

“One day I shall be Queen?”

Thinking About Rational Choices in the Face of Epistemic Transformations

Chiara Stenico

Elizabeth, daughter of King George VI, happens to find herself touring the Commonwealth when her father passes away. The king of England has died, and she is next in the line of succession. This is an event she has not consciously anticipated. When she was born, Elizabeth was only third in line and thus not even remotely close to becoming queen. Her accession to the throne is an abrupt yet very real possibility. She now faces the following decision: Should she give up the throne or become queen?

How should she decide? This is a high-stakes decision, and it better be based on solid ground. Rational choice theory—the standard decision-making framework for most economists and many social scientists and philosophers—would give Elizabeth clear advice. First, she should consider the available options: become queen or abdicate. Second, she should identify the possible outcomes for each option. She should then assign probability and utility values that reflect her beliefs and preferences. Third, she should choose the option with the highest expected value. But can we plausibly expect Elizabeth, or any other decision-maker, to reason this way about such a decision?

Rational choice theory assumes that you can always determine expected utility values and always have an idea of what the space of all the possible outcomes looks like.¹⁰³ And yet, Elizabeth is not facing some ordinary decision, such as whether to go groom the horses or stay inside to read books. Her decision is not merely uncertain: it is *profoundly* uncertain on the values she can compute to being queen.

Elizabeth faces an extraordinary decision that will change her “in a deep and fundamental way” (Paul 2015, 761). Whether as a princess or presumably, as queen, her future self will likely believe and desire different things. Her uncertainty is then *more* profound than uncertainty on the outcomes of her decision insofar as her future self will change—due to one decision or the other—and with it the value that she would assign to such outcomes.¹⁰⁴ This sets her decision apart from any ordinary decisions. For being queen is not just any experience, but a life transformative one, or so says L.A. Paul.

In a well-known argument, L.A. Paul (2014; 2015) claims that these kinds of decisions—“life transformative decisions”—are a *challenge* that rational choice theory cannot meet.¹⁰⁵ She argues that this profound uncertainty on the outcomes from transformative decisions cannot be solved: as a decision-maker, you cannot assign values to such outcomes. Elizabeth cannot tell whether she will enjoy leading the country, nor what she will believe once in her father’s shoes.

Rational choice theory needs outcome values to formulate advice for Elizabeth. This means that she cannot use a decision rule—such as choosing the option with the highest expected utility—to make a principled rational decision for transformative decisions.

This conclusion could not but spark debate on the matter.¹⁰⁶ Perhaps you may not be ascending the throne

103 Savage’s decision theory for decision-making *under certainty* appeals to a spatial dimension. In his model, a decision is defined for a world, described by the states (or events) of that world (i.e. the ‘states space’) and the outcomes that obtain in that world (i.e. the ‘outcome space’). According to Savage’s theory, these worlds are necessarily well-specified. One idea is that outcomes are specified for each state of the world and under each act. For example, in the event of rain, whether I decide to take or not to take an umbrella, the outcome will be specified, and the same applies in the event of no rain (Steele and Stefánsson 2020).

104 Note that rational choice theory standardly works as a framework for decisions *under uncertainty*: that is, it requires that decision-makers assess their probabilistic beliefs that one outcome will obtain from an option. In Elizabeth’s case, the uncertainty also refers to how her beliefs and desires will have changed once the outcomes have been obtained.

105 Note that Paul challenges rational modelling on much weaker bases than Expected Utility Theory would require. Even allowing for imperfect conditions, for example, approximated values or ignorance, you cannot expect individuals—such as Elizabeth—to follow a rational decision rule for her decisions (Paul 2015, 152).

106 The problem of transformative experiences and decisions has been, in fact, widely discussed. Most replies to Paul’s argument are available in the special Issue “What you can’t expect when you’re expecting” in *Res Philosophica* 92 (2): 149-170.

anytime soon. Still, you may have to take analogous life transformative decisions: for example, deciding on having a kid or deciding on your career path. In that case, you would want a decision rule to guide you on properly thinking about your choice.¹⁰⁷

In this essay, I will consider a reply to Paul made by Jennifer Carr. This reply is interesting for it offers a promising argument to rescue rational choice theory from Paul's challenge. To do so, Carr uses epistemic utility theory and constructs a model such that rational evaluations for life transformative decisions are possible. Her idea is to look at what she calls 'conceptual resources' (hereafter: CRs). By clarifying their value to the agent's beliefs,¹⁰⁸ she uses them for a rational assessment which helps the agent decide. Even though CRs are central to Carr's model, she provides no substantial definition for them. You can think of them as *cognitive abilities*, such as drawing distinctions or inferences in the logical space of possible outcomes. One plausible and relatable example could be the following. Having and using the concept of 'honour' means having the ability to distinguish A (every outcome that requires an understanding of honour) from B (every other outcome that does not require so). Consider how A could be 'being honoured as a divine figure' in Elizabeth's case. In Carr's view, what matters is how she can draw this distinction between such an outcome and another. She seems to be satisfied with defining CRs on the basis of these distinctions only.

