Rationally Speaking #183: L. A. Paul on "Transformative Experiences"

Julia Galef: Welcome to Rationally Speaking, the podcast where we explore the borderlands between reason and nonsense. I'm your host Julia Galef, and with me is today's guest, L.A. Paul. She is a professor of philosophy at the University of North Carolina, Chapel Hill. One of her areas of focus is what she calls "transformative experiences", which is a topic I've been interested in for a long time, although I didn't have this excellent handle to describe it.

Choices about transformative experiences are especially difficult to think about conceptually, philosophically for reasons we'll get into, and they also include some of the most important decisions that we have to make in our lives. With that introduction, Laurie, welcome to Rationally Speaking.

L.A. Paul: Thank you. Thanks for inviting me to have the conversation.

Julia Galef: Laurie, you're a philosopher, and so hopefully you won't mind if I ask you to kick things off by precisely defining your terms. What is a transformative experience?

L.A. Paul: Okay, so the way that I like to characterize transformative experiences are experiences that have both an epistemic and a personal component. A transformative experience is something that changes you epistemically. That's a sort of technical philosophical term, but you could think of it as it changes your conceptual framework or it changes your mind, it changes your understanding of the nature of reality. And a personal transformation changes who you are.

A transformative experience, the way that I like to talk about it, and I think it's important to talk about it this way, which we can get into, is an experience that changes what you know about yourself, or about the world, or about your relationship to the world. That change either scales up into a change about the way you think about yourself, or it causes a change in the way you think about yourself or understand yourself or in who you are.

The epistemic transformation, which is this change or sort of mental discovery that you make, then brings about a personal change, and it's a personal change that's not just minimal. It's not like getting your hair cut, but it changes what you care about, how you think of yourself or who you care about.

Julia Galef: Great. Maybe you could give an example of a characteristic change.

L.A. Paul: Sure. Okay, so the way that I like to think about these cases, it often involves starting out with a pretend case, a fictional case, because that way it's nice and clean, and we can imagine ourselves into the situation.

In particular, the case that I like is one where we imagine that you have the chance to become a vampire. It's a one-time only chance. Dracula himself has

discovered you and he tells you, "Look, with one swift painless bite, you'd be permanently transformed into an elegant and fabulous creature of the night."

Julia Galef: Do I glitter, though? That's really crucial for my decision.

L.A. Paul: Oh, yeah. Yeah, yeah. You might not glitter, but you will look really good in black.

Julia Galef: Okay. Excellent.

L.A. Paul: Yeah. That's a plus. You'll be really strong and powerful, and you'll have like incredible sensory experiences, but it's also irreversible. It's an irreversible transformation, and the problem is that you can't bathe in the sun anymore, you have to drink blood ... Maybe it's blood of humanely-farmed animals or it's artificially-created blood ... And the way that you'll regard the world and people around you is going to change.

So if you were presented with that choice, and it was a one-time only chance, the question is, would you do it? Now you might think, "Well, I'll just go and ask some other people." So you go and ask your friends or your mom, and they confess to you that they've already become vampires.

Julia Galef: Well if I couldn't *tell*, that's kind of interesting...

L.A. Paul: Yeah, exactly. Vampires are incredibly skilled at covering themselves up in some ways, so it's really interesting.

You talk to them and you ask them: what should you do? And they laugh at you, because they tell you that you can't possibly understand what it's like to be a vampire until you become one.

Okay, so why this matters is: the thought is that, you're faced with a choice where you could become a vampire, but as a mere human, you just can't understand what it's *like* to become a vampire, and so if you become a vampire that's gonna change the way that you understand reality. And your relationship to reality's gonna change how you experience the world. It's also gonna change you personally.

So that's a good example I think about a fictional case of, if you were thinking about undergoing a massive transformation. Real life cases, I think involve having a woman's first child, or ...

Julia Galef: Yeah, in fact I have to say -- I read your book, and you kick off the book with the vampire example, which is great. And the whole time you were describing it I was thinking to myself, "Boy, this really has a lot of parallels to deciding whether or not to become a parent."

...And then later in the book you introduce the idea of becoming a parent, and after introducing it you say, "And I should point out this has some parallels to the choice of whether or not to become a vampire." I was like, "Exactly!"

I mean not on the object level, it's just the parallels in the sense of being transformative. Although I guess you do kind of become undead for the first couple years of parenting, so there's some object level parallels.

L.A. Paul: That's right. You stay up a lot at night... That's right, that's right.

A couple other cases that are really interesting are: if you go into the military and fight on the front lines. I haven't done that, but I've been doing quite a bit of reading about that kind of case. And that's one where people testify to having basically a kind of transformative experience, especially if they have to kill someone or they're enmeshed in a really violent altercation. That's a bad kind of transformative experience, but it's transformative nonetheless.

Julia Galef: Yeah. I mean I think the most interesting cases are ones in which the people who have been through the transformative experience wax rhapsodic about it and say they're so glad they did it, but to you it seems like not clearly a good experience. But they have been through it, and they say that it's great and you just can't understand.

It's a little less interesting I guess when the person says, "You can't understand how terrible this was. You just have no way of imagining how terrible this is." And I'm like, "Great, I will take your word for it and not undergo that transformative experience."

