become vampires. They tell you, "It is amazing. It is the best thing ever. It's absolutely fabulous. It's incredible. You get these new sensory capacities. You should definitely become a vampire." Then you say, "Can you tell me a little more about it?" And they say, "You have to become a vampire to know what it's like. You can't, as a mere human, understand what it's like to become a vampire just by hearing me talk about it. Until you're a vampire, you're just not going to know what it's going to be like."

The question you need to ask yourself is how could you possibly make a rational decision about whether or not to become a vampire? You don't know,



and you can't know what it's like. You can't know what you'd be choosing to do if you became a vampire, and you can't know what you're missing if you pass it up. This would be a problem if we faced these choices on a regular basis because what it suggests is that there is a principled, philosophical reason why, when faced with this big choice, we would be unable to reach our epistemic gold standard.

If that were the only case in which this situation arose, most of us probably wouldn't have to worry about it, but I don't think it's the only situation in which this kind of thing arises. Now I want to talk about a case that is different in important ways from the vampire case because it's a low-stakes case. It's a little closer to real life, so we can see how this philosophical problem is one that we grapple with even if we're not always recognizing that we're grappling with it on a regular basis.

I've never tried a durian fruit. If you've tried a durian before, then bear with me. You can probably remember back to before you'd tried durian, and for those of you who haven't tried durian, we're in the same epistemic boat. The thing to know about durian is it's an exotic Southeast Asian fruit; it's very distinctive. One important chef says, "The only way to describe its taste is 'indescribable.'" The thought is, until you've tasted a durian fruit, you can't know what it tastes like. There are various evocative descriptions people have: "Eating vanilla ice cream by a sewer" or "French kissing a dead rat." These evocative descriptions are interesting, but they're not going to give you the information that you might like to have, namely, what it's like to taste a durian. The only way that you can know what it's going to be like for you is to taste one.

It's not about being sophisticated or liking exotic things because, as I already mentioned, even those with sophisticated palates, like chefs, differ widely on how they respond. Some people find it absolutely repulsive; other people call

it the king of fruits. Ambrosia would be the description. In this situation, when someone asks, "Do you like the taste of durian?" you don't know. You would have what I think of as an epistemic transformation if you tasted durian. Once you taste durian for the first time, you know what it's like.

The philosophical example in the literature that parallels this, about the value or what experience can teach you, is an example that was developed by Frank Jackson. He talks about black and white Mary. Mary, we suppose, has grown up in a black and white room. She's never seen color, she's just seen shades of gray and black and white. When she's finally let out of her black and white cell and sees a red fire engine, she learns something. She learns something that she couldn't have learned by reading all the literature about color science or about how we see or hearing testimony of other people. She learns what it's like to see red. The thought is that we can all recognize that there is something important that we gain by experience and by experience alone. We gain an ability to grasp a certain phenomenal concept. We gain a certain imaginative ability—the ability to imagine redness in various contexts.

This is important because if we think about what experience teaches us, then we can see how the puzzle that I was sketching with the vampire comes up again in the case of the durian fruit. Imagine that you're in Thailand. It's breakfast time, you're looking at the menu and you're trying to decide what you're going to have for breakfast. You have a choice between having some ripe pineapple for breakfast, or having some ripe durian. I'm going to assume we've all had ripe pineapple, and let's just assume you like pineapple, you think it's pretty good, but you've never had ripe durian.

The problem, when you're looking at your breakfast menu, is that you can't make a decision about what to have for breakfast based on which taste you prefer. Why? Because you've never had durian. You can't assign a value to the outcome of what it's like for you to taste durian. In a certain sense, the utility

of that outcome is not defined. If that's the case, then there's no way to make sense of determining how best to maximize your utility, or how best to respect your preferences in terms of picking whatever you would like best to have for breakfast that day. Because you can't assign a value to what it's like for you to taste durian, you can't, in a sense, have a preference, at least based on the way that we're thinking of the options. You can't step back and think about what the epistemic gold standard would be for you, that you should apply to yourself, when you're thinking about how best to choose what you want to have for breakfast.

When we're in context where we face epistemically transformative experiences, there's a way to make the decision that's just not accessible to us because we lack certain information, or we lack a certain ability. Why does this matter? As I said before, one way in which we assess our different options is by imaginatively projecting ourselves forward into different possible scenarios: "There's me having durian for breakfast," or "There's me having ripe pineapple for breakfast." We decide which scenario meets our desires in a more satisfying way, which scenario we assign a higher utility, then act so as to maximize that utility. That's the gold standard route.

In a low-stakes case, like deciding what you want to have for breakfast, there are other things we might want to do. We might say, "I value discovery. I'll just flip a coin. I'll just try durian for the heck of it. It's not a big deal." That's just fine in low-stakes cases. What matters are high-stakes cases. The vampire case I was describing to you is a high-stakes case. What makes it high-stakes is that it's both epistemically transformative and personally transformative. It's the personal transformation, the fact that it's going to affect the rest of your life and your very being, that makes it important.

These high-stakes cases are the cases where we care most about meeting the epistemic gold standard, or at least we should care most, because the decision

has big effects on you or maybe your loved ones. If any of those personally transformative decisions are also epistemically transformative, then the same problem we faced with the durian case resurfaces with the big decision case.