I will contend that CRs are not adequately defined by Carr. More specifically, using CRs to address the challenge that life transformative choices pose on the value-assignments on Elizabeth's outcomes requires to qualify CRs beyond a distinction-based definition. In other words, the solution that Carr proposes to make rational evaluations possible for transformative decisions is only workable if CRs are *qualified*. By a qualification of CRs, I mean being able to specify the CRs that Elizabeth would need to make a decision in the way Carr's epistemic model suggests, thereby solving Paul's challenge. My aim is two-fold: to (a) show how Carr's proposal has limits in the way she characterizes CRs, and (b) propose a qualified notion of CRs such that it does not have those same limits and can properly respond to the challenge and make use of her proposed epistemic model.

I will proceed as follows. In Section 2, I will shortly present Paul's argument on life transformative decisions. In Section 3, I will introduce Carr's reply. This section will outline the model Carr proposes for transformative decisions in fair detail and describe how she characterizes CRs. In Section 4, I will juxtapose her characterization with what I call a '*qualified*' notion of CRs. Building on such a qualified notion, I will show how Carr's proposal can be reconciled with Paul's challenge. Finally, in Section 5, I will conclude by returning to Queen Elizabeth II.

107 The interpretation or use of rational choice theory that I am assuming is more than just descriptive: it assumes that rational choice theory can be used to give normative prescriptions to an agent on how to act in the form of a decision rule.

108 Note that an agent's beliefs are expressed in probabilistic terms in epistemic utility models. This refers to an epistemic norm called 'probabilism' in epistemic utility theory (Pettigrew 2019).

2. The challenge of Laurie Paul

You and I make the most important decisions of our lives without knowing how the future will change us. I have decided to move to the Netherlands, not knowing the person I would become. Paul introduces her argument on life transformative decisions by questioning how normative decision theory could have helped me, could help Elizabeth or any of us with these decisions.¹⁰⁹ Paul goes on to claim: if you analyse these transformative decisions, you will find a problem with assigning values to your outcomes. Eventually, you will look at a challenge that cannot be solved.

As I have mentioned, applying rational choice theory requires that one assigns values to the potential outcomes of one's possible actions. However, in the face of life transformative decisions, one is unable to assign values to options that are transformative in their nature or, more precisely, to the outcomes of such options. The challenge to rational choice theory comes from the inability of an agent to apply a decision rule that can help her evaluate her options.

Now, how does Paul characterize these decisions? Were Elizabeth to choose to become queen, she would undergo what Paul calls a 'life transformative experience'. Having a kid, taking up an academic career, and becoming a vampire are all analogous experiences. What they have in common is a revelatory aspect at the moment when you enter this experience. That is, once you become a vampire, you are revealed what this experience is like: what the experience of immortality is like or what not being hungry for food feels like to you. The revelation or, if you want, epiphany¹¹⁰ brings about a new insight that primarily refers to *beliefs* and *desires*. Paul explains (2015, 155) experiences are life transformative when, before transformation, that is before any revelation or epiphany, you cannot reason what your experience *will be like* post-transformation, that is when the epiphany is concluded. This is because you cannot know how your future self will experience and think about that same experience. You cannot know what it will be like for decisions with a transformative character, such as having a kid, choosing an academic career, becoming a vampire, a priest or a queen. You cannot know what you will feel or what you will *believe* and *desire* being a mum, a scholar, a different person, a queen. At the same time, you cannot know what it would be like if you had chosen differently.

Paul's argument distinguishes two kinds of transformations for transformative experiences: an epistemic transformation and a personal transformation. These correspond to two challenges.¹¹¹ In this essay, I will only address the former epistemic challenge that arises from an epistemic transformation.¹¹² The epistemic transformation refers to your newly transformed epistemic situation, that is, the kinds of beliefs you hold. The idea is that post-transformation, you obtain '*what it's like*' knowledge. This is, again, knowledge in terms of your beliefs and desires. Once you have made your decision, you obtain knowledge and hence learn about the values of being queen or not being queen. Before then, you cannot know nor significantly learn what it is like to be queen or not to be queen. In this sense, you are in a different—"impoverished"—epistemic position than prior to transformation (Paul 2015, 155). The epistemic transformation describes one kind of uncertainty that rational modelling *cannot*, according to Paul, overcome or solve: the uncertainty about Elizabeth's epistemic situation. How should Elizabeth address such uncertainty?

