

# Transformative experience and the knowledge norms for action: Moss on Paul's challenge to decision theory

Richard Pettigrew  
Richard.Pettigrew@bris.ac.uk

November 22, 2017

Experience can change you in many different ways. To name just two: it can teach you new things, thereby changing your epistemic state; and it can lead you to change what you value and the extent to which you value it, thereby changing your conative state. If an experience changes you in the first way, L. A. Paul dubs it an *epistemically transformative experience* (henceforth, *ETE*). If it changes you in the second way, it is a *personally transformative experience* (Paul, 2014, 2015). Paul argues that the possibility of both sorts of experience poses serious and novel problems for the orthodox theory of rational choice, namely, expected utility theory.<sup>1</sup> In this paper, I will focus only on Paul's argument that the possibility of ETEs raises a challenge for expected utility theory — I will call her objection *the Utility Ignorance Objection*. In a pair of earlier papers, I responded to Paul's challenge (Pettigrew, 2015, 2016), and a number of other philosophers have responded in similar ways (Dougherty et al., 2015; Harman, 2015) — I will call the argument that we have each put forward *the Fine-Graining Response*. Paul (2014) has her own reply to this response, which we might call *the Authenticity Reply*. But Sarah Moss has offered an alternative reply to the Fine-Graining Response on Paul's behalf (Moss, ms) — we'll call it *the No Knowledge Reply*. This appeals to the knowledge norms for action, together with Moss' novel and intriguing account of probabilistic knowledge. In this paper, I consider Moss' reply and argue that it fails. I argue first that it fails as a reply made on Paul's behalf, since it forces us to abandon many of the features of Paul's challenge that make it distinctive and with which Paul herself is particularly concerned. Then I argue that it fails as a reply independent of its fidelity to Paul's intentions.

---

<sup>1</sup>While Paul targets expected utility theory with her objection, if it works, it affects the whole gamut of non-expected utility theories as well (Kahneman & Tversky, 1979; Quiggin, 1993; Buchak, 2013; Wakker, 2010).

Before we can state Paul’s challenge to decision theory, we have to make clear exactly which version of that theory she wishes to challenge (§1). Having done that, we can state Paul’s Utility Ignorance Objection to that theory (§2), the Fine-Graining Response (§3), and Paul’s Authenticity Reply (§4). Then, we introduce Moss’ No Knowledge Reply (§5). We argue that it fails both as a reply on behalf of Paul (§6) and as a reply on its own (§7).

## 1 What is decision ?

Paul addresses her Utility Ignorance Objection to a particular interpretation of decision theory, namely, the realist-deliberative interpretation. Let me introduce this now by drawing two distinctions: the realist/constructivist distinction and the deliberative/evaluative distinction.

First, the realist/constructivist distinction. Realist and constructivist understandings of decision theory agree on the building blocks of decision theory; they differ on which of these building blocks are more fundamental. On both, we have the following:

- a set of possible actions the agent might perform — call that set  $\mathcal{A}$ ;
- a set of possible states of the world — call it  $\mathcal{S}$ ;
- a preference ordering  $\preceq$  over the possible actions in  $\mathcal{A}$ ;
- a credence function  $P$  over combinations of actions from  $\mathcal{A}$  and states from  $\mathcal{S}$  — given  $a$  in  $\mathcal{A}$  and  $s$  in  $\mathcal{S}$ ,  $P(s||a)$  records how strongly the agent believes that the world is in state  $s$  under the supposition that they perform  $a$ ;<sup>2</sup> and
- a utility function  $U$  over combinations of actions from  $\mathcal{A}$  and states from  $\mathcal{S}$  — given  $a$  in  $\mathcal{A}$  and  $s$  in  $\mathcal{S}$ ,  $U(a \& s)$  records how strongly the agent values or desires or wants or endorses the outcome of performing action  $a$  when the world is in state  $s$ .

According to the constructivist understanding of decision theory, the preference ordering is primary — for some constructivists of a more behaviourist persuasion, only the preference ordering is psychologically real; for others, credence and utility functions are real too, but the preference ordering is more fundamental than both. Typically, constructivists then show that, providing that the preference ordering  $\preceq$ , set of states  $\mathcal{S}$ , and the set of possible actions  $\mathcal{A}$  satisfy particular structural constraints — the Savage (1954), Jeffrey (1983), or Joyce (1999) axioms, for instance — there

---

<sup>2</sup>If the supposition in question is indicative supposition, the resulting theory is *evidential decision theory* (Jeffrey, 1983); if it is subjunctive supposition, the result is *causal decision theory* (Joyce, 1999). The difference between them will not matter here.

is a credence function  $P$  and a utility function  $U$  such that the preference ordering places one action at least as high as another iff the expected utility of the first is at least as great as the expected utility of the second; that is,  $a_1 \preceq a_2$  iff  $\text{Exp}_P(U(a_1)) \leq \text{Exp}_P(U(a_2))$ , where  $\text{Exp}_P(U(a)) = \sum_{s \in \mathcal{S}} P(s|a)U(a \ \& \ s)$ . That is, the preference ordering, which is primary, can be represented by the credence function and the utility function, which are thereby secondary. Such a result is known as a *representation theorem*, and they are central to the constructivist understanding of decision theory. On the other hand, the realist understanding of decision theory says that the credence function and utility function are primary, while the preference ordering on actions given by their expected utility relative to that credence function and utility function is secondary. Thus, for the constructivist, we begin with the preference ordering and derive the credence function and utility function; for the realist, we begin with the credence function and utility function and derive the preference ordering. For both, faced with a decision between a range of possible actions, an agent is rationally required to choose an action that sits at the top of the preference ordering — that is, it is irrational to choose an action  $a$  in  $\mathcal{A}$  when there is an alternative action  $a^*$  in  $\mathcal{A}$  such that  $a \prec a^*$ . Paul’s challenge applies primarily to the realist understanding of decision theory.

Next, let’s turn to the deliberative/evaluative distinction. Deliberative and evaluative understandings of decision theory differ on which elements of a decision are relevant to its rationality. For those who favour a deliberative understanding, decision theory governs not only the choice that an agent makes in a given situation, but also the deliberation by which she comes to make that choice. In contrast, those who favour an evaluative understanding say that decision theory evaluates the choice only. Thus, for instance, suppose I must decide whether or not to take an umbrella when I leave my house. As it happens, I would maximise my expected utility by taking the umbrella — I think it’s pretty likely to rain, I hate getting wet, and it doesn’t much bother me to carry the umbrella. Now suppose that I do indeed end up taking the umbrella. But suppose that my reason for doing so was not that it would maximise my expected utility — it was not by calculating which action would maximise expected utility and then picking it that I reasoned to my conclusion. Rather, I chose the action I did simply using the rule *Always pick the action that involves approximating most closely the sartorial choices of Mary Poppins*. Then, according to the evaluative understanding of decision theory, I am fully rational, because I chose the option that maximises expected utility, while according to the deliberative understanding, I am not, because I did not deliberate correctly concerning my choice — my decision was not sensitive to the expected utility of the actions between which I had to choose. Paul’s challenge applies primarily to the deliberative understanding of decision theory.