L.A. Paul: With the case of going to battle it's interesting because people will say, "Oh yeah, it's terrible to be in battle," but then they'll also tell you things like, "I was forged in that battle," and that they're stronger now as a result, or their positive characteristics have come out of it.

So sometimes you get this kind of interesting, weird, "Oh it was terrible, but I'm glad that I went through it." Or, "I wouldn't be the person I am now if I hadn't." But you're right, yeah.

Julia Galef: There is a bunch of interesting stuff to unpack there, but before we touch on that, I want to ask whether think of the difference between a transformative vs. a non-transformative experience as a difference in kind, or a difference in degree.

Do you see a sort of spectrum of *how* transformative an experience is? Where on the one end we have, like, eating a sandwich -- and maybe I'm a slightly different person after I eat that sandwich, but not much -- and then at the other end would be like, maybe having your IQ multiplied by a billion. Or stepping outside of the simulation, or something like that, where it's just, basically impossible to imagine what that would be like.

| | And you know, having kids or being a vampire is somewhere on that spectrum.

Or do you see a sort of clean bright line that separates transformative from non-transformative experiences? |

L.A. Paul: Good question. There are some things as a philosopher I want to sort of try to distinguish there.

First, I don't think that usually having a sandwich, for example, would count as a transformative experience the way that I'm thinking about it. Maybe it's a little more tasty than any sandwich you've had, or it's a little colder, or it's a little different in some way -- so it kind of epistemically changes you a little bit, but it doesn't lead to any personal change. And so on that way of thinking about things, you made a minor discovery by tasting this new kind of sandwich or tasting this new kind of food generally, but it doesn't scale up into something that's transformative.

Julia Galef: Are you agreeing that it transforms you a tiny bit, but you just think that for practical purposes, we should round that down to zero? Or do you really think it changes you zero?

L.A. Paul: What I'm saying is I agree that it epistemically transforms you but it doesn't personally transform you, and the way that I think about transformative experience, it's gotta do both. So it's not a transformative experience, it's an epistemic change experience, is the way that I'd put it.

Julia Galef: I guess I'm imagining that my values, my personality, my worldview, change very gradually over time just as the cumulative result of all these little decisions that I make, all these little experiences that I have that add up. I just picked eating a sandwich as a random example, but surely all of those little actions that I take on an everyday basis are part of that gradual change in my personality.

L.A. Paul: Yes, yes. That's right. No no, so you can ... What matters here though is how we frame the events in question.

I do think that the kind of gradual change that individuals undergo, especially when they change from small children into adults, involve conceptual and personal change. So over, say, a ten-year period when you age from 10 to 20, then conceptually you're gonna add certain kinds of epistemic resources to your repertoire, and you're gonna change as a person, and those conceptual changes will lead to personal changes.

So yes, you've got an extended and slow transformative experience -- it's transformative not moment to moment, but we have to say it's a transformative experience from age 10 to age 20. But even if it's not transformative from moment to moment, it's still transformative starting at T1, which is at age 10, then by age 20 it's transformative.

Julia Galef:	Okay, so maybe eating a sandwich wasn't the most useful example I could have picked. What about something like going to a party, or meeting a new person, where there's at least some non-negligible chance that my values or worldview will be transformed a little bit by meeting these new people and talking to them? I can point to examples where I see at least a slight difference in who I am based on having met a new person.
L.A. Paul:	Right. So we need to distinguish again ... Sorry, distinctions, that's philosophers.
Julia Galef:	No, great. Hey, I love distinctions.
L.A. Paul:	So we want to distinguish the one possibility which is that some small minor event causes a change in you, or causes other sorts of small changes that can, let's say they *add up to* change; and also the small chance that some small event like having a sandwich will cause a dramatic change in who you are over a short period of time. Both are possible.
	Whether or not it's transformative in the sense that I'm interested in depends on partly on what you bring to it. If I've had say six children, I think it's unlikely that having another child is gonna be epistemically transformative ... It might be, but it's familiar to me in certain ways, so adding another child to the brood ...
Julia Galef:	Got it. You're distinguishing the expected magnitude of the transformation from the probability of the transformation?
L.A. Paul:	What I'm saying is that there are a number of different dimensions of this, and it's only when certain, one might say, "parameters" are set that you might get the kind of transformative experience that I think is interesting and that we want to pursue. But there can be messy cases around the edge, sort of edge cases, that are also interesting. It's not so much about the probabilities in particular or about the magnitude of it. You have to look at it on a case-by-case basis and say, "Oh, this bit's the natural kind of transformative experience, or the psychological kind of transformative experience." For one person it might, for another one it doesn't.
Julia Galef:	So it's sort of like saying, yes there's a difference in kind between games and buildings, but that doesn't mean that there aren't some cases where you're like, "Does this count as a game?" Like, playing hopscotch, is that really a game? There's no "win" in hop -- Well, maybe there is. Never mind. Cat's cradle or something. Or dancing. Is that a game? Kind of; not really.
	So it's fuzzy, but it's still a distinct category. It's not a sort of continuous spectrum from games to buildings.
L.A. Paul:	That's right, that's right.