109 This is, as mentioned above in footnote n.2, a decision theory which has more relaxed assumptions than standard rational choice theory: Paul wants to capture norms for ordinary successful reasoning, and hence, account for approximated values or ignorance of the decision-maker.

110 This is a term Paul has sometimes used to give an idea of what she means by radical transformation without any definitional commitment.

111 Note that Paul herself does not use this terminology. She only talks about 'epistemically' and 'personally transformative experiences', which I translate as the two challenges to be met if we want to use rational modelling for these transformative experiences. Carr's argument does not tackle the personal transformation challenge. Instead, she assumes stability of preferences. In this sense, Carr's argument is not an exhaustive counterargument to Paul's transformative experience problem. It does not solve every issue that Paul raises about transformative experiences and decisions.

112 There is a second challenge which completes, in fact, the challenge that transformative decisions raise on outcome values: the personal challenge. The upshot of it is that an agent's core preferences may change post-transformation (ibid., 155). In Elizabeth's case, the preferences that she holds over her options when touring the Commonwealth do not necessarily match those she would hold post-transformation, either as a princess or as queen. Note that core preferences are higher order preferences. The idea is that such core preferences will unlikely be stable in the face of a transformative experience: this means that the agent cannot know at the time of the decision-making whether she will satisfy her preferences post-transformation (Paul 2015, 521).

3. Jennifer Carr and a reply to L.A. Paul: a model for Epistemic Transformations and Conceptual Resources

When rational modelling seems to be of no help with transformative decisions, J. Carr in *Epistemic Expansions* (2015) shows a way out. Her article argues that the *epistemic challenge* need not be a challenge to rational choice theory. To do so, she proposes an *epistemic utility model* for transformative decisions.

Broadly, the model rests on the idea of conceptual resources—the agent’s cognitive *abilities* to draw distinctions in the logical space of possibilities. First, CRs are centrally relevant to Carr’s model because they explain what the problem of uncertainty on the outcomes is all about. Second, Carr translates such a problem into one of conceptual resources: Elizabeth is incapable of assigning the outcome values because she does not have the CRs to do so.

In the following sections, I will clarify Elizabeth’s problem of CRs by describing the epistemic model that Carr has in mind. In Section 3.1, I will first describe the idea of partial credences and how they relate to credence functions and CRs in Carr’s model. Then, in Section 3.2, I will discuss how CRs are used in the model: these are what you *compare* to evaluate your beliefs.

3.1. J. Carr’s Epistemic Utility Model

Carr starts from the premise that you do not always have an opinion, in terms of (dis)beliefs or upholding judgement, towards every proposition (Carr 2015, 219). A proposition is the object of an agent’s beliefs. Together with others, a proposition describes a possible world (Genin and Huber, 2021), for example, a world in which you are the Queen of England—in other words, Elizabeth’s outcomes. Sometimes, you are not in the right position to say something about this possible world or even see a proposition in such a world. For example, you have never been queen: you may not be able to form a belief, nor to uphold judgment about ‘waking up in the morning *as* a queen’. Maybe you do not even know there is such a thing as ‘waking up as a queen’. This means that for sets of propositions like ‘what it is like to be a queen’—a set of propositions including ‘waking up in the morning as a queen’ and ‘having lunch with Jackie Kennedy’—you may only have partial credences. Let me clarify.

Carr uses a *partial credences’ framework*,¹¹³ in which a *credence* represents the strength of your beliefs in a proposition. Epistemic utility theory constructs functions for credences. Formally speaking, a credence function is defined over a set of propositions, returning values (real numbers $[0,1]$) for each individual proposition (at a given time t). Now, epistemology commonly does not account for those cases in which you do not have an opinion on a proposition (Carr 2015, 2019). But, as I have mentioned, there can be such cases, and these are cases of transformative decisions and experiences. Carr wants to open up epistemology to them. She introduces the possibility for a *partial* credence function in her model, which is an epistemic utility function that is *only partially* defined. This means that she explores an epistemic utility function for Elizabeth such that *it does* account for the fact that she may not have an opinion on some propositions.