## 2 Paul's Utility Ignorance Objection

What is an ETE? It is an experience that teaches you something that you couldn't come to know without having that experience. Thus, for Frank Jackson's scientist, Mary, who has lived her whole life in a monochrome black-and-white room, the experience of stepping outside and seeing the colour red for the first time is an ETE. However much Mary learned about the physical properties of red objects, she could not know what it is like to see red (Jackson, 1986). Similarly, for some people, becoming a parent for the first time is an ETE. However much they attended to the testimony of people who already have children, however much they read novels about parenting, they couldn't know what it was going to be like to be a parent until they became one themselves (Paul, 2014).

In Mary's case, what she learns from her ETE is a phenomenological fact — she learns what it is like to see red. In the case of the new parent, there is likely a phenomenological component to what they learn from the experience as well — they learn what it is like to feel a particular sort of bond with another person; and they learn what it is like to have responsibility for another life. But there may well be other components — the experience might teach you some moral facts, for instance. For Paul's objection, she needs only this: ETEs teach you something that you cannot learn any other way and that you need to know in order to know the utility that you assign to the outcomes of certain actions that are available to you.

For instance, suppose I must decide whether or not to apply to adopt a child and become a parent — this will serve as our illustrative example throughout the paper. If I choose to apply and my application is successful, I will become a parent. In order to calculate the expected value of choosing to apply, I must therefore know the utility I assign to the outcome on which I apply and my application is successful. But in order to know that, I need to know what it will be like to be a parent. And that, for some people, is something that they can know only once they become a parent. For such people, then, it seems that the ingredients that they require in order to calculate their expected utility for applying to adopt a child are not epistemically available to them. And thus they are barred from deliberating in the way that the realist-deliberative understanding of decision theory requires of them. They are unable to make the decision rationally.

Using the ingredients of decision theory introduced above, we can state the problem as follows: there are two actions between which I must choose — apply to adopt a child (*Apply*); don't apply (*Don't Apply*). And let's say that there are two states of the world — one in which I would be successful if I were to apply (*Succeed*), and one in which I would be unsuccessful (*Fail*). So the decision table is this:

	<i>Succeed</i>	<i>Fail</i>
<i>Apply</i>	??? $U(\textit{Succeed} \ \& \ \textit{Apply})$ ???	$U(\textit{Fail} \ \& \ \textit{Apply})$
<i>Don't Apply</i>	$U(\textit{Succeed} \ \& \ \textit{Don't Apply})$	$U(\textit{Fail} \ \& \ \textit{Don't Apply})$

According to the realist, to choose whether or not to apply, I must determine whether I prefer applying to not applying — that is, whether  $\textit{Apply} \preceq \textit{Don't Apply}$  or  $\textit{Apply} \sim \textit{Don't Apply}$  or  $\textit{Apply} \succeq \textit{Don't Apply}$ . To determine that, I must calculate the expected utility of those two actions relative to my credence function and my utility function. And to calculate that, I must know what my credence is in each of the two possible states of the world under the supposition of each of the possible actions — that is, I must know  $P(\textit{Succeed} | \textit{Apply})$  and  $P(\textit{Fail} | \textit{Apply})$ , as well as  $P(\textit{Succeed} | \textit{Don't Apply})$  and  $P(\textit{Fail} | \textit{Don't Apply})$ . And I must know my utilities for the different possible outcomes — that is, I must know  $U(\textit{Apply} \ \& \ \textit{Succeed})$  and  $U(\textit{Apply} \ \& \ \textit{Fail})$  and  $U(\textit{Don't Apply} \ \& \ \textit{Succeed})$  and  $U(\textit{Don't Apply} \ \& \ \textit{Fail})$ . The problem that Paul identifies is that it is impossible to know  $U(\textit{Apply} \ \& \ \textit{Succeed})$  prior to making the decision and becoming a parent; and thus it is impossible to deliberate about the decision in the way that the realist-deliberative understanding of decision theory requires. Paul concludes that there is no rational way to make such decisions. This is the Utility Ignorance Objection to the realist-deliberative understanding of decision theory.

Before we move on to consider how we might respond to this objection, let us pause a moment to consider its scope. First, note that the challenge targets only the realist understanding of decision theory, not the constructivist. For the constructivist, my credence and utility functions are determined by my preference ordering. Thus, to know them I need only know my preference ordering. And for many constructivists I can know that simply by observing how I choose between given sets of actions. Paul's challenge applies only when we take the utility function to determine, at least in part, our preference ordering, as the realist does. Second, note that the challenge targets only the deliberative understanding of decision theory, not the evaluative. On the realist-evaluative understanding, I do not need to know my credences or my utility function in order to be rational. On this understanding, in order to be rational, I need only choose the action that *in fact* maximises expected utility; I need not choose it *because* it maximises expected utility. Thus, Paul's argument has no bite for the evaluative understanding.

Now, we might try to extend Paul's argument so that it does apply to the realist-evaluative understanding. To do that, we need to argue not only that I do not *know*  $U(\textit{Apply}, \textit{Succeed})$  prior to my choice between *Apply* and *Don't Apply*, but indeed that  $U(\textit{Apply}, \textit{Succeed})$  is not even *determined* prior to that choice. If that is the case, then there is no way to make the choice rationally, even according to the realist-evaluative understanding. Similarly, if my preference between *Apply* and *Don't Apply* is not even defined

prior to my choice between them, then I cannot make that choice rationally, even according to the constructivist. But of course this is not what Paul's argument establishes. The argument is explicitly epistemic.

### 3 The Fine-Graining Response

There is a natural response to Paul's challenge. Expected utility theory is designed to deal with decisions made in the face of uncertainty. Usually that uncertainty concerns the way the world is beyond or outside of the agent. For instance, suppose I'm uncertain whether my adoption application would be successful if I were to apply. Then, when I'm making my decision, I ensure that the set of possible states of the world includes one in which my application succeeds and one in which it fails. I then quantify my uncertainty concerning these two possibilities in my credence function, and I use that to calculate my expected utility — perhaps I know that only 12% of adoption applications succeed, and I set my credence that mine will succeed to 0.12 in line with that. However, there is no reason why the uncertainty quantified by my credence function should concern only the way the world is beyond me. What Paul's argument shows is that I am uncertain not only about the world, but also about the utility that I assign to becoming a parent; I am uncertain not only about whether *Succeed* or *Fail* is true, but also about the value  $U(\text{Apply}, \text{Succeed})$ . Thus, just as I ensured that my decision problem includes possible states of the world at which I succeed in my application and possible states where I fail, similarly I should respond to Paul's challenge by ensuring that my decision problem includes possible states of the world at which I become a parent and value it greatly, possible states at which I become a parent and value it a moderate amount, states at which I become a parent and value it very little, and so on. Having done this, I should quantify my uncertainty concerning the utility I assign to being a parent in my credence function, and use that to calculate my expected utility as before.

