

L.A. Paul

The First Time as Tragedy, the Second as Farce

It's widely accepted that experience teaches you things you can't learn in other ways. When you see the colour red, you learn what red experience is like. When you undergo an extremely painful procedure, you learn what it feels like. If you've never had those kinds of experiences, you can't know what they are like.¹ There's something important and distinctive that experience, and experience alone, will teach you.

But what about when it doesn't?

Montero (this issue) argues that, despite the fact that we can learn something special and distinctive from experience, we shouldn't overestimate what we gain from this. Sometimes experience doesn't teach us, in any lasting way, what something is like. Somehow, even after having the right sort of experience, we don't seem to learn as much as we'd expect.

One thing that experience can't teach us is how to explain what something is like to another person. I can experience red, and know what it is like as a result, but this isn't enough to teach me how to describe it to another person. Even though I know what it's like, I couldn't describe what it's like to see red to save my life. Knowing what it's like won't save me. The problem seems to be that language just can't function effectively as a vehicle for the expression of this sort of content. We aren't able to capture the qualitative nature and character of phenomenal states using descriptions or testimony.

Correspondence:

L.A. Paul, Yale University, New Haven, CT 06520, USA.

Email: laurie.paul@yale.edu

¹ At least when you face real world constraints on what you'd know beforehand.

Metaphor and poetry can do a slightly better job but, in general, language fails us here.

But, as Montero points out, at least sometimes, there's another thing experience can't teach us. There's a way that discovering the nature and character of an experience might not imprint on us in a way that allows us to qualitatively represent it to ourselves for our future use. This isn't just about language. It's about not being able to represent things in the right way, even to oneself, after the fact. For certain sensations, especially when they were intense, we can't recall *what they were like* in a way that will enable us to accurately represent them. Ordinarily, once we experience a token of an event type, we can recall enough about it to projectively assess the nature and character of future tokens of that type. But if Montero is right, in certain cases we fail to form this sort of representation. Somehow, in these cases, after having the experience and (in the moment) learning what is like, we don't retain knowledge of what we've learned.

This is important, and it rings true. Montero focuses on examples involving extremes. The deep hunger that comes from long-term denial. The excruciating pain of childbirth. Passing a kidney stone. But her insight extends to much more mundane cases as well. Try as you might, it may be impossible to recall the precise taste of your favourite lemongrass pandan tea, or the scent of the crepes your grandmother used to make.

Montero uses her insight to develop the notion of a 'qualitative memory', a memory of something that captures what it's like, and to contrast that to 'non-qualitative memory'. An absence of qualitative memory seems to arise with many kinds of sensations, especially if they are unusual or far enough in the past. If this type of memory is not available to us, when we recall events, we can only remember them 'non-qualitatively', that is, we do not remember the nature of their phenomenal quality or character. If she is right about the phenomenon, and I suspect she is, it has interesting and important implications. One particularly interesting implication concerns the way we reason about events and actions and choices.

Sometimes we reason about events and actions using 'model free' reasoning (Crockett, 2013.) Model free reasoning includes actions like using trial and error to come to a conclusion. A robot exploring a new surface might use a randomized sampling technique to select points from which to explore properties of its location. A person in an unfamiliar place might do something similar, especially if they are lost or disoriented.

Another way we reason about events and actions uses ‘model-based’ reasoning, which employs planning strategies. Perhaps you want to remodel your kitchen or you need to find the best way to get to Times Square. For this sort of reasoning it’s better if you think about it differently from when you are lost and disoriented. You don’t use trial and error. Instead, you use what you know to cognitively model, simulate, or explicitly represent your situation, and then use this representation to decide on the course of action that best fits your goals.