Now, a (partial) credence function assigns an epistemic utility (dis)value to Elizabeth’s credences. Meaning that it evaluates Elizabeth’s *degree of belief* over all the propositions that make up a world in which she is queen (e.g. having lunch with Jackie Kennedy), and over all the propositions that make up a world in which she is just a princess (e.g., having lunch with her kids). Once again, these worlds’ descriptions apply only to those propositions that Elizabeth can “speak about” or even “see”.

¹¹³ Epistemic utility theory is in line with theories of partial beliefs (for example, Bayesianism). For more information refer to Pettigrew (2019).

Here is where CRs, the core concept for the epistemic model, comes into play: CRs. In fact, CRs determine whether I can or cannot see some propositions, and so they determine the (partial) credence function for a world. Remember how Carr conceives of CRs, merely as the abilities “to draw logical distinctions in the space of possibilities” (ibid., 225). Plausibly, what she means by this is being able to both see and say something about the propositions of an (impoverished) prior to transformation world, and about a post-transformation one: for example, propositions of a world in which Elizabeth is just a princess and those of a world in which she has ascended the throne, and she is effectively experiencing being queen. More precisely, a CR allows to distinguish between propositions prior to transformation, and post-, for example, between a proposition (e.g., living in Buckingham Palace) and its negation (i.e., not living there). Think back to the initial example of the concept of honour. Having such a concept allows us to distinguish between the proposition ‘being honoured as a divine figure’ and ‘not being honoured as such’. In other words, having this concept makes it so that you can see both the propositions just mentioned.

3.2. Conceptual Resources and Comparability

Provided that the conceptual tools determine credences and the related (partial) credence functions at hand (ibid., 220), an epistemic model should be able to account for and (epistemically) assess conceptually distinct situations. These are situations in which the decision-maker does not have the CRs to assess all the propositions of a world—hence, establish the value of an outcome; and situations where she does hold all CRs to have a complete credence function. Carr’s model aims to address the issue of transformative decisions and experiences by accounting for the conceptual change that occurs via the underlying transformation: from not having what it is like knowledge—that is, not being able to say something about or even see some propositions of a world—to the opposite situation.

Carr’s model accounts for such a conceptual change as follows. The model compares partial credence functions (c), characterized by limited conceptual possibilities, with complete credence functions (c+).¹¹⁴ These are, respectively, the partial credence function describing an epistemic state before transformation (e.g., the state in which Elizabeth is touring the Commonwealth), and the complete credence function describing the post-transformation state (e.g., the state in which Elizabeth is either queen or princess upon her decision). Consider that, traditionally, epistemic decision theory presupposes credence functions’ comparability over the *same* set of propositions. This means that it only compares situations where the decision-maker “sees” all the propositions. Carr does not want to accept this limit for epistemic decision theory, and she aims to contribute by making sure her model “accommodates that conceptual change has epistemic (dis)value” (ibid., 222).

To sum up, Carr’s model aims to address the issue of transformative decisions and experiences by accounting for the conceptual change that occurs with the transformation: from not being able to say something about or even see some propositions of a world to having what it is like knowledge. This, in formal terms, means to allow for comparability of credence functions which are defined over different sets of propositions. How these credence functions are assessed (in terms of epistemic (dis)value) depends on the properties of the credence function: I will not address this second—normative—aspect of the model. Instead, the elements of the model so far described will suffice to follow my argument for why the definition that Carr provides for CRs is problematic to answer the epistemic challenge.

¹¹⁴ Note that a complete credence function describes a situation in which the agent has sufficient CRs to assign values to each proposition in a post-transformation world. It follows that this is a situation in which one is able to see and speak about all the propositions that describe that world.

4. An alternative to CRs: *Qualified CRs*

This section will introduce why Carr's definition of CRs is unsatisfactory and why she should refer to a qualified version of CRs instead. In what follows, I will first, in Section 4.1, describe my proposal for qualified CRs. Qualified CRs will provide more information about CRs themselves, against a minimal understanding of those, and better clarify the value of CRs to an agent's credences. Second, in Section 4.2, I will illustrate what motivates my proposal and why qualified CRs are, in fact, able to meet the challenge that Paul raises.

4.1. A proposal to qualify CRs

Carr does not offer a substantive definition of CRs. In her view, CRs have two characteristics: (i) they are abilities to see and say something about the logical space of possibilities; and (ii) they allow you to see and distinguish A from $\neg A$, infer that A or $\neg A$, and negate that A or not $\neg A$. Thus, CRs can allow you to distinguish between 'living in Buckingham Palace' and ' \neg living there', inferring, and negating other propositions that you see. The problem with this definition for CRs is that it does not contain any specification on the cognitive tools to draw the mentioned distinctions.