More precisely, and simplifying greatly, suppose the possible utility values that I might assign to being a parent are  $-12$ ,  $3$ , and  $10$ . Then, while my original set of possible states of the world is  $\mathcal{S} = \{\text{Succeed}, \text{Fail}\}$ , my new expanded set of possible states of the world is

$$\mathcal{S}^* = \{ \text{Succeed \& utility of being a parent is } -12, \\ \text{Succeed \& utility of being a parent is } 3, \\ \text{Succeed \& utility of being a parent is } 10, \\ \text{Fail} \}.$$

Now, recall the problem that Paul identified. Given the original way of setting up the decision problem, in order to deliberate rationally between *Apply* and *Don't Apply*, I need to know the utilities I assign to each possi-

ble outcome of each of the possible actions. In particular, I need to know  $U(\text{Apply}, \text{Succeed})$ . But I can't know that until I make the decision and become a parent. However, on the new formulation of the decision problem, with the expanded set of states  $S^*$ , I do know the utilities I assign to each possible outcome of each of the possible actions. For I know that:

- $U(\text{Apply}, \text{Succeed} \ \& \ \text{utility of being a parent is } -12) = -12,$
- $U(\text{Apply}, \text{Succeed} \ \& \ \text{utility of being a parent is } 3) = 3,$
- $U(\text{Apply}, \text{Succeed} \ \& \ \text{utility of being a parent is } 10) = 10,$

Next, I quantify my uncertainty in these new possible states to give:

$$\begin{aligned} P(\text{Succeed} \ \& \ \text{utility of being a parent is } -12|a), \\ P(\text{Succeed} \ \& \ \text{utility of being a parent is } 3|a), \\ P(\text{Succeed} \ \& \ \text{utility of being a parent is } 10|a), \\ P(\text{Fail}|a), \end{aligned}$$

where  $a$  is either *Apply* or *Don't Apply*. And, given this, I can calculate my expected utility and discharge the obligations of rationality imposed by the realist-deliberative understanding of decision theory. Paul's Utility Ignorance Objection, it seems, is answered. Call this the Fine-Graining Response, since it involves expanding, or fine-graining, the set of possible states of the world.

## 4 Paul's Authenticity Reply

Paul herself is not satisfied with this response. She allows that I can expand the set of possible states of the world in the way described. And she allows that I can form credences in those different states of the world. But she worries about the sort of evidence on which I might base those credences.

Let's start with an ordinary decision that does not involve an ETE. Suppose I am deciding whether to have chocolate ice cream or strawberry ice cream. I have tasted both in the past, so I know what both experiences will be like — neither experience would be transformative. As a result, when I come to make my decision, I know the utility I assign to the outcome in which I eat chocolate ice cream. I know it by imaginatively projecting myself forward into the situation in which I am eating chocolate ice cream. And I can do this because I have tasted chocolate ice cream in the past. And similarly for the utility I assign to the outcome in which I eat strawberry ice cream. I know what it is, and I know it because I've tasted strawberry ice cream in the past and so I can imaginatively project myself forward into the situation in which I'm eating it.

When I consider the utility I assign to becoming a parent, I can't imaginatively project in this way, since I'm not a parent and becoming a parent is

an ETE. As described above, I respond to this epistemic barrier by expanding the set of possible states of the world I consider in my decision problem. I expand them so that they are fine-grained enough that each specifies my utility for becoming a parent at that world; and my credences in these different possible states quantify my uncertainty over them. But how do I set those credences? I cannot do anything akin to imaginatively projecting myself into the situation of being a parent, as I did with the chocolate ice cream, because becoming a parent is an ETE. What can I do instead?

Well, the natural thing to do is to seek out the testimony of people who have already undergone that transformative experience.<sup>3</sup> Perhaps I cannot discover from them exactly what it is like to be a parent — since it’s an ETE, the only way to learn what it’s like is to undergo the experience. But perhaps I can learn from them how much they value the experience. And after all, that’s all that I need to know in order to make my decision rationally, according to the realist-deliberative understanding of decision theory — expected utility theory doesn’t require that you know what an outcome will be like; it requires only that you know how much you value it and thus how much it contributes to the expected utility calculation. However, as we all know, different people value being a parent differently. For some, it is an experience of greater value than all other experiences they have in their life. For others, it is a positive experience, but doesn’t surpass the value of reciprocated romantic love, or extremely close friendships, or succeeding in a career, or helping others. And for yet others, it is a negative experience, one that they would rather not have had. Simplifying greatly once again, let’s assume that all parents fall into these three groups: members of the first assign 10 utiles to the outcome in which they become a parent; members of the second assign 3; and members of the third assign -12. And let’s assume that 10% fall into the first group; 60% into the second; and 30% into the third. Now, suppose that I learn this statistical fact by attending to the testimony of parents. Then I might set my credences as follows (where we assume for convenience that I am certain that my adoption application will be successful, so  $P(\text{Succeed}||\text{Apply}) = 1$ ):

- $P(\text{Succeed} \ \& \ \text{utility of being a parent is } -12 \ || \ \text{Apply}) = 0.3,$
- $P(\text{Succeed} \ \& \ \text{utility of being a parent is } 3 \ || \ \text{Apply}) = 0.6,$
- $P(\text{Succeed} \ \& \ \text{utility of being a parent is } 10 \ || \ \text{Apply}) = 0.1,$

With these in hand, I can then calculate the expected utility of *Apply* and *Don’t Apply*, I can compare them, and I can make the choice between them in the way that the realist-deliberative decision theorist requires.

---

<sup>3</sup>See (Dougherty et al., 2015) for two further ways in which I might set these credences. I focus on testimonial evidence here since it is the sort of evidence that Moss considers.

However, Paul claims that if I choose in this way then my decision is badly flawed. She holds that an agent who made the decision to become a parent in this way would be “alienated” from that decision; the choice thus made would be “inauthentic”:

A [...] problem with leaving your subjective perspective out of your decisions connects to the Sartrean point that making choices authentically and responsibly requires you to make them from your first personal perspective. A way to put this is that if we eliminate the first personal perspective from our choice, we give up on authentically owning the decision, because we give up on making the decisions for ourselves. We give up our authenticity if we don’t take our own reasons, values, and motives into account when we choose. To be forced to give up the first person perspective in order to be rational would mean that we were forced to engage in a form of self-denial in order to be rational agents. We would face a future determined by Big Data or Big Morality rather than by personal deliberation and authentic choice. (Paul, 2014, 130)

For Paul, then, the problem lies in the way that I set my credences in the fine-grained states of the world. I set my credences concerning my own utilities by deferring to statistical facts about how others assign different utilities. My evidence does not sufficiently concern *my* utilities; and thus I am alienated from any decision based on the credences that I form in response to that evidence. I am like the agent who makes a moral decision by deferring to societal norms or the value judgements of the majority group, rather than making those decisions herself. Paul contrasts this statistical method of forming opinions about my own utilities with the method described above in the case of the chocolate and strawberry ice cream, where I imaginatively project myself into the situation in which I have the experience based on my own memory of previous similar experiences. In those cases, the opinions formed do not give rise to the same sort of alienation and inauthenticity, since they are connected in the right way to my own utilities. They are more akin to the agent who makes the moral decision for themselves.