We use this sort of model-based reasoning for retrospection, for example, when we look back to past actions or events to assess, judge, and draw inferences about what happened. Perhaps you like to play soccer or run races. After a game or a race, you might mentally review your performance, thinking about ways you could have kicked the ball or trained for the final heat to have gotten better results. Such reflections can help you to improve future outcomes. Or perhaps, after a difficult discussion in a faculty meeting, you might reflect on ways to improve collegiality and encourage more productive debate. Or you might retrospectively assess your performance in a job interview in order to improve how you handle aggressive questions. After giving a talk, you might reflect on what you learned from the audience response. Model-based retrospective reasoning, then, is especially useful for reflection, for assessment, and for discovering what one has accomplished.

Another important way we use model-based reasoning is to prospectively assess different actions and events. We do this to psychologically prepare for future events, to make plans for what might come, and to make choices about what to do. For example, you prospectively reason in a model-based way when you use a map to plan a hike to the top of a mountain or when you figure out how to get from an unfamiliar airport to a new hotel. You use it when you plan a lecture, structure a class discussion, or organize a series of meetings for a busy day.

A particularly important type of prospective reasoning involves the imaginative anticipation of future possibilities, especially when we want to anticipate possible consequences of our acts. One reason we might want to do this is to prospectively assess our preferences. If we are to act in a rational way, we need to act in order to maximize our expected value. This requires us to assess our preferences about the different possible consequences for different possible actions, and to

choose in accordance with them. Do we prefer consequence A over consequence B? If so, how should we act?

To assess our preferences, we need to represent the states of affairs that they concern. Assessment can take different forms. Sometimes we model to *identify* our preferences with regard to possible consequences, and sometimes we model to *construct* our preferences with regard to possible consequences. If you are going to remodel your kitchen, you might have a clear idea, right from the start, of what you'd like with respect to the different construction options. But if you don't, you need to assess your preferences. To do so, you might find it extremely helpful to consult drawings or mock-ups of different architectural plans. Perhaps consulting these drawings will help you to recognize your preferences about what you'd prefer: your preferences might be latent in some way, but brought to your awareness through consulting the visual representation. Or, perhaps you do not even have preferences, latent or otherwise. Rather, consulting the drawings allows you to form preferences about these consequences. In either case, whether your preferences are latent or don't exist, before you act, it is important to have a representation (the mock-up) with the right content. You need to know this content to accurately assess your preferences and make your choice about how to act; in this case, about what kind of kitchen renovation to undertake.

We do the same kind of assessment with retrospective reasoning. After you give your talk or leave the faculty meeting, and you reflect on what happened in order to form a judgment about the outcome, you need to be able to consult an accurate representation. You need a representation with the right sort of content in order to accurately assess the event. What does such an assessment involve? In order to have preferences about possible outcomes, you need to know what matters to you. In other words, part of assessing your preferences about an event requires the assessment of its value (or disvalue).

For these reasons, Montero's point about our inability to recall the nature of an experience has implications for the way we draw on prospective and retrospective model-based reasoning in the determination of value assignment and forming preferences.

In many practical contexts, in order to know how to prospectively assess, we turn to retrospective assessment. Using model-based reasoning, we draw on memories to determine anticipations. In cases where we are planning for a future that involves a repeat experience, we want to use our reflections and judgments on what we've already

experienced in order to make informed choices for the next time around.

In the cases of interest, we assess our values for events in our past to predict our values for qualitatively similar events in our future. We use these assessments to make comparisons in order to discover or construct our preferences. If we lack access to our value function for the qualitative character of a past event, we cannot use it to define our value function for the qualitative character of the future event. This means we can't use our experience to inform our prospective assessments or anticipations in the way we'd like to. If we can't compare the value of event A to event B, we can't have preferences about A vs. B. That is, we can't discover or construct our preferences about which event would be better to bring about. In a case where the qualitative nature of the events matters to us, we have a problem.