It seems that if you consider the ability to draw logical distinctions between the various propositions that you can entertain in different worlds, you may want to know what *kind* of abilities allow you to draw distinctions. In other words, how do you distinguish between A or $\neg A$? In this sense, you may want to qualify those abilities to draw distinctions. One relevant reason for this is precisely an informative one: you may want to know more about how CRs work. Recall how CRs determine a credence function, returning values for those propositions that you see or can speak about. By specifying those abilities to draw distinctions, qualifying CRs can provide information on how an agent's credences relate to the CRs she holds or not in a world or outcome.

Now, consider how Elizabeth could possibly qualify CRs in her specific case. The CRs she may have at hand—determining a credence function—may be rather varied. Among these CRs there could be the following: (1) Testimony reliance (Dougherty, Horowitz and Sliwa 2015, 308): Elizabeth could consult her uncle who, in the recent past, abdicated from the throne and whom she trusts as a testifier. (2) Inferences from relevant subjects' behaviour (ibid., 311): Elizabeth could appeal to what she observed about her late father in the past years. (3) Inferences from similar (and resembling) experiences (ibid., 312): in this case, Elizabeth may refer back to her experiences as royal, particularly when people were giving her a distinct treatment.¹¹⁵ Note how these instances of CRs are in line with Carr's definition for CRs, each one describing an ability to see and say something about the space of possibilities. Take testimony reliance: this can be interpreted as the ability to engage with the space of possibilities, based on someone else's word. Furthermore, it seems plausible to think that these kinds of CRs, qualified CRs, would help determine epistemic value of a credence function. Elizabeth could realistically form a probabilistic belief on what it will be like to govern based on her uncle's testimony.

Consider now the following to see how there are advantages to qualifying CRs beyond 'cognitive abilities to draw distinctions'. Based on the same testimony, Elizabeth could plausibly form a probabilistic belief on what kinds of marital duties and concerns she will have once she is queen. By qualifying CRs the way I did, Elizabeth could reasonably be able to assign values to propositions other than A ; $\neg A$, that is a *binary* distinction between logical possibilities. Elizabeth could assign a value to an additional proposition to the set

¹¹⁵ Here I am assuming all these resources are reliable. The probabilistic credences we form about them do not deal with reliability as such, but with the expectations on truth values.

of propositions, for example, ‘having marital infidelity concerns’. A qualified CR—in this case, testimony reliance—seems to be able to uncover other propositions than A and $\neg A$ (e.g., ‘being respected and honoured by the British citizen’ and ‘ \neg being respected and honoured’).

If what I just described is plausible, qualified CRs seem to have two positive implications: having *more credences*—probabilistic beliefs over single propositions—and, at least in some cases, uncover *additional* propositions to the overall set of propositions.¹¹⁶ One CR such as testimony reliance will help Elizabeth develop probabilistic beliefs over ‘what it will be like to govern’ and other single propositions making up a world in which she is the sovereign—that is, the set of all propositions. What is more, Elizabeth has the chance to uncover some propositions which may be significant to her decision, for example, propositions about marriage. Generally, the idea is that qualified CRs seem to describe the space of possibilities and one’s ability to navigate this space or world in more detail.

These two positive implications from qualified CRs show how qualifying CRs clarifies the relationship between CRs and an agent’s beliefs. First, it tells you that one can form a belief on one or more propositions based on the same CR. The uncle’s testimony reliance will plausibly help Elizabeth develop beliefs over different propositions. Second, qualifying CRs tells you that more propositions can be uncovered such that you can form a belief about them. The same testimony reliance can uncover propositions other than A and $\neg A$: for example, ‘being respected and honoured by the British citizen’ and ‘ \neg being respected and honoured’. Not qualifying CRs tells you no such things, in fact, it leaves the relationship between CRs and an agent’s beliefs rather vague.

I will now turn to how a *qualified* notion for CRs can address Paul’s epistemic challenge. That is, I will investigate why a qualified notion should be preferred: to adequately account for the epistemic challenge.

4.2. *Qualified CR as Means to Truth-ends*

The reason why Carr’s definition of CRs is unsatisfactory relates to the epistemic challenge and one adequate response to it. First, remember how Paul raises the epistemic challenge on account of the incapacity of an agent to assign values to her outcomes due to a lack of the necessary ‘*what it’s like*’ knowledge. The agent can only access this knowledge post-transformation, making it impossible for her to rationally assess outcomes from transformative decisions.

Now, I want to argue that only a qualified version of CRs can adequately respond to this challenge and recover the agent’s ability to assign values to the outcomes from her decisions. To do so, I propose to examine what this challenge aims at.