I have responded to Paul’s concern elsewhere, where I argue that there is a crucial difference between these cases (Pettigrew, 2015, 770). When I set my credences concerning my own utilities by appealing to the statistical evidence concerning the utilities of others, I do so because I think that this statistical evidence tells me something about *my own utility*; it is good evidence concerning *my own utilities*. In contrast, when I defer to societal norms to make a moral decision, I do so not because I think that those norms tell me anything about my own values; I do not think they provide good evidence concerning what I think is the correct moral action. I do so

because I can't decide what I think is the correct moral action, or I do not have the courage to follow my own moral compass. I mention Paul's Authenticity Reply here partly for the sake of completeness, but also because Moss' No Knowledge Reply to the Fine-Graining Response also argues that the problem with such decisions lies in the nature of the evidence on the basis of which I form my credences about my utilities. Let's turn to Moss' reply now.

## 5 Moss' No Knowledge Reply

Suppose I set my credences in *Succeed & utility of being a parent is -12*, etc., as above. That is, I set them on the basis of statistical evidence concerning the utilities that existing parents assign to being a parent. For Paul, the problem is that such evidence does not sufficiently concern my utilities in particular; it is too much concerned with the utilities of other people. For Moss, the problem with those credences is not that they are not sufficiently concerned with me, or at least that is not the primary problem. Rather, the problem is that those credences do not constitute knowledge, and rational decisions must be based on credences that constitute knowledge (Moss, ms, Section 9.5).

To those unfamiliar with Moss' work, it might sound as if she is making a category mistake. Credences, you might think, are simply not the sort of thing that can constitute knowledge. Full beliefs can — if I believe that it's raining, then that belief might count as knowledge. But credences, or partial beliefs, cannot — if I have credence 0.6 that it's raining, then it makes no more sense to say that that credence counts as knowledge than it does to say that a colourless idea sleeps furiously. Or so you might think. But Moss denies this (Moss, 2013, ms). Let's see why.

First, it is worth saying what Moss takes credences to be. Suppose I say that I'm 50% confident that Kenny is in Hamburg. On the standard interpretation, this means that I have a precise graded attitude — a credence — towards the standard, non-probabilistic content *Kenny is in Hamburg*, where the latter might be represented by a set of possible worlds. In particular, I have a 0.5 credence in that non-probabilistic content. For Moss, in contrast, a credence is not a graded attitude towards a standard propositional content; rather, it is a categorical attitude towards what she calls a *probabilistic content*. For instance, to say that I'm 50% confident that Kenny is in Hamburg is to say that I have a categorical attitude — in fact, a belief — towards the probabilistic content *Kenny is 50% likely to be in Hamburg*.

What are these probabilistic contents? Well, just as a standard propositional content, such as *Kenny is in Hamburg*, can be represented by a set of possible worlds, so a Mossian probabilistic content, such as *Kenny is 50% likely to be in Hamburg*, is represented by a set of probability spaces, where a

probability space is a set of possible worlds together with a probability distribution defined over those worlds. Thus, the probabilistic content *Kenny is 50% likely to be in Hamburg* is represented by the set of those probability spaces in which the probability distribution assigns 50% to the proposition *Kenny is in Hamburg* — that is, the set  $\{P : P(\text{Kenny is in Hamburg}) = 0.5\}$ .

Another example: Suppose I say that I'm more confident than not that Kenny is in Hamburg. On the standard interpretation, this means that I have an imprecise graded attitude towards the propositional content *Kenny is in Hamburg*. Imprecise graded attitudes are also represented by sets of probability spaces — these are usually called *representors*. In this case, my imprecise graded attitude is represented by the set of those probability spaces in which the probability distribution assigns more than 50% to the proposition *Kenny is in Hamburg* — that is, the set  $\{P : P(\text{Kenny is in Hamburg}) > 0.5\}$ . That set is my representor. For Moss, in contrast, I do not have a graded attitude towards the propositional content *Kenny is in Hamburg*, but rather a categorical attitude towards the probabilistic content *Kenny is more likely than not to be in Hamburg*. The probabilistic content towards which I have that categorical attitude is in fact represented by the same set of probability spaces that is used to represent the imprecise graded attitude that is usually attributed to me — that is, my representor,  $\{P : P(\text{Kenny is in Hamburg}) > 0.5\}$ .

Now, citing a large body of examples, Moss argues that we often say that, just as beliefs in standard, non-probabilistic contents — viz., propositions — can count as knowledge, so can beliefs in probabilistic contents — viz., the contents represented by sets of propositions. For instance, I might say that Patricia knows that Kenny is 50% likely to be in Hamburg, or that Jason knows that Kenny is more likely than not to be in Hamburg.

As well as citing intuitive examples in which we ascribe probabilistic knowledge, Moss also gives examples that show that there are distinctions between categorical beliefs in probabilistic contents that are analogous to the distinctions that we mark between different categorical beliefs in propositions by categorising one as merely justified and the other as knowledge. For instance, suppose that I know that the objective chance of this coin landing heads is 60%. And my credence that it will land heads is 0.6 — that is, in Moss' framework, I believe that the coin is 60% likely to land heads. Next, suppose that you also set your credence in heads to 0.6 — that is, you also believe the coin is 60% likely to land heads. But you set your credence in this way not because you know the objective chance, but because you know that Sarah's credence in heads is 0.6 and you have good reason to take Sarah to be an expert on the bias of coins. However, while you are right that Sarah is generally expert on such matters, in this case she hasn't actually inspected the coin and instead just plucked a number from thin air. In such a case, it seems that, while both of us have justified credences that are correct in a certain sense, yours is merely justified, while

mine counts as knowledge.

Moss furnishes us with a splendidly detailed account of probabilistic knowledge, which includes a Bayesian expressivist semantics for probabilistic knowledge ascriptions as well as an account of the factivity, safety, and sensitivity conditions on probabilistic knowledge. But her No Knowledge Reply to the Fine-Graining Response does not depend on the more sophisticated or radical elements of her account. Rather, it depends on just three claims about probabilistic knowledge.

The first, we have met already: it is the claim that credences — and, more generally, beliefs in probabilistic contents — can count as knowledge, just as beliefs in non-probabilistic contents can.