There are some real-life cases that have this structure. Let me just sketch two. The first case involves sensory capacities. Imagine a congenitally deaf person who has never been able to hear contemplating whether or not he should have a cochlear implant. Let's say he's built a lot of his life around being a member of the Deaf community. He's contemplating the possible outcome of getting a cochlear implant, and then presumably after he's learned how to interpret the signals from his implant, knowing what it's like to hear. Because he doesn't know what it's like to hear he can't, in principle, know what it's like to hear until he becomes a hearing person. There's a certain sense in which he can't assign a value to the outcome of what it's like to hear. It's a high-stakes case because, presumably, what it's like for him to hear is going to have a huge effect on the way he lives his life, and a lot of the features of his life.

It's not a matter of thinking more carefully or reflecting in a deliberate manner. For principled reasons, there's something that's epistemically inaccessible to him, and we can't expect him to make a decision based on information that he can't have access to. There has to be a different way to make that decision. Part of what I want to say is we need to recognize that agents can find themselves in that kind of epistemic situation.

There are lots of other cases involving disability and similar issues, but there's another case that is maybe a little more familiar to those of us who never had to face the possibility of having a cochlear implant: The choice of whether or not to have one's first child. Having one's first child is also an epistemically transformative experience. One of the most important and salient features of becoming a parent is what it's like to experience the attachment to the actual child that you produce—the loving, satisfying, attachment relation that you

stand in to the child that you produce. In order for you to stand in this attachment relation, first you have to produce the child. Second, the character of that attachment relation is going to be highly defined by the particular characteristics of the actual child that you produce. Until you stand in that relation, you can't know what it's like. You might know some very general features, but it's the particular features that matter and that are going to have the biggest impact on your experience of being a parent.

When you make the choice or you think about whether or not you want to become a parent, and you cognitively evolve yourself forward and imagine holding your baby and what it would be like to be a mother or a father, performing that act might be an interesting exercise in imaginative fiction, but it's not going to give you information about what it's going to be like for you to become a parent. That means that the utility of that outcome is not defined for you. And of course, this is a high-stakes decision. Becoming a parent is one of the classic cases where people's preferences and other things about their situation change dramatically. Often people do take themselves to be a different person. Some people say they're less selfish, they care about different things, they don't party as much. There are lots of different things that happen.

This is another case, one that many people face, and that is when they think about whether or not they want to have their first child, there's something important that's epistemically inaccessible to them. It's the thing that we care about, and the thing that's going to personally transform you if you have a child. When you contemplate whether or not to have a child, if you want to do it by assessing what it would be like for you to be a parent then it's, in principle, not possible for you to make that decision while reaching the epistemic gold standard. In other words, by acting so as to maximize your utility in the way that you understand to be doing it.

It's important to recognize this philosophical issue, and to recognize how a natural and intuitive way that we want to deliberate and introspect and think about who we are might be in conflict with the thought that rational decision-making defines our epistemic gold standard. (Decision theory does define our epistemic gold standard, so there's a real tension here.) There's a lot of value in introspecting. It's important for us to try to think about who we are and who we want to become when we make these big decisions. Yet there might be an in-principle conflict between this desire we have to be authentic in this sense, and the desire we have to reach the epistemic gold standard.

I want to close with another problem that comes up because there are a cluster of issues here. The other problem, which stems from the conflict between authenticity and the epistemic gold standard is the following decision theoretic issue. A natural thing to do is to say, "Let's just do some empirical research. Let's look at this question from a scientific perspective." I'm in favor of that. Doing empirical research on these questions is absolutely the best way to go, but imagine that we find ourselves in the following situation. Imagine yourself as a child-free person, as somebody who takes themselves to be essentially someone who is child-free. You're a person who has no kids, you love your life the way it is, and you think of yourself as intrinsically child-free; you have no desire to have children. When you think about what it would be like to be a parent, you think, "I wouldn't be happy, that's just not the life that I want to live."

You go around and you talk to people. Let's pretend that all the empirical research out there tells you once you become a parent, the way that you're going to evaluate the quality of life as a parent, the utility of becoming a parent, is going to skyrocket. Once you become a parent, you're going to think that being a parent is fabulous, that it's the best way to live your life, far better than it would be to live your life child-free. Let's say that all the description and testimony—from your parents, your friends who have children—all say

the same thing. Now you're in a situation where you value who you are as a child-free person, your preferences are to remain child-free. You also, in principle, cannot introspect into what it would be like for you to be a parent. If you were to be rational, you should replace your assessment, your imaginative projection, with this empirical information. All the empirical information and testimony that you have tells you that once you've undergone this experience, your preferences will change so that you'll be much happier—you'll be maximizing utility as a parent.

If we're to meet the epistemic gold standard, obviously we're supposed to be utility maximizers, right? If the way to do that is to listen to the empirical research in this case, then the right thing to do is to reject your current self and replace it with the future self that's a parent. There's a problem here. The way that I've set the case up, your current self assigns a reasonably low utility to becoming a parent, but because you undergo a transformative experience in virtue of becoming a parent, your future self as a parent assigns a very high utility to being a parent. Because you want to maximize utility, you should give up your present self and replace it with your future self.

That again illustrates the philosophical tension that comes out here. Some people think that to be rational, you need to respect your current preferences. To be authentic, you have to respect who you are now. It looks like if we want to meet the epistemic gold standard in this case, we have to violate the preferences of our current self—violate who we take ourselves to be now, who we take our current self to be—and replace it with a different self. That suggests that rationality can entail a kind of self-alienation that I find worth exploring.

## **THE REALITY CLUB**

**MOLLY CROCKETT:** That was so fascinating and engaging and touches on a