A reasonable interpretation of the ends of the challenge would be an agent’s access to outcome values,¹¹⁷ post-transformation (Paul 2015, 157). Particularly, Paul thinks of *subjective* values (ibid., 152). These are the values of being in a mental state experiencing ‘what it’s like’ (ibid., 153). For Elizabeth, this is a mental state that she can be in experiencing being queen. Note how these are ‘first-personal values’ (Paul 2014, 6), derived from how you—personally—experience a certain experience. Carr aims at a much less demanding result: *epistemic values*. Specifically, values that tend to be as accurate as possible, maintaining an objective truth as the standard of accuracy. This means that credence functions would ideally return true values, according to an objective notion of truth: “a credence has greater epistemic utility at a world if it’s closer to the truth in that world, i.e. if it has greater accuracy” (Carr 2015, 513). Thus, what is desirable is proximity

¹¹⁶ How to sum or subtract single credences is an important matter. The different strategies to do that describe the properties or epistemic norms of an epistemic model. Refer to Carr (2015), Section 3, for a discussion on such properties.

¹¹⁷ Note that Paul argues such an end can never be met. You must live through that specific experience to get to know ‘what it’s like’ for your values. This is a limit to human nature and you cannot get a hold on these phenomenal values but when you have experienced it.

to truth, or conversely, distance from falsity (ibid., 221).¹¹⁸

Now, you can find symmetry between what both the epistemic challenge and Carr's epistemic model aim at: roughly arriving at some kind of truth values. Consider the first kind of values, namely subjective values. You would want to access such values for pragmatic purposes to make the highest expected utility decision. You would want to know how you personally will live a transformative experience, to make a good assessment of your options prior-transformation and make the best decision possible. Paul's relentless demand for subjectivity reveals the relevance of subjective values.¹¹⁹ Subjective values play a pivotal role in her argument, for any proposal that rules out a first-person perspective is not worth considering because not authentic (Paul 2015, 761). Such values should match the *truth* of your beliefs, desires, dispositions post-transformation. The values you would want to access are those that map *truthfully* onto your subjective values post-transformation. It follows that, reasonably, you can interpret these subjective values as containing a claim to truth and this is reaching the truth in a post-transformation world. Such truth is independent of your current desires and beliefs. In this sense, it is a truth *in* a post-transformation world parallel to the truth that the epistemic model aims at: a truth *in* that world, however dependent on your future beliefs and desires formed only due to your transformative experience.

You may wonder: Why expose the symmetry of ends? After all, the model already assumes that truth is what you care about when it comes to epistemic transformations. Yet, Carr is saying little about what truth means to the epistemic challenge posed by transformative experiences. Only once you have established this—that such a truth is, in fact, meaningful to the epistemic challenge—can you claim to be proposing a model for transformative decisions that adequately accounts for a fundamental aspect of transformative decisions. In fact, you can clarify the relationship between truth and CRs by showing how the notion I propose can reconcile Carr's model with Paul's challenge.

How do *qualified* CRs meet the ends of the epistemic challenge? In fact, how do *qualified* CRs help us assign values to the outcomes from a transformative decision? Assuming that CRs are the abilities that you need to approach truth values, what helps you get there—moving closer to meet such values—should reasonably be given more *emphasis*. In other words, the *ends* to the epistemic challenge demand more emphasis on the *means* to such ends. CRs should reasonably be paid more attention in Carr's attempt to have rational assessments for transformative decisions.

Consider the following. CRs allow you to be further or closer to the truth (value 1), compared to non-means, that is not using any CR at hand. A CR, conceived of as the ability to see or say something about a proposition, having an opinion on it, seems comparatively *better* than not having it. The idea is to show how having more concepts and credences precisely puts the necessary emphasis on the means to truth-ends. This is because you better explore how CRs relate to an agent's beliefs.

First, note that not having an opinion or belief *does not* in any way relate to truth or falsity. You could decide to treat this situation, which is formally called a "non-attitude", as the furthest you can be from truth or, for example, impose a rule on how to value non-attitudes.¹²⁰ Formally, this means that you would decide on the properties for your epistemic utility functions. Assume for now that you do not—you do not decide how to treat non-attitudes. This means that you only consider those propositions you can form an opinion about for partial credence functions—the propositions you can see and speak about. You can show that qualified CRs are the means to a desirable truth-end, no matter non-attitudes.

Consider one qualified CRs example. You can plausibly claim that Elizabeth is in a *better* epistemic position by relying on the testimony of her uncle and reminding herself how she, as a princess, was never allowed to speak her mind. She has more knowledge of the truth (and falsity) *than* if she had none of the

¹¹⁸ Most epistemic utility models assume gradational accuracy which is the closeness of a credence function to truth.