The second claim concerns a certain sort of case in which the credences you form don't count as knowledge. Suppose we meet. Noting that I am a living human being, and knowing that about 0.7% of living human beings will die in the next year, you form a credence of 0.007 that I will die in the next year. Then, for Moss, your credence does not count as knowledge. The problem is that you cannot rule out relevant alternative reference classes to which I belong and amongst which the frequency of death within the next year is quite different. For instance, you know that I am 35 years old. And you can't rule out that the likelihood of death amongst living 35 year olds is quite different from the likelihood amongst all human beings. You know that I am male. And you can't rule out that the likelihood of death amongst living males is different from the likelihood amongst all human beings. And so on. You believe that it's 0.7% likely that I will die in the coming year, but you can't rule out that my death is  $X\%$  likely, for a range of alternative values of  $X$ . Moss likens the case to Goldman's fake barn scenario (Goldman, 1976). I am travelling through Fake Barn County, and I stop in front of a wooden structure that looks like a barn. I form the belief that the structure in front of me is a barn because that's what it looks like. But my visual experience cannot distinguish a barn from a barn facade. So I cannot rule out the alternative possibility that the structure is a barn facade. And this alternative is relevant because Fake Barn County lives up to its name: it's full of fake barns. Therefore, my belief cannot count as knowledge. Similarly, since you cannot rule out certain alternative reference classes amongst which my likelihood of death within the next year is quite different from 0.7%, your credence of 0.007 that I will die in the next year cannot count as knowledge. Or so Moss says.

Now, recall our response outlined above to Paul's Utility Ignorance Objection to decision theory. Since I cannot know the utility I assign to being a parent, I expanded the set of possible states of the world so that, in each, my utility is specified; and then I quantified my uncertainty concerning these different utilities in my credences. Since I could not set those credences by imaginatively projecting myself into the position of being a parent, I had to set them by appealing to the statistical evidence concerning the utilities

that existing parents assigned to being parents. Since the evidence for my credences is statistical, if it is to count as knowledge, I must be able to rule out relevant alternative reference classes to which I belong on which the statistics are quite different. For instance, suppose I set my credences in the different possible utilities by appealing to the statistics amongst *all* existing parents. Then there are certainly relevant alternative reference classes that I should consider: the class of all male parents; the class of all gay male parents; the class of adoptive parents; the class of all parents with family and social support network similar to mine; and so on. Given the evidence on which I based my credences, I cannot rule out the possibility that the distribution of the three candidate utilities for being a parent is different in these reference classes from the distribution in the reference class on which I based my credences. Thus, according to Moss, my credences cannot count as knowledge.

Finally, the third claim upon which Moss bases her No Knowledge Reply to the Fine-Graining Response is a conjunction of a probabilistic knowledge norm for reasons and a probabilistic knowledge norm for decision — together, we refer to these as the *Probabilistic Knowledge Norms for Action*, following Moss.

**Probabilistic Knowledge Norm for Reasons** Your credal state can only provide a reason for a particular choice if it counts as knowledge.

**Probabilistic Knowledge Norm for Decisions** Suppose the strongest probabilistic content you know is represented by a set  $\mathbf{P}$  of probability functions; and suppose you are faced with a choice between a range of options. It is permissible for you to choose a particular option iff that option is permissible, according to the correct decision theory for imprecise credences, for an agent whose imprecise credal state is represented by  $\mathbf{P}$ .

For instance, suppose you must choose whether to take an umbrella with you when you leave the house. The strongest proposition you know is represented by the set of probability spaces,  $\mathbf{P} = \{c : 0.4 < P(\text{Rain}) < 0.9\}$ . If rain is 90% likely, then taking the umbrella maximises expected utility; if it is only 40% likely, then leaving the umbrella maximises expected utility. Now imagine an agent whose credal state is represented by  $\mathbf{P}$  — in the language introduced above,  $\mathbf{P}$  is her representor.<sup>4</sup> Which actions are permissible for this agent? According to some decision theories for imprecise credences, an action is permissible iff it maximises expected utility relative to *at least one member of the representor*. We might call these *liberal* decision

<sup>4</sup>For more on the correct decision theory for imprecise credences, see (Seidenfeld, 2004; Seidenfeld et al., 2010; Elga, 2010; Joyce, 2010; Rinard, 2015).

theories, since they make many actions permissible. On this decision theory, it is permissible to take the umbrella and permissible to leave it. Thus, according to the Probabilistic Knowledge Norm for Decisions, both actions are also permissible. According to other decision theories, an action is permissible iff it maximises expected utility relative to *all members of the representor*. We might call these *conservative* decision theories, since they make few actions permissible. On this decision theory, neither taking nor leaving the umbrella is permissible for the agent with representor **P**, and thus, according to the Probabilistic Knowledge Norm for Decisions, neither is permissible for me.

Thus, putting together the various components of Moss' No Knowledge Reply, we have:

- (i) the only precise credences I could form concerning the utility I assign to being a parent do not count as knowledge, because my statistical evidence doesn't allow me to rule out alternative reference classes that are made salient, or relevant, by the high stakes decision I wish to make based on those credences;
- (ii) by the Probabilistic Knowledge Norm for Reasons, these credences can therefore not provide a reason for me to act in any particular way, so that if I choose to do whatever maximises expected utility relative to those credences, my reason for choosing in that way cannot be that the choice maximised expected utility for me, since that invokes my credences as a reason;
- (iii) by the Probabilistic Knowledge Norm for Decisions, I am not necessarily required to choose the action that maximises expected utility relative to those credences — they do not correspond to the strongest probabilistic content I know, and thus what is permissible for me is not determined by maximising expected utility with respect to them.

What, then, am I required to do? That depends on what my statistical evidence allows me to know, and what the correct decision theory is for imprecise credences. As I mentioned already, there are many candidate theories, including the liberal and conservative versions described above. And on the question of what my statistical evidence allows me to know, we will have more to say below.

## **6 Assessing Moss' No Knowledge Reply: the Paulian view**

We have now seen Paul's Utility Ignorance Objection to decision theory, the Fine-Graining Response, Paul's Authenticity Reply, and Moss' No Knowledge Reply. Given this, we can ask two questions: Does Moss' reply work

from Paul's point of view? Does Moss' reply work independently of Paul's point of view? Paul emphasises four important features of her objection. As we will see, Moss' reply to the Fine-Graining Response preserves two of those to some extent and two not at all. We begin with those it doesn't preserve.

First, Paul claims that the challenge to decision theory raised by ETEs is unique to those experiences. Whatever problem they raise, it is not raised by any other sort of phenomenon. And yet that isn't true on Moss' interpretation. Consider the doctor who must choose a treatment for her patient. She has the following statistical evidence: in 98% of trial cases, the treatment cures the illness; in 2% of trial cases, the patient deteriorates severely. She sets her credences in line with that. The illness is serious, so this is a high stakes decision. Thus, other reference classes are relevant, and the doctor's evidence cannot rule out that the frequency of successful treatment is very different in those. So, by Moss' lights, the doctor's credence of 0.98 that the treatment will succeed and 0.2 that it will fail do not count as knowledge and so cannot provide a reason for action. Now, you might consider that the wrong conclusion or the right one — you might think, for instance, that the doctor's credences can provide reason for action, even if the doctor would prefer to have better evidence. But that is not the issue here. The issue is only that this other decision faces exactly the same problems that, for Moss, any decision faces that involves ETEs. That is, ETEs do not pose any new or distinctive problem for decision theory. And thus, on Moss' account, we lose this crucial feature of Paul's account.