Julia Galef:	Okay.
L.A. Paul:	And also, some of the fuzziness can be with the temporal extension. So I tend to look at very sharp transitions because those are easy to pinpoint, but you were right to say, "Well, you could have an extended one." I mean, the socialization process that happens when people pick out a professional career is often one that's a kind of temporally-extended temporal experience. Then you look back and you realize 10-20 years down the line, that you've become a different kind of person.
Julia Galef:	Yeah. I guess one reason that I was asking about whether it's a difference in degree or difference in kind is that, when you talk about transformative experiences, you use phrasing like, "You can't know what it's like." And it seems to me that that's ... I mean maybe that's just shorthand, but if it's not just shorthand it seems like: I know more than nothing. Maybe I've never been a vampire before, but I have been in the nighttime before, and I have been part of sort of exclusive or secret societies before. Or I've felt sort of like a discriminated-against minority before. I don't know, I can pick out components of the experience that I have some experience with. And unless you want to stipulate that they combine in such a way that is completely unrelated -- that the whole is *completely* different from the sum of the parts -- it still feels like I have more than zero information about what that experience would be like. There's not just a complete singularity between me and that experience.
L.A. Paul:	I think that's right. I think on the very edge it may be massive technological change or yeah, the singularity or something like that, you might get very extreme. The quality of life is fundamentally different and there's no way that you could know what that's gonna be like. But in most kinds of cases ... First, there's stuff that you know. For example, you might have evidence or testimony or descriptions from others, and it might be that you know something about, like you're saying, some of the parts of the experience. Couple things. That book, the book *Transformative Experience*, although it's pretty accessible to people who aren't totally enmeshed in philosophical debates about these questions, it still is written for those people, and so when I use the "what it's like" phrasing, I'm actually really trying to kind of pull out and say, "Look, there's a subjective quality of the experience that is really important to us, or at least potentially really important to us when we're thinking about making the decision, and we can't grasp that." So it's not that there aren't other elements of the subjective quality of the experience. Like becoming a parent. Well, I know what tired is now, I've been jet lagged, and there's a way in which I knew before I became a parent that it was kind of gonna be like being jet lagged. Now, okay, jet lagged for

two years instead of two days, and that difference is actually significant. But it's not like I didn't know anything.

So now, your point about, "Well, does the sum of those things that I can know make a different ..." Maybe there's three different things. One is: sometimes, those parts that you know about, when they sum together, even if you did know all the parts, I actually do think it can be distinctively different often. Which is very interesting with these complex experiences.

It's not just being tired and just feeling joy and things. But it's also that I think in the most interesting cases, it's not that you don't know anything about the subjective nature of the future experience, but you don't know the most important part.

Let me take the baby case because that's one that I like to talk about. The interesting thing about having a child... It's not that everybody who has a child experiences this kind of change, but for a lot of us, when we think about becoming a parent, one very important thing is, "Well, is this who I want to become? Is this how I want to live my life? Is this the kind of life change that I want to bring on myself? Because it's irreversible, obviously, and requires sacrifice, and so it seems like a big decision. You bring new life into the world, that's your thing.

And the thought is that the most salient part of having a baby and raising that child to adulthood involves the forming of the attachment relation between you and the child. So the mother-child bond... Or the parent-child bond.

And what's interesting about this is that this attachment ... It's a bit like having a close friend or loving your parent, but it's not quite the same. It's a very distinctive experience.

So while you're undergoing all these other things like running around like mad, feeling tired, feeling joy and amazement at the fact that you've created this life if you physically produced the child, feeling tired ... You also form this amazing bond to this other individual, and it's different from the love relation to a spouse, although it's intense like that in a lot of ways. That has a really distinctive character.

And I guess I think that that character then, both kind of dominates the whole experience in a really interesting way, and also *changes how you experience the other stuff*. So changing nasty poopy diapers is just not the same, when it's this child that you love with all your heart and you're responding to, but you're also exhausted and blah blah blah. That makes a difference.

So this is a long answer, but ... You can babysit and you can know how to do some of these things, but if you're not standing in that attachment relation to the child that you're babysitting, it just is a different kind of experience.

Julia Galef:	Yeah, and I'm very aware of sounding sort of laughably, naively, reductive here. But the claim that I'm making is not that having a child is basically just like babysitting plus being jet lagged, or something. I'm just trying to claim that having the subjective knowledge of what these experiences that are similar to some of the components of being a parent are like -- having that subjective knowledge, it at least *slightly* reduces your uncertainty about what the subjective experience of being a parent would be like, and that ...
	Well, I suppose I should put this question in context. Where I was hoping to go next was to the question of whether it's possible to make a rational decision about whether or not to undergo a transformative experience, given the fact that we can't know *exactly* what it will be like to have undergone that transformative experience.
	And I think ... We should talk about this, but my sense from having read your book was that you feel that the fact that we can never know *exactly* what the subjective experience will be like means that we can't make a rational decision about whether to make that choice.
	And I was going to claim that, as long as we can reduce our uncertainty *to some extent* with our past experiences, that are in some way similar -- or with outside information that seems relevant -- as long as we can reduce our uncertainty to some extent about whether we will be satisfied with that choice, then that's where rationality comes in. And that counts as rational decision-making.
	Sorry, there's a lot stuff there to unpack. You probably want to disagree with like three things that I said at least, but ...
L.A. Paul:	No, I'm totally great. So first, I think you're right that we can, shall we say, "refine" our uncertainty about the outcomes that we can grasp using some of the subjective information we might have, like knowing what it's like to be tired or knowing what it's like to work hard and those kinds of experiences.
	But I think the real root of the problem just isn't one that involves uncertainty. When you're uncertain about what's going to happen, then you have a clearly defined sense of the different possible outcomes, but you're not sure how, let's say ... I'm just gonna put this roughly, how "likely" each different outcome is to occur. Or maybe you're not sure about just what kind of a value you want to assign to these different possible outcomes.
	But the case that I'm interested in is one where, there's a a way in which you can't represent the outcomes that matter.
	Maybe you can get descriptions of them, but because you yourself can't imaginatively put yourself in the situation that you need to, you can't represent them in the way that you need to in order to either form defined preferences, or -- as a professor I would say, "assign them value" in the way that we need to.