¹¹⁹ This is a recurrent theme in Paul's Replies (2015) to the other authors.

¹²⁰ This is what Carr (2015) does when discussing the properties of her model. See Section 3 of her text.

mentioned concepts. This does not seem overly contentious. Having more concepts and more credences, at least in this case, seems better than having fewer credences. Specifically, the idea of summing positive credences' utilities (however small) seems to benefit a truth-end. Because you are silent on non-attitudes, having more credences, more beliefs over propositions world—the world in which Elizabeth is queen, for example—seems better than having less of those. Having more credences is then desirable to a truth-end.

My proposal for qualified CRs shows how only by qualifying CRs can you clarify the positive relationship between credences and CRs, which is not otherwise substantially clarified with non-qualified CRs. Recall how Section 4.1 concludes on the two positive implications of my proposal: having more credences over single propositions and uncovering additional propositions in the set of all propositions. Consider the testimony from Elizabeth's uncle. Given this CR, she is in a better epistemic position because she cannot only judge a proposition 'x', but also judge an additional proposition 'y'. The same testimony of Elizabeth's uncle allows her to judge, say, other people's inauthentic behaviour, but also form a credence about the crown's interference in her marital life. Qualifying CRs this way entails having more credences, not only in virtue of specifying a higher number of CRs but also in virtue of having such specified concepts to help you with a richer array of propositions, a more detailed world of propositions. Thus, only exploring the relationship between credences and CRs, and observing the positive implications just mentioned, seriously emphasizes the truth-ends of the epistemic challenge.

To clarify why qualifying CRs adequately considers CRs as means to truth ends, suppose that you do not qualify CRs. Suppose, in fact, that you are only using CRs as Carr intended in her model. Elizabeth can now likely only distinguish between, say, 'being treated differently' from '–being treated differently', and not ever come to uncover a proposition such as 'having marital issues'. With qualified CRs she can do more than a binary propositional distinction. She is in a better epistemic position *than* with unqualified CRs because she can assign truth-values to more propositions *than* she could have foreseen. She can assign a truth value to 'having marital issues' too and better define the truth of that world or outcome: being queen, post-transformation. Failing to explore the relationship between credences and CRs, and identify such a better epistemic position, means failing to have truth-ends seriously considered within the model. If you want to do so—seriously considering truth-ends—you should explore and better clarify the relationship between CRs and credences. *Qualified* CRs precisely take this role and allow you to see how more credences better serve truth-ends.

Now, recall that I assumed non-attitudes to be undefined on truth-values. Leaving non-attitudes undefined or untreated would not do justice to an argument based on a positive relationship between credences and CRs, and the idea of having more credences. Now, you can show that even when defining non-attitudes on truth-values, that is when you say how one should value a non-attitude, one should choose to qualify CRs.

Suppose there can be non-attitudes—that is not having an opinion or a belief over a proposition—that are better than credences that are 'no better than chance' (Carr 2015, 232): for example when Elizabeth can attach a small probability (less than 0.5) to what represents her belief about having marital issues. In this case, you may think a non-attitude is better. In fact, what use could Elizabeth make of such a credence about a world? Would this mean that having more beliefs is not always better in transformative decision-making, after all? From a practical rationality perspective, it seems so. If Elizabeth could make a choice, a credence with a small probability attached would be useless to her evaluation. It would not help an assessment of the expected utility values of her decision options. How should she be helped by knowing that she most probably will not have marital issues? From an epistemic rationality perspective,¹²¹ it is not as clear-cut. Consider that if a qualified CR, for example, relying on your uncle's testimony can apply to multiple propositions, at the very

¹²¹ This is when a decision maker worships the epistemic features of the outcomes of a given decision, rather than their non-epistemic ones, as it would normally occur in a pragmatic rationality framework (Paul 2015, 219).

least, it increases the chances to be able to obtain a credence that is better than mere chance. In fact, it could be uncovering a proposition over which you have a credence that is better than mere chance. For example, you could have such a credence over ‘governing the country’. This is, by uncovering more propositions and exploring the space of possibilities where better credences could be. A non-attitude—conversely—can only apply to a proposition and its negation. It seems that even if you specify how to treat non-attitudes and you, in fact, reject that these can always be helpful to a rational evaluation, an emphasis on credences is still justified.