For example, let's say you need to make a judgment about some event that you have had personal experience with; for example, a painful surgery. Perhaps you need to decide whether it is worth it to you to have this type of surgery again in order to address a chronic condition. Perhaps you need to decide whether another person should have this type of surgery, such as your child or your aging parent, and to do so, you need to balance their psychological ability to handle pain with other considerations about their quality of life. Ordinarily, you'd expect that, as the result of your previous experience with that kind of surgery, you'd be in an epistemically favourable position with regard to making this judgment. You'd expect to occupy a position of epistemic privilege in this kind of practical reasoning context.

However, if your qualitative memory of your past surgery is elusive, you may lack the ability to make an accurate (qualitative) judgment about the disvalue of having the surgery. For example, if you cannot recall the nature of your painful surgery, yet you must choose between various types of surgeries, you may be unable to accurately compare or contrast their value against the true (qualitative) cost of the surgery that you've already had. If you cannot accurately recall the nature of the pain, you may not accurately anticipate its disvalue in the future. You are also likely to be especially prone to temporal discounting and other forms of bias in the value judgements that we constantly make as part of the exercise of our ordinary decision-making capacities (see, for example, Sullivan, 2018, and Ainslie, 1991).

More generally, if you cannot recall the qualitative nature of an experience, this can affect your ability to assess its value. If you can't accurately recall its qualitative character, you can't retrospectively

assess it to discover or inform your judgments. If you cannot represent the event in a way that allows you to accurately assign value, you may not be able to realize latent preferences or create new ones in model-based reasoning contexts.

I suspect this happens often in contexts of practical reason. To go back to Montero's example of feeling hungry, even relatively ordinary cases of qualitative recall seem to fail in an utterly mundane and regular way. Who isn't familiar with the experience of making ambitious diet plans right after a filling dinner, or early-morning exercise plans in the comfort of an evening's repose?

As a boy I saw a movie about Admiral Byrd's first Antarctic expedition and was impressed that as a boy he had gone outdoors in shirtsleeves to toughen himself against the cold. I decided to toughen myself by removing one blanket from my bed. That decision to go to bed one blanket short was made by a warm boy; another boy awoke cold in the night, too cold to go look for another blanket, cursing the boy who removed the blanket and swearing to return it tomorrow. But the next bedtime it was the warm boy again, dreaming of Antarctica, who got to make the decision, and he always did it again. (Schelling, 1984, p. 8)

My own nemesis seems to be the healthy, lean, homemade lunch that always goes uneaten once the taco truck is right there in front of me. So I get the joke: warm Schelling just consistently fails to appreciate just how miserable cold Schelling will be.

Part of the joke is that it isn't just weakness of will that's the culprit here. It's lack of access to one's future qualitative judgments. If you misidentify the cost of an experience, you may choose to undergo it again despite the fact that it does not maximize your expected value, and so despite the fact that it does not or will not accord with your preferences when the event occurs. You'll have made the wrong choice, because you failed to accurately assess the value of the event and thus failed to accurately develop and assess your preferences. In my own case: I fail to recall just how appealing that taco will be when it's right in front of me. And even though it isn't the first time I've done this, I never seem to learn.

All of this, as Montero points out, is deeply related to issues that arise in the discussion of transformative experience. In particular, part of what makes an experience transformative is that its subjective (qualitative) value is only revealed to the subject through experience. This relates to the transformative decision problem developed in Paul (2014), concerning the epistemic changes grounded in having new types of experiences. In contexts where we face a decision about

whether to undergo a transformative experience, if we want to decide rationally, we need to be able to prospectively model future events in order to assign values to our options. And the trouble is, if the experience will be transformative, we can lack the ability to imagine in the way we need to in order to prospectively model and assign values in the right way. If we can't assign the needed values, our subjective value function goes undefined. A second part of the transformative decision problem is that act–state independence is violated, in the sense that the agent's preferences change as the result of undergoing the experience. This can result in a conflict between the preferences of the person before they have the experience (the *ex ante self*) and the preferences of the person after they've had the experience (the *ex post self*). This part of the problem raises the issue of whose preferences should determine the choice: the *ex ante self*? Or the *ex post self*? (See Paul, 2015, and Paul and Quiggin, 2018, for discussion.)