Julia Galef:	So it's sort of like, you're distinguishing between on the one hand, "There's a 40% chance that I will feel jet lagged and a 60% I won't feel jet lagged and I know what both of those things feel like, I'm just uncertain about which one I'll land in," versus "I know for sure that I will feel flubbity-buh, but I don't know what feeling flubbity-buh is like," or something.
L.A. Paul:	Yeah, so compare, "There's a 40% chance that I'll feel jet lagged and a 60% chance that I'll feel great," to, "There's a 20% chance that I'll feel jet lagged, and a 40% chance that I'll feel great, and then there's a 40% chance that I don't know what the hell's gonna happen."
	Because there's just a big gap there. Or maybe there's a 20% chance of something –I-know-not-what, and another 20% chance of a different something-I-know-not-what. And so then, I think about decision making as like building a model, and it's like, you have a map that doesn't have all the spaces represented. There's just this big mushy gray area in a big chunk of the map, so you don't have any path to get there. You don't even know whether there are mountains, if there's rivers, what's going on there.
Julia Galef:	Yeah, it's a really interesting question, how to represent ... I do think there's some way in which it makes sense to talk about reducing your uncertainty about "what an experience would be like." Maybe the way to quantify that would be in like, *similarity* to what you're envisioning? I don't know.
	Maybe a different way to cash that out would just be in terms of how *satisfied* you will be in that experience, and so I can reduce my uncertainty about how satisfying being a parent will be, even if I can't imagine exactly what it will feel like.
L.A. Paul:	Okay, so, right. So here's what I'm hearing you say, and this is an approach I think that some people might find satisfying: Say, "Look, whatever. I don't have to be able to imagine myself in these different situations." Let's say somebody just tells me, "Oh, you're gonna love it" and they're right.
	I mean, one of the problems I have is that I just don't think these life experiences are that reductive. I think that they're complex and there's lots of varied sorts of responses we have, and that the subjective character of the experience isn't just reducible to satisfaction level, for example.
Julia Galef:	I agree, but there's still some mapping that could be useful to make, even if it doesn't capture all of the rich texture of the experience.
L.A. Paul:	Yeah.
Julia Galef:	I was just reading about how back before we had thermometers, people argued that the concept of measuring temperature made no sense, because you could never capture all the rich texture of humidity, and cool breezes ... And that's all true, that the actual number on the thermometer doesn't

capture that. But that doesn't mean thermometers don't measure something useful.

L.A. Paul: I think that's right, and I think a useful thing to do to kind of try to get to the root of this, is imagine that ... We're not in this situation, but imagine you had some kind of super scientist, I don't know, Google DeepMind builds some kind of super scientist that can tell people what their satisfaction levels will be if they become parents.

Julia Galef: I love that these kinds of thought experiments used to involve God or a supreme being. Now they just involved Google DeepMind creating a super powerful algorithm.

L.A. Paul: I know, I don't know if I love that, but I get ...

Julia Galef: For some value of "love," I love that..

L.A. Paul: I know, I know. I like to think of HAL from 2001: A Space Odyssey, because he was the original kind of Google DeepMind ...

Julia Galef: Right, right. Go on, I'm sorry.

L.A. Paul: So you get some kind of, right now not available but in our technologically-advanced future, somebody builds something that can tell you how you're gonna respond in terms of, "Oh, you'll be this satisfied." If we were in that situation, we could make the choice rationally.

Now, part of what I argue in my book is that we're not actually in that situation now, so we can't rely on that kind of specific information. We can talk more about, "Well what do you when you don't have that specific information?" And I said, "Well, we try to imagine ourselves in the situation and see how we would respond, and try to kind of correct for any errors."

And I think that we're telling ourselves a story when we try to imagine how we're gonna respond. We overestimate what we know about our inner selves in this weird way.

But even if we did have this kind of perfect algorithm to tell us how we're gonna respond, so we made the choice rationally -- then I think yeah, you make it rationally but then there's something else that we lose, and that is we lose the sense we have of musing through the possibility of what we're gonna do with our lives, and trying to think ourselves into these situations and making a decision based on who we think we want to become.