The second distinctive feature of Paul's account is that, in decisions that involves ETEs, the problem is first-personal. When I am choosing whether or not to become a parent, the problem arises, according to Paul, because I am trying to make a decision for myself about my own future and yet I cannot access a part of my self that is crucial to the decision, namely, my utilities. This is why Paul turns to concepts like *alienation* and *authenticity* to account for the phenomenon: they apply to first-personal choices in a way that they don't to third-personal ones. However, as the example of the doctor from above shows, there is nothing distinctively first-personal in Moss' diagnosis of the problem with decisions that involve ETEs — the problem arises just as acutely for a doctor making a major decision for a patient as it does for me when I try to choose whether or not to adopt.

The first feature of Paul's account that Moss' No Knowledge Reply does preserve and explain, though for quite different reasons, is the importance of what is at stake in the decision that we wish to use our credences to make. As Paul and Moss both acknowledge, there are trivial ETEs and important ones. When I choose whether to spread Vegemite or Marmite on my toast — having tried neither — I am choosing which ETE to have. But neither thinks that this poses a problem for decision making in the way that choosing to become a parent does. Both think it is quite acceptable to

use statistical evidence about the utilities that others assign to eating those two condiments as reasons I might cite when making my decision. Paul's explanation: only in significant life decisions do alienation and inauthenticity threaten. Moss' explanation: in low stakes cases, there are no alternative reference classes that are relevant, and so my credences will constitute knowledge even if my evidence cannot rule out any alternative reference classes. Different explanations, but both agree that stakes matter.

The second feature of Paul's account that Moss' reply preserves, though again for quite different reasons, is the attitude to decision theory. It is important to note that neither Paul nor Moss wish to abandon the machinery of decision theory in the face of the Utility Ignorance Objection; neither wishes to reject expected utility theory. Rather, in the case of significant life decisions that might give rise to ETEs, they advocate changing the decision problem that we feed into that decision theory. For instance, on the Fine-Graining Response, when I am deciding whether or not to adopt a child, I formulate the following decision problem:

- the possible acts are
  - *Apply,*
  - *Don't Apply;*
- the possible states are
  - *Succeed & utility of being a parent is –12,*
  - *Succeed & utility of being a parent is 3,*
  - *Succeed & utility of being a parent is 10,*
  - *Fail;*
- the doxastic states are my precise or imprecise credences over those states, on the supposition of those acts;
- the conative states are my utilities over the conjunctions of acts and states, which encode the overall value I attach to these conjunctions, incorporating the quality of the phenomenal experience they give me, the moral and aesthetic values they boast, and so on.

I then feed this decision problem into the machinery of decision theory, which then tells me which of the possible acts are permitted by rationality and which are not.

For Paul, the new decision problem that we feed into the machinery of decision theory is this:

- the possible acts are
  - *Apply,*

– *Don't Apply,*

as before;

- the possible states are
  - *Succeed,*
  - *Fail;*
- the doxastic states are my precise or imprecise credences over the states, on the supposition of the acts;
- the conative states are my utilities over the conjunctions of acts and states, but instead of encoding the overall value I attach to these conjunctions, which Paul has shown we cannot access prior to making the decision, they encode *only the value I assign to the revelatory experiences involved in those conjunctions.*

Thus, the conative state specified in the decision problem is different from that in the orthodox version, while the doxastic state remains the same.

In contrast, for Moss, the new decision problem is this:

- the possible acts are
  - *Apply,*
  - *Don't Apply;*
- the possible states are
  - *Succeed & utility of being a parent is –12,*
  - *Succeed & utility of being a parent is 3,*
  - *Succeed & utility of being a parent is 10,*
  - *Fail;*
- the doxastic states are not my precise or imprecise credences over the states, but rather *the strongest imprecise states that count as knowledge for me;*
- the conative states are my utilities over the conjunctions of acts and states, which encode the overall value I attach to these conjunctions, as in the orthodox approach.

Thus, the doxastic state specified in the decision problem is different from that in the orthodox version, while the conative state remains the same.

So, again, Paul and Moss agree — the orthodox decision problem should be replaced. But they agree for different reasons — Paul thinks that the conative state should be specified differently, while Moss thinks the doxastic state should be specified differently.

## 7 Assessing Moss' No Knowledge Reply: the independent view

In this section, we continue to consider Moss' No Knowledge Reply to the Fine-Graining Response to Paul's Utility Ignorance Objection to orthodox decision theory. But this time we consider it independently of its relationship to Paul's own reply to that response to her objection. We can read Moss' No Knowledge Reply in one of two ways: on the one hand, granted the possibility of probabilistic knowledge and the accompanying probabilistic versions of the knowledge norms for action — Moss' Probabilistic Knowledge Norm for Reasons and Probabilistic Knowledge Norm for Decisions — we can read it as trying to establish that the Fine-Graining Response is wrong; on the other hand, granted that the Fine-Graining Response is wrong, the need to appeal to probabilistic knowledge to explain why it is wrong is supposed to furnish us with an argument in favour of probabilistic knowledge, its possibility and its use as a concept in epistemology.

The first concern I wish to raise concerns the second reading. I will argue that a notion of probabilistic *knowledge* is not, in fact, required in order to explain the problem with decisions involving ETEs in the way Moss wishes to. The explanation can be given better, in fact, using only the familiar notion of probabilistic *justification*. The central point is this: the feature of first-personal utility credences based on statistical evidence that prevents them from counting as knowledge on Moss' account also prevents them from counting as justified. In the Fine-Graining Response outlined in §3 above, I have credence 0.3 in *Succeed & utility of being a parent is -12*, 0.6 in *Succeed & utility of being a parent is 3*, and 0.1 in *Succeed & utility of being a parent is 10*. I base these credences on my statistical evidence that 30% of parents assign utility -12 to being a parent, 60% assign utility 3, and 10% assign utility 10. Moss claims that these credences do not count as knowledge. I claim that, if they don't, they also don't count as justified.

Moss claims that these credences don't count as knowledge because my evidence doesn't allow me to rule out alternative reference classes that are rendered relevant by the high stakes of the decision I am making. I claim that they don't count as justified for the same reason. After all, the ability to rule out relevant alternatives is important for justification too. Suppose Charlie and Craig are identical twins. I know this; I've known them for years. I also know that I can't tell them apart reliably. I see Craig in the supermarket and I form the belief that Craig is in front of me. Now, while true, my belief does not count as knowledge because I can't rule out the relevant alternative possibility that it is Charlie in front of me, not Craig. But equally my inability to rule out this possibility also renders my belief unjustified. In general, if I believe  $p$  and there is an alternative possibility to

$p$  such that (i) I'm aware of it, (ii) I'm aware that it's relevant, and (iii) I can't rule it out, then my belief in  $p$  is not justified. The cases in which my inability to rule out an alternative precludes knowledge but not justification are those where either I am not aware of the possibility or not aware that it is relevant. For instance, in Goldman's Fake Barn County example, either I am not aware of the possibility of barn facades — perhaps I've never heard of such a thing — or, if I am aware of that possibility, I am not aware that it is relevant — because I don't know that I am in Fake Barn County. Thus, while I might be justified in believing that the structure in front of me is a barn, my belief doesn't count as knowledge. However, as soon as I learn about the possibility of barn facades and learn that I'm currently in Fake Barn County, my belief is neither justified nor knowledge. And the same goes for my credences about my utilities in the case of ETEs. Almost whatever statistical evidence I have about my utilities for becoming a parent, there is some relevant alternative reference class in which there are different frequencies for the various possible utility assignments such that (i) I'm aware of that reference class, (ii) I'm aware it's relevant, and (iii) I can't rule it out. Thus, any precise credence that I assign on the basis of that statistical evidence is not justified.