To sum up, in this section, I have explained why having qualified CRs, implying forming beliefs about more propositions, is comparatively better than not having qualified CRs and thus not forming such beliefs. Having qualified CRs better clarifies the value of CRs to an agent’s credences. This shows how these are genuine means to truth-ends, the ends of the epistemic challenge. As such they should receive emphasis. CRs have a guiding value that is purely epistemic: they can guide Elizabeth in her knowing the truth. The idea is then relying on the full capacity of the CRs at hand, and granting them to be, in fact, resources.

5. Conclusion

Let me go back to Elizabeth dealing with a life transformative decision on her future as queen. I have given you a rough idea of how rational choice theory would think about Elizabeth’s decision, laying out her options and making certain value calculations. Following this framework, she would have a clear, most preferred option. Now, Paul would ask: How can Elizabeth know how she will value ‘being the most powerful woman in the world’ before experiencing being queen? Paul would say that these values that you have computed are not *truthful* to your values post-transformation. I have clarified how this is, in fact, a plausible formulation of the epistemic challenge posed by transformative decisions. Carr’s model for epistemic transformations allows you to think about truth as something that you—conversely—*can* aim at. This is possible if you think of Elizabeth’s transformative decision in terms of CRs. To Carr’s mind, these are the abilities to draw logically binary distinctions of the kind: ‘being treated differently’ and ‘–being treated differently’. Carr would model Elizabeth’s decision as a comparison between two epistemic situations—abdicating or becoming queen—which are significantly defined by the CRs you hold in each one of them. CRs are what Elizabeth should consider when assessing the epistemic value of abdicating or becoming queen. The problem is, I have argued, that Carr’s definition for CRs is not satisfactory. CRs need to be qualified.

My proposal is *qualifying* Carr’s notion of CRs beyond abilities to draw logically binary distinctions. Accepting my proposal would mean for Elizabeth to consider conceptual abilities that allow her to discover new propositions by multiple conceptual means: for example, by relying on somebody’s testimony on a proposition. I have argued (in Section 4.1) how this qualifying move can expand the credences about a set of propositions and expand the set of propositions itself. This is meaningful to *truth*, which means that *qualified* CRs are relevant to truth-ends. You need *qualified* CRs to meet the epistemic challenge that transformative experiences pose. Eventually, I have shown (in Section 4.2) how qualified CRs can fulfil truth-ends on account of having more credences and having the relationship between CRs and an agent’s beliefs better clarified. In my refined version of Carr’s solution to the epistemic challenge, this is how Elizabeth should now think about her decision: she should appeal to qualified CRs to best define the epistemic value that the options hold. Should she ever decide to become queen—I conclude—rational guidance should be coming from the use of *qualified* CRs in a framework where truth is the end-goal.

References

- Carr, Jennifer. 2015. "Epistemic expansions." *Res Philosophica*, 92 (2): 217–236.
- Dougherty, Tom, Sophie Horowitz, and Paulina Sliwa. 2015. "Expecting the unexpected." *Res Philosophica*, 92 (2): 301–321.
- Genin, Konstantin, and Franz Huber, "Formal Representations of Belief." In *Stanford Encyclopedia of Philosophy*. Stanford University, 1997-. Article published October 22, 2008; last modified November 13, 2020, <https://plato.stanford.edu/archives/spr2021/entries/formal-belief/>.
- Gibbard, Alan 2007. "Rational credence and the value of truth." In *Oxford studies in epistemology*, edited by Tamar Szabo Gendler, and John Hawthorne, 2: 143–164. Oxford: Oxford University Press (2007).
- Margolis, Eric, and Stephen Laurence, "Concepts." In *Stanford Encyclopedia of Philosophy*. Stanford University, 1997-. Article published November 7, 2005; last modified June 17, 2019, <https://plato.stanford.edu/archives/sum2019/entries/concepts/>.
- Pettigrew, Richard. 2013. "Epistemic Utility and Norms for Credences." *Philosophy Compass* 8 (10): 897–908.
- Pettigrew, Richard, "Epistemic Utility Arguments for Probabilism." In *Stanford Encyclopedia of Philosophy*. Stanford University, 1997-. Article published September 23, 2011; last modified November 6, 2019, <https://plato.stanford.edu/archives/win2019/entries/epistemic-utility/>.
- Paul, L. A.. 2015. "What you can't expect when you're expecting." *Res Philosophica* 92 (2): 149–170.
- Steele, Katie, and H. Orri Stefánsson, "Decision Theory." In *Stanford Encyclopedia of Philosophy*. Stanford University, 1997-. Article published December 16, 2015; last modified October 9, 2020, <https://plato.stanford.edu/archives/win2020/entries/decision-theory/>.