Because notice that, when the computer tells you whether you're gonna be happy or not, that has nothing to do with what you think. It's more like, "Well, I have assessed your personality type, and even though you think

you'd be a great parent and it's gonna be wonderful for you, your actually satisfaction is gonna drop, so no, it's not rational for you to have a child."

Julia Galef: Yeah, but -- okay, first I want to separate two things. One is the question of whether it makes sense to have preferences about what kind of person you become, independently of how satisfied you are once you are that person. I'm gonna set that aside for now, and just focus on this other thing –

-- which is whether you should care about the *manner* in which you make choices about your life. Because you seem to put a lot of weight on *how* you make a decision, independently of how the decision turns out.

That like, even if making a decision using empirical evidence about how satisfied other people like yourself were with their decision -- even if that turns out really well, and predictably so -- that's still somehow worse than making the decision by relying on your own internal simulations of what it would be like. Why would it be worse?

L.A. Paul: No, no no.

Julia Galef: No, I've misunderstood?

L.A. Paul: No, I don't think it's worse.

Julia Galef: Okay.

L.A. Paul: I probably just didn't say it very well. The thought is not that you should take information into account, the thought is that ordinarily ... So I want to distinguish the situation we're in, from the ordinary situation that we're in when we're not thinking about a transformative experience.

And part of my point is, we run those together. And it's really important not to run them together. So when I'm thinking about, "Oh, would I like to move to this part of this country?" And I get some empirical evidence about moving to California, let's say, and I know a little bit about California because I've been there before.

So I imagine myself living in California, and I sort of consider the statistical information that I have, in conjunction with my own reflections on what I think I like, and I come to a conclusion. And of course I have to do that because the evidence is not about me, the evidence is about people who are very roughly like me, like white women in my kind of demographic, it's actually pretty coarse-grained, at best.

Julia Galef: At best. Maybe you don't even have it broken down like that.

L.A. Paul: Yeah, exactly. Might not even have racial ... Whatever, you know, age, it's very very coarse-grained. In fact, when it's coarse-grained like that, there are

lots of properties I could have that could mean that in fact I'm not well-matched to that population, and so maybe I'd get a good response but there's probably a pretty good sized error bar there.

So, who knows where I fall on the distribution there. Hopefully it would be good, but there's always a gamble -- and we try to close that gap by thinking about ourselves.

But the problem with the transformative experience case is that that second bit where we close the gap is not available to us, and we forget that because we're used to making our decisions by reflection on who we are as an attempt to kind of inform our decisions when we're assessing statistical evidence. So what I'm worried about is making a decision purely on the base of statistical evidence, that's a big life decision like that. Now, maybe that's the only thing we can do, to be honest.

Julia Galef: Would you count it as making the decision based purely on statistical evidence if the way that you're closing that gap is not by running internal simulations, but instead making some kind of judgment calls about relevant ways in which you might be different from the average person in the population you have data about?

Like I could say, "I notice that most people really like playing with children, and I don't. And so the fact that most people end up happy as parents, that's a sort of baseline I can use to estimate how happy I would be, but I should adjust downward from that, given that I seem different in this relevant way."

L.A. Paul: I don't think that you can do that.

Julia Galef: You don't? Why not?

L.A. Paul: No, because what happens when you have a child is in fact the way that you respond to children, especially playing with your own child, is different. That's one of the bizarre things.

Julia Galef: I mean, somewhat different, but *completely* different? Do you really think that if you took 100 people who enjoy playing with children, and 100 people who don't enjoy playing with children, and you forced them all to have kids -- because in this thought experiment there's no IRB -- you force them all to have kids... do you really think that you wouldn't expect any average difference in how those two groups would feel about being parents?

L.A. Paul: That's not what I'm saying. You're asking whether you would ... You're saying, "Do I know enough about myself to know that my properties are enough like the people who stay the same throughout a child, or do I know enough about myself to know if I'm relevantly similar to the people who don't stay the same?" And that's the information that you don't have. The problem is that-

Julia Galef: Well I don't know *definitively*, but you don't think that I could reduce my uncertainty in a helpful way by looking at these traits, these similarities or differences?

L.A. Paul: I didn't say you couldn't reduce your uncertainty a little bit. Maybe you can -- although I think relying on anecdote and internal reflection is a pretty iffy and unscientific way to try to reduce that -- but even if you do reduce it, the problem is you're not gonna be able to reduce it very much.

                 I'm not really a fan of relying on anecdote, and local, what my friends say. On the other hand, I've had plenty of friends say, "You know, I really didn't like kids, but now that I have one of my own, I at least like my kid. I like playing with my kid. I might not even like playing with any of those other kids, but I can spend hours reading storybooks to my own child." This is not an uncommon experience.

Julia Galef: Okay, so maybe that's not the best -- But surely there are some-

L.A. Paul: But this is the weirdness of it. Let me add one more thing. We haven't covered this, but this is ... One of the interesting things about certain kinds of transformative experiences, and I think being a parent is one of them and I touched on it with the military case as well, is that weirdly I think the experience itself can form preferences in you.

Julia Galef: Yes, I completely agree. Yes.

L.A. Paul: Okay, so the problem is, is that ... Let's say that you don't really like kids, and you're a woman who's really interested in doing lots of things with her life and one of them is not having children. And you see these women, they're exhausted, and they clearly have taken career hits, and you think, "God, I would never want to do that."