Thus, it seems to me that Moss' diagnosis of the problem with the Fine-Graining Response is wrong. The problem is not that the credences based on statistical evidence are not *knowledge*, it's that they're not *justified*. If that's right, then the argument in favour of the possibility of probabilistic knowledge that Moss bases on that diagnosis fails.

But this seems a Pyrrhic victory. If I am right, surely this only makes the problem worse for the Fine-Graining Response itself. After all, the possibility of probabilistic knowledge and the putative norms that link it with reasons and decisions are controversial, whereas the possibility of probabilistic justification and the norms that link it with reasons and decisions are not. I think most decision theorists would agree that, while there is sense in which an agent with unjustified credences should maximise expected utility with respect to those credences, that agent will nonetheless not be fully rational. Thus, we seem to be left with a stronger reply to the Fine-Graining Response than we had before: we might call it the *No Justification Reply*.

But this is too quick. All that the considerations so far have shown is that, if I take a single statistical fact based on the distribution of utilities amongst people in a single reference class, and set my credences about my own utilities exactly in line with that, without considering anything else, then those credences will typically neither be knowledge nor justified. But there are other, better ways to respond to statistical evidence, and these can give justified credal states that can then be used to make our ETE decisions.

For instance, suppose I have the statistical evidence from above: 10% of all parents assign 10 utiles to being a parent, 60% assign 3 utiles, and 30% assign -12. But I also realise that I have properties that I share with some

but not all parents: I enjoy spending time with my nieces and nephew; and I am a moderately anxious person. Let's suppose I think that the latter is the only property I have that affects the utilities I assign to being a parent. That is, I think that the distribution of utilities in the reference class of people who enjoy being around children is much the same as the distribution of utilities in the reference class of all parents, but the distribution amongst the reference class of moderately anxious people is quite different from the distribution in the class of all parents. And let's suppose that this belief is justified by my background evidence. Now, I don't know exactly what the latter distribution is, since that isn't included in my body of statistical evidence, but I have credences in the various possible distributions that are based on my background evidence. Let's assume again that those credences are justified by my background evidence. I then use these credences, together with my statistical evidence concerning the distribution of utilities in the reference class of all parents, to set my credences concerning my own utilities for being a parent. The resulting credences will be justified.

Now notice: these credences will be justified not because I've *ruled out* the alternative distributions of utilities amongst the alternative reference classes, but rather because I've *incorporated my uncertainty* about those different distributions into my new credences concerning my utilities for parenting. And indeed that is the natural thing to do in the probabilistic setting. For many Bayesian epistemologists, nothing that is possible is ever completely ruled out; we just assign to it very low credence. This is the so-called *Regularity Principle*, and there are various versions determined by the various different notions of possibility (Shimony, 1955; Stalnaker, 1970; Lewis, 1980; Jeffrey, 1992). If the Regularity Principle is true, it is too demanding to require of an agent with probabilistic attitudes that they rule out alternative possibilities before they can know anything. Rather, we might say: in order for a probabilistic attitude to be justified, the agent must have *considered* all relevant alternative possibilities and must have determined their attitude by *incorporating their attitudes towards those possibilities*. And we can do that in the case of credences concerning ETEs, even when those credences are based on statistical evidence, as we can see from the example of my adoption decision described above.

Now, I imagine that Moss might reply: while such credences might be justified, they will rarely count as knowledge. In order to count as knowledge, she might say, I must not only consider the properties I have *that I think might* affect the utility I assign to being a parent, and incorporate into my credences concerning that utility my uncertainty about the distribution of utilities for being parent amongst the reference classes defined by those properties; I must also consider the properties I have *that will in fact* affect that utility, and incorporate my uncertainty about the distribution of utilities for being a parent amongst the corresponding reference classes.

Failing to consider those other properties might not preclude justification — I might be perfectly justified in not having considered those properties, and indeed justified in not even being aware of them. But it does preclude knowledge. Thus, just as I am perfectly justified in ignoring the possibility that the structure in front of me is a fake barn, but will be unable to know various propositions if that possibility is relevant in my situation, similarly, I might be justified in not considering various reference classes and the distribution of utilities within them, but nonetheless unable to know various probabilistic content if those reference classes are relevant in my situation. And thus, Moss might claim, by the Probabilistic Knowledge Norms for Action, the justified credences that I formed by incorporating my uncertainty about distributions amongst alternative reference classes cannot be used in rational decision making in the usual way.

The problem with this claim is that it asks too much of us. If, in order to know a probabilistic content concerning an event in a high stakes situation, you must have considered all of the causal factors that contribute to it being likely to a certain degree, there will be almost no probabilistic contents concerning complex physical phenomena that we'll know. In a high stakes situation, I'll never know that it's at least 50% likely to rain in the next ten minutes, even if it is at least 50% likely to rain in the next ten minutes, since I simply don't know all of the causal factors that contribute to that — and indeed knowing those factors is beyond the capabilities of nearly everyone. There are many situations where, through no fault of our own, we just do not have the evidence that would be required to have credal states that count as knowledge. And this is not peculiar to credences concerning utilities for ETEs, nor even to credences based on statistical evidence.

Now, Moss might reply again: yes, it's difficult to obtain probabilistic knowledge; and perhaps we rarely do; and it's true that people shouldn't be held culpable if they violate the Probabilistic Norms of Actions; but that doesn't mean that we shouldn't strive to satisfy them, and it doesn't mean that the norms are not true. On this reply, Moss considers the Probabilistic Norms of Action as analogous to the so-called Truth Norm in epistemology, which says that we should believe only truths. Certainly, no-one thinks that those who believe falsehoods are always culpable. But nonetheless the Truth Norm specifies an ideal for which we should strive; it specifies the goal at which belief aims; and it gives us a way of assigning epistemic value to beliefs by measuring how far they fall short of achieving that ideal. Perhaps that is also the way to understand the Probabilistic Knowledge Norms for Action. They tell us the ideal towards which our actions should strive; and they give a way of measuring how well an action has been performed by measuring how far it falls short of the ideal.