The problem with defining the value function arises whether we need to qualitatively represent an event to discover our latent values and preferences, or whether we need to represent an event in order to create our values and preferences. In either case, if the value function goes undefined, the model is undefined, because we can't calculate expected value. In short, if we can't accurately assess the subjective (dis)value of a future experience, we can't form accurate preferences concerning it.

In my 2014 book, I assumed that, at least *ex post*, experience could teach us what the relevant values were, and in this sense, past transformative experience could guide us when confronting future events of the same type. But if Montero is correct, then even having the experience may not get us out of the value-assignment problem. Certain types of experiences may just be gifts that keep on giving.

That is, Montero's work suggests that there are new problems with the assessment of subjective values in practical reasoning contexts. As with standard transformative decision contexts, in certain cases of repeat reasoning, the decision model could fail for principled epistemic (or psychological) reasons involving the lack of a defined subjective value function.

There's another interesting psychological wrinkle here. The intensity of the experience may also contribute, in some cases, to our practical inability to reason rationally. In the throes of an intense experience, such as childbirth, we may not be able to reason or assess in the ordinary way, adding a further practical constraint on the rational assessment of our subjectively intense experiences. We may

only be able to assess and judge properly once we've gotten past the pain, because it's only then that we recover our reasoning abilities.

But, of course, by then it's too late. By the time we are able to do the model-based reasoning needed to discover or construct our judgment about the horrifically painful event we've just recovered from, we are unable to do so. We can know *that* it was bad. But just how bad, and thus how we'd trade off that pain (in the moment) for other options, is inaccessible. During the event, when we can know what it is really like, we are unable to perform the complex reasoning required for rational decision making. And afterwards, we are no longer able to know what it is like.

As Montero's paper suggests, we can apply this to certain well-known examples. Take the case of choosing to have a child. Giving birth, for many women, is transformative. But what about having a second child, or a third? An important part of the calculation for many women is the physical toll of the gestation and birth process. The pain and exhaustion can be significant. If you've already had the experience of gestation and childbirth, it would seem that you could factor in the disvalue of the process in your overall assessment of the positive and negative utility of having another child. So you might conclude that it's only the first time that having a child is transformative. But as Montero suggests — think again! Perhaps, each time you reflect on the life-changing possibility of having another child, your brain plays a cruel epistemic joke on you, keeping you from accurately judging just how much it's going to hurt. If so, then you can't build a model that takes that pain into account, no matter how much experience you have. If so, Montero is indeed correct when she says that 'the standard philosophical accounts of experience, rational choice, and the sources of moral action may all need revision'.

References

- Ainslie, G. (1991) Intertemporal choice: Derivation of 'rational' economic behavior from hyperbolic discount curves, *American Economic Review*, **81**, pp. 334–340.
- Crockett, M.J. (2013) Models of morality, *Trends in Cognitive Sciences*, **17** (8), pp. 363–366.
- Montero, B. (this issue) What experience doesn't teach: Pain-amnesia and a new paradigm for memory research, *Journal of Consciousness Studies*, **27** (11–12).
- Paul, L.A. (2014) *Transformative Experience*, New York: Oxford University Press.
- Paul, L.A. (2015) Transformative choice: Discussion and replies, *Res Philosophica*, **92**, pp. 473–545.

- Paul, L.A. & Quiggin, J. (2018) Real world problems, *Episteme*, **15** (3), pp. 363–382.
- Schelling, T.C. (1984) Self-command in practice, in policy, and in a theory of rational choice, *The American Economic Review*, **74** (2), in *Papers and Proceedings of the Ninety-Sixth Annual Meeting of the American Economic Association*, pp. 1–11.
- Sullivan, M. (2018) *Time Biases: A Theory of Rational Planning and Personal Persistence*, Oxford: Oxford University Press.