                 But then you fall pregnant through some kind of happenstance that, it just happened. Maybe your birth control failed or whatever, and you think, "Okay, I'm gonna do this," and you have the child. That experience of having the child and forming that identity-defining attachment relation that I was talking about can disrupt and change your preferences, so much so that you are grateful to have the very child that you had. And this is bizarre.

Julia Galef: I guess when you really look at it, it is bizarre. I'm so used to this fact about the world that it doesn't immediately seem bizarre to me until I stare at it, but yeah.

L.A. Paul: So then go back to statistical evidence, and the question is... So the computer tells you, "Oh, you're gonna be so happy if you have a child," and you think, "No way. There's no way I want to have a child." But what's really going on is the computer actually knows in some weird sense, is calculating for the fact that when you have the experience your preferences are gonna change so that you are satisfied at the other end of it.

And then the question is, should we just kind of do what the computer says because the experience is gonna morph our preferences in a way so that I'm happy at the end, or should I be true to myself right now? Because right now I don't want to have a kid.

Julia Galef: Well, in my mind ... I'm gonna be kinda simplistic here, but just roughly speaking it seems like the rational thing to do is to trust the computer's prediction to the extent that the computer seems to have a lot of data that's relevantly similar to you. So trust the computer's prediction about how happy your future self will be.

But then *also* separately ask, "What kind of person with what kind of preferences do I want to be?" Because I do think that it makes sense to have preferences *about* your future self's preferences.

L.A. Paul: Yes, that's right.

Julia Galef: Like you could tell me, "Hey, take this pill. It'll make you insane, but happy." Well, I guess that doesn't really count as preferences, but personality, let's say.

Certainly if I had to choose I'd rather be happy insane than miserable insane. But I have very strong preferences about not being insane, independently of how happy I am in that state.

And I think it can similarly make sense to say, "I like the idea of becoming the kind of person who just cares far more about her child than about her career, even if that's not the kind of person I am now." Or it can also make sense to say, "I *don't* like the idea of becoming the kind of person who has this strong attachment to a child at the expense of her career."

And I don't think one of those preferences is more correct, but you probably will notice that you have those preferences about your preferences, if you introspect. And I think it's totally rational to take that into account in making your decision. No?

L.A. Paul: No, the problem is not when you want to change your preference and the computer tells you you'll be happier because your preferences will have changed. The problem is when, at the time when you're reflecting on whether you want to have a child, you prefer to remain the kind of person you are. You prefer to remain-

Julia Galef: Do you? I don't always prefer to remain the kind of person I am.

L.A. Paul: No no no, I wasn't saying "always."

Julia Galef: Oh.

L.A. Paul:	We're talking about a particular kind of case, so it's not that ... You don't have a rationality conflict in every kind of case.
	But there is a particular class of cases where there is a problem. And the problem exists when you're at T1 when you're thinking about your choice, and you prefer to remain the kind of person that you are, you don't want preferences to be changed. But the computer tells you, "I've done all the empirical work, and actually you'll be more satisfied afterwards, because replacing your preferences will make you more satisfied." And you think, "I don't want to replace my preferences!"
	So what I'm saying is, there are lots of interesting transformative experience cases where something like that is the situation that people find themselves in. And so if you think about, "Oh, well I'll just decide whether or not to have a baby by... trying to figure out if really by having this child, I'll reveal an unknown preference to be a mother"...
	And I think no, that's not actually what it's about. It's actually, you go into this experience, it changes you from the outside and implants a preference in you, so you become a different kind of person, and that's the way lots of life experiences really are.
Julia Galef:	So that certainly seems to make it a *trickier* decision, but do you think that that makes it a decision that's impossible to think about rationally? Or a decision that's impossible to make rationally?
L.A. Paul:	No, what I think is that if you try to make the decision based on simulating your future self, before you've done something that involves a transformative experience, that you don't have the information you need to be able to compare the different options and make a decision rationally, because you can't define the outcomes.
	If the computer tells you the result, then you can make the decision rationally if you just follow what the computer says -- but the problem is, is if you find yourself in a case like the one we were discussing before, where you actually don't want to change your preferences. Then I think there's a conflict because we don't have any higher-order way to solve whose preferences matter more, the person that you start with or the person that you end with?
Julia Galef:	Why can't I just use my current self's preferences about both how much I value being happy in the future, and also how much I value keeping my current identity? There's probably some degree of happiness at which I would be willing to change who I am fundamentally in a way that I don't think.
L.A. Paul:	I think that there are ways of trying to resolve ... In the book I talk about this a little bit. I say, "Look, if we want to try to make these choices rationally, what we need to do is understand these different sort of deeper questions at

play here," and then step back and see if in fact, we have preferences -- let's say either to discover what it's like to be a parent, so then my preference is to allow my preferences to be changed, and I prefer that... Let's say the computer tells me I'll testify then to having this higher level of life satisfaction.

Or I can have a preference to basically remain who I am, because that's the life that I know and that's the life that I value, and so then I have a kind of higher-order preference to weight that over the other.

I think that's all fine. That's really the point of the book, is that these questions I think are weirder and more puzzling and kind of deeper than people have explored.