But that can't be right. Consider the following Non-Probabilistic Knowledge Norm for Reasons. It says that a proposition  $p$  can count as your reason for performing an action just in case you know  $p$ . That can legitimately

be said to set an ideal — there really is no extra feature of a categorical attitude towards  $p$  that we would want to add once we know  $p$ ; it just doesn't get any better than that. But that isn't the case for probabilistic knowledge. Suppose I know that it is at least 50% likely to rain. And suppose I am deciding whether or not to take my umbrella. The higher the likelihood of rain, the higher my expected utility. If it's over 40% likely to rain, I maximise my utility by taking my umbrella. Thus, since I know it's at least 50% likely to rain, I should take it. But this piece of knowledge is not as good as it could be. If it's going to rain, it would be better if I were to believe that it is 100% likely to rain; if it's not going to rain, it would be better if I were to believe that it is 0% likely to rain. What's more, suppose I believe that it's at least 50% likely to rain; and suppose that belief is justified but not yet knowledge. It isn't obvious that I do better by gaining evidence that turns my justified belief that it's at least 50% likely to rain into knowledge than by gaining evidence that justifies a belief that it's at least 90% likely to rain, but which doesn't make it knowledge. And if that is not the case, then knowledge isn't the goal at which we always aim.

Before we wrap up, I'd like to draw attention to one final point, which is apt to be neglected. On the orthodox version of decision theory, an agent is bound to choose in line with her credences and her utilities — in the precise version of decision theory, for instance, she must pick an act that maximises expected utility by the lights of her current precise credences. Both Moss and Paul argue that this is too demanding in the case of an agent who has adopted the Fine-Graining Response and who sets her credences in the fine-grained states in line with the statistical evidence. Requiring that she chooses in line with her credences, Paul argues, is tantamount to requiring that she makes her decision by deferring to the utilities of others — and that way inauthenticity and alienation lie. For Moss, on the other hand, it is not reasonable to demand that an agent choose in line with beliefs in certain probabilistic contents — which is, after all, what her credences are — when she cannot rule out other probabilistic contents.

However, it is worth noting that the demand that orthodox decision theory makes is in fact rather weak. Suppose  $\mathbf{P}$  is the set of credence functions that represents the strongest probabilistic content that you know. Then, in many cases, and certainly the cases under consideration here,  $\mathbf{P}$  is also the set of all and only the credence functions that you are justified in adopting. Then, while it is true that, once you have picked your credence function  $P$  from  $\mathbf{P}$ , you are bound to maximise expected utility with respect to  $P$ , you are not bound to pick any particular credence function from  $\mathbf{P}$  — you might pick  $P$ , but equally you might pick any other  $P' \neq P$  from  $\mathbf{P}$ , and you would be equally justified whichever you picked. Thus, the set of permissible choices for you is in fact exactly the same according to the orthodox view and according to Moss' Probabilistic Knowledge Norm for Decisions, when that is coupled with a liberal decision theory for imprecise credences.

In each case, an act is permissible if there is a credence function  $P$  in  $\mathbf{P}$  such that the act maximises expected utility from the point of view of  $P$ .

I conclude, then, that Moss' No Knowledge Reply to the Fine-Graining Response does not work. I agree with Moss that credences that are based directly on sparse statistical evidence do not constitute probabilistic knowledge. But I argue that they are not justified either. And it is their lack of justification that precludes their use in decision-making, not their failure to count as knowledge. What's more, there are ways to set credences in the light of purely statistical evidence that gives rise to justified credences. Moss may say that these do not count as knowledge, and I'd be happy to accept that. But if she then also demands that credences used in decision making should be knowledge, I think the standard is set too high. Or, if she thinks that probabilistic knowledge simply serves as an ideal towards which we ought to strive, then there are times when I ought to abandon that ideal — there are times when I ought to pass up getting closer to knowledge in one probabilistic content in order to get justification in a more precise and useful probabilistic content.

## 8 Conclusion

In the end, then, I conclude that the Fine-Graining Response to Paul's Utility Ignorance Objection to decision theory is safe. Paul's Authenticity Reply does not work, as I have argued elsewhere. And nor does Moss' No Knowledge Reply, as I have argued here.

## References

- Buchak, L. (2013). *Risk and Rationality*. Oxford University Press.
- Dougherty, T., Horowitz, S., & Sliwa, P. (2015). Expecting the Unexpected. *Res Philosophica*, 92(2), 301–321.
- Elga, A. (2010). Subjective Probabilities Should Be Sharp. *Philosophers' Imprint*, 10(5), 1–11.
- Goldman, A. (1976). Discrimination and Perceptual Knowledge. *Journal of Philosophy*, 73, 771–91.
- Harman, E. (2015). Transformative Experience and Reliance on Moral Testimony. *Res Philosophica*, 92(2), 323–339.
- Jackson, F. (1986). What Mary Didn't Know. *Journal of Philosophy*, 83(5), 291–295.

- Jeffrey, R. (1992). *Probability and the Art of Judgment*. New York: Cambridge University Press.
- Jeffrey, R. C. (1983). *The Logic of Decision*. Chicago and London: University of Chicago Press, 2<sup>nd</sup> ed.
- Joyce, J. M. (1999). *The Foundations of Causal Decision Theory*. Cambridge Studies in Probability, Induction, and Decision Theory. Cambridge: Cambridge University Press.
- Joyce, J. M. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives*, 24, 281–322.
- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263.
- Lewis, D. (1980). A Subjectivist's Guide to Objective Chance. In R. C. Jeffrey (Ed.) *Studies in Inductive Logic and Probability*, vol. II. Berkeley: University of California Press.
- Moss, S. (2013). Epistemology Formalized. *Philosophical Review*, 122(1), 1–43.
- Moss, S. (ms). *Probabilistic Knowledge*. Oxford University Press.
- Paul, L. A. (2014). *Transformative Experience*. Oxford: Oxford University Press.
- Paul, L. A. (2015). What You Can't Expect When You're Expecting. *Res Philosophica*, 92(2).
- Pettigrew, R. (2015). Transformative Experience and Decision Theory. *Philosophy and Phenomenological Research*, 91(3), 766–774.
- Pettigrew, R. (2016). Book Review of L. A. Paul's *Transformative Experience*. *Mind*, 125(499), 927–935.
- Quiggin, J. (1993). *Generalized Expected Utility Theory: The Rank-Dependent Model*. Kluwer Academic Publishers.
- Rinard, S. (2015). A Decision Theory for Imprecise Probabilities. *Philosophers' Imprint*, 15(7), 1–16.
- Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons.
- Seidenfeld, T. (2004). A contrast between two decision rules for use with (convex) sets of probabilities: Gamma-maximin versus E-admissibility. *Synthese*, 140, 69–88.

- Seidenfeld, T., Schervish, M. J., & Kadane, J. B. (2010). Coherent choice functions under uncertainty. *Synthese*, 172, 157–176.
- Shimony, A. (1955). Coherence and the Axioms of Confirmation. *Journal of Symbolic Logic*, 20, 1–28.
- Stalnaker, R. C. (1970). Probability and Conditionals. *Philosophy of Science*, 37, 64–80.
- Wakker, P. P. (2010). *Prospect Theory: For Risk and Ambiguity*. Cambridge, UK: Cambridge University Press.