Julia Galef: Oh, that I totally agree with, yeah.

L.A. Paul: Okay, okay. And so the-

Julia Galef: Where I thought we disagreed, which doesn't in fact seem to be the case, was that I thought that you were saying that the fact that we can't know for sure ... the fact that we can't quantify our uncertainty *precisely* until it collapses into risk, like rolling a die, that we can't therefore make the decision rationally. But that's not what you're saying?

L.A. Paul: Right, so I framed some discussion in the book in a way that fits exactly what you said.

                 But what I do is I have this kind of big caveat, which is to say, "Look, ordinarily when we think about these cases, we just don't go that far." We just think you're supposed to ... I read around some of the self-help literature or some of the stuff on, "Should you have a baby?" And you just are supposed to kind of think yourself into your future self and make a decision that way, and I say, "That's not rational." That's what I'm arguing against.

Julia Galef: Okay, that I agree with. And in fact I think there's this opposite failure mode, that maybe is more common among sort of analytical engineers in the Bay Area that I hang out with, that just says, like, "Well, just look at the empirical evidence about whether parents are happy. Looks like they're less happy than non-parents -- well there you go, that's your answer." And that also seems like a very poor way of making a decision.

                 In part because ... I guess we didn't quite get into this directly, but I think that the question of what your subjective experience will be like after you make a choice does not have to be the end of the considerations you take into account in your decision-making.

                 You can totally have preferences that don't cash out in terms of your subjective experience. Like a preference to pass on your genetic material. Or

a preference to have accomplished something big in your life, even if you don't feel satisfied during most of the moments that you're accomplishing it, you know?

L.A. Paul: I think that's exactly right. I think there's this other dimension of some of these life experiences where sometimes people choose suffering, and they choose suffering because they-

Julia Galef: Knowingly.

L.A. Paul: Yeah, knowingly, because they also value the kind of life experience that's involved in that, or the performance of that whatever, achieving that goal or doing something for others.

                So sometimes, especially some of the empirical evidence seems to kind of be underdeveloped. That's not a criticism, really, of the empirical evidence, because it's extremely hard to measure these things, and so it's not like there's some obviously easy way to measure this.

                But that's just another reason for not just unthinkingly taking on some of these results -- and kind of thinking either that you can just do it by imaging yourself into the situation, and figure it all out that way, or you can just take on the scientific, the statistical evidence, and not have to think about it that way either.

Julia Galef: Right, right. Do you want to talk at all about this concept of ... I think I'm remembering correctly, the "revelatory" factor in decision making? Because that was sort of an interesting new angle on the question.

L.A. Paul: Yeah. I mean it does connect to something we were talking about, because -- and my thinking is still evolving on this, but -- I like to think that there's value in discovering the nature of reality as I experience one. Maybe this can be captured, but, you go to a new country. Like I've never visited Japan, I've always wanted to visit Japan. I hope I'll get there one day, and *when* I get there (because I'm gonna go) --

Julia Galef: That's the spirit, I like it!

L.A. Paul: Yeah. I'm gonna have lots of contact with people and parts of culture that are different from anything I've had before, and I will discover what it's like for my mind to come into contact with those bits of reality. This is a way to put it.

                And you can think of it as, "I'll discover reality," but some of us might think, "Well, it might be a little more complicated than that." And so what you're really doing is discovering what it's like for you to come into contact with other parts of the world, and then kind of creating a really interesting rich

life experience that can involve other people. And that's part of what life is about.

So there's this kind of discovery element to a rich and full satisfying life, that involves revelation. Namely, allowing the world to reveal itself to you as you go out and experience it. And I think-

Julia Galef: And is that ... Yeah, go on. Sorry.

L.A. Paul: I think having a kid can be the same kind of thing, that's the short answer.

Julia Galef: Okay, cool. So is the idea then that the revelatory value of a choice could be an additional factor to weigh into your decision-making? In addition to how much you predict you'll enjoy yourself after having made that choice, and how much you like the idea of the kind of person you'll be after making that choice, et cetera -- that the revelatory value should be added in as well?

L.A. Paul: Yeah, I think you're right.

I say something a little more radical in the book. I say something I think is kind of unsatisfactory, which is: in the edge cases, maybe that's all we have, is the revelatory value. Because if we're really considering like a dramatic change where we can't think ourselves into the future in a relevant sense, and there's no empirical evidence, then all we've got is, like -- we're standing on the edge of a cliff. And we have to decide if we want to leap into the abyss and discover what's there.

(That makes it sound bad! Where it's not supposed to be either bad or good.)

...Or stay where we are.

Julia Galef: To step out of the simulation, then.

L.A. Paul: Yeah, yeah. To step out of the simulation. Do you take the red pill or do you take the blue pill?

But I guess, again just coming back to what I was saying, I think that real life is like this in a lot of ways, and what I'm trying to do, in the work that I'm doing, the sort of analytic philosophical work, is to identify the ways in which real life is like this. And find the kind of philosophical structure and press on those parts, so that we can discover more about it. And not underestimate the importance of the experience, but also not underestimate the kind of challenges that we face when we think about these decisions.

Let me just get one more thing in, because something I'm really interested in is, when we do advanced directives, or when, say you're facing the possibility of cognitive decline over the next three years; you find out that you've got early onset Alzheimer's and you have to make decisions about

your future life, or you have to make decisions for someone that you care about...

I think we're in a really difficult epistemic position there. And recognizing how difficult how it is for people to know what to do in those situations, I think, is a part of respecting people. And being tolerant of mistakes that people make and not blaming people... And also understanding what they're facing, so that we can support them in the right way.

And I guess I've been concerned for a long time that a failure to recognize some of the really weird structural facts about transformative experience has kind of, maybe we've underestimated those issues.

Julia Galef: Yeah, that does seem very important. That does seem like a very important practical takeaway from this area of philosophy.

I'm wondering if there are other ways in which, just practically speaking, you think having this framework in mind changes the way someone should approach having children.

Like, usually *I'm* the one saying, "Well, things are actually much more complicated than you think and we should have much more uncertainty than you think." And so I have a lot of sympathy for that perspective.

But at the end of the day, you have to make a choice somehow. And I feel bad just leaving people there, and saying, "Look how complicated and uncertain everything is! There you go, have fun, I'm off."

So for example, you might say that, "Well in practice, people ..." -- I don't know that *you* would say this, but something someone could say in response to my question is -- "Well in practice, people tend not to allow themselves to take into account the fact that they will be a different person. And if they *did* allow themselves to take that into account, and really stared at it, they might not take the leap as often as they otherwise would. Because in fact they do have preferences about their preferences change, and if they are ignoring that, then they'll end up making choices that they maybe rationally shouldn't."

Or something. So that would be a difference you would expect to see in people's decision-making after having this framework in mind.

Maybe that's horribly wrong, but ...

L.A. Paul: Well, sadly I don't have a lot of positive suggestions. I seem to be better at finding problems than having them.

Julia Galef: Well, that is the role of the philosopher! So I'm asking you to step out of your comfort zone, to some extent. Or out of your prescribed role.

L.A. Paul:	I guess the practical thing, as someone who's had children -- and was, to be honest, uninterested in having them before I had them, so I was fascinated by this whole kind of weird process -- and yet I love my life now. And I just don't know what to say about ... I don't want to tell other people how to live their lives.
	That said, I guess I want to say that recognizing that it isn't straightforward, and that there can be these problems that people face, and there is -- right now, I think -- no obvious way, or no satisfactory way, to resolve them... means that we should all be more tolerant of each other.
	And if people can make these decisions, maybe not in a perfectly rational way, but a perfectly adequate ... There's a lot permissible here. And people shouldn't blame themselves if they make a mistake. And they also shouldn't congratulate themselves if it turned out well, because I think it had a lot less to do with foresight than ... It has a lot more to do with luck. And a lot more to do with maybe structural facts -- that just kind of happened to work out, or happened not to -- than us soldiering on as carefully planning things, and looking at things, and realizing ourselves in various sorts of ways.
	I just think that's a kind of interesting picture that puts too much individual responsibility on us. That's the American way, I guess, but ...
Julia Galef:	Yeah. There's something liberating about that, about making that move. Yeah.
L.A. Paul:	That's my hope, that people would feel a little bit more liberated.
	Especially those of us that in a sense have to buck the system. Like women in today's society who are really committed to their careers often have to make really tough choices when they decide to have a child. And so the world hasn't made it easy for us to do this, and I think that one should be forgiven for not always being able to sort of prescriptively predict what the future's gonna be like.
Julia Galef:	Yeah. Okay, well maybe that's a good place to wrap up. Before we end the episode, I would invite you to give the Rationally Speaking pick of the episode, which is a book or a website or a blog or something like that that has influenced your thinking in some respect. So Laurie, what's your pick for today's episode?
L.A. Paul:	Thomas Nagel's *The View from Nowhere*. Thomas Nagel's a great philosopher and his work is fascinating and thoughtful, and I think that's a great book.
Julia Galef:	Is there any particular aspect of the book that you want to highlight that influenced you?

L.A. Paul:	Yeah. One thing I liked about it is: I've always been a person who cares a lot about exploring the nature of experience. And contemporary philosophers have not always been so interested in that question. But Nagel is an exception.
	And he worries about whether or not it can make sense to ever have a view from nowhere, and sort of eloquently defends-
Julia Galef:	And by "view from nowhere," you mean a view that isn't colored and shaped by the person having that view?
L.A. Paul:	Yes, that's one way to put it. There's this kind of picture that you could have, some kind of birds' eye, or God's eye, view of objective reality that isn't colored by the subjective lens, one might say.
	And he does a really nice job of articulating these issues and defending the thought that there *always* is a view from *somewhere*. And he's a good writer. I think it's just a great book for people if they want to start exploring some of these questions.
Julia Galef:	Great. Well, we'll link to *The View from Nowhere* on the podcast website, as well as to your personal site with links to your book, and a bunch of your popular writings on transformative experiences and other areas of philosophy as well.
	Laurie, thank you so much for being on the show, it's been a pleasure having you.
L.A. Paul:	Thanks for having me, it was really fun.
Julia Galef:	This concludes another episode of Rationally Speaking. Join us next time for more explorations of the borderlands between reason and nonsense.