

EPISTEMIC EXPANSIONS

Jennifer Carr

Abstract: Epistemic transformations—changes in one’s space of entertainable possibilities—are sometimes rational, sometimes irrational. Epistemology should take seriously the possibility of rationally evaluable epistemic transformations. Epistemic decision theory compares belief states in terms of epistemic value. But it’s standardly restricted to belief states that don’t differ in their conceptual resources. I argue that epistemic decision theory should be expanded to make belief states with differing conceptual resources comparable. I characterize some possible constraints on epistemic utility functions. Traditionally, it’s been assumed that the epistemic utility of a total belief state determines the epistemic utility of individual (partial) beliefs in a simple, intuitive way. Naive generalizations of extant accounts generate a kind of repugnant conclusion. I characterize some possible alternatives, reflecting different epistemic norms.

I’ve never had a child. I’ve never tasted an oyster. I’ve never experienced war. I don’t have the slightest idea what any of these experiences is like. I can’t even entertain the possibilities for what they are like. How can I make rational decisions about whether to have a child, taste an oyster, or go to war, when I have so little idea what kind of outcome I’d generate? Normally I could decide on the basis of thinking about the possible outcomes of my actions: how valuable each would be, and how likely it is. But I can’t do that.

This is L. A. Paul’s (2015) challenge: what could rationalize a decision about whether to perform any of these actions? How can there be a decision theory for partial credence functions, when decisions hinge on possibilities the agent can’t entertain? The problem is not uncertainty: it’s not simply that the agent is unsure of the outcomes of her actions. Rather, the problem is limited conceptual resources: there are some possibilities that the agent can’t “see,” propositions she isn’t in a position to entertain.

We can formulate the puzzle, then, in overtly decision theoretic terms. Suppose an agent has only a partial credence function, one that doesn’t range over some of the possible outcomes of an action available to her. Then

1 the expected utility of that action is undefined. So what could rationalize
2 the choice of whether to perform it?

3 An epistemic analogue to Paul’s challenge: how can there be an epistemic
4 decision theory for an agent with a partial credence function, when her
5 epistemic “decisions” hinge on possibilities the agent can’t entertain?¹

6 Suppose I’m deciding whether to ϕ , where ϕ ing involves changing my
7 credences. But I don’t know the epistemic utility of ϕ ing. Its epistemic
8 utility depends on what state of the world I’m in: in particular, on which
9 credence function ϕ ing will lead me to adopt. What’s worse: I can’t even
10 entertain these credence functions, because they involve concepts that I
11 don’t possess. They distinguish among possibilities I can’t distinguish. So
12 how am I supposed to choose whether to ϕ ?

13 The epistemologized version of Paul’s challenge, then, is: can we extend
14 epistemic decision theory for agents with partial credence functions? For
15 agents whose credences can change domain, so that the agent can come to
16 see different possibilities? I’ll argue that we can. Moreover, we should: it’s
17 sometimes irrational to change the domain of one’s credence function.

18 The plan of the paper is as follows: I argue that epistemic decision theory
19 should be expanded such that credence functions with different domains
20 are sometimes comparable. The argument is based on a weak conservative
21 principle: that it’s sometimes irrational to lose conceptual resources. Then
22 I characterize some possible constraints on epistemic utility functions that
23 compare credences with different domains. Traditionally, the epistemic
24 utility of a total credence function is understood as a function of the
25 epistemic utility of credences in individual propositions. The most natural
26 ways of generalizing to partial credences generates a kind of repugnant
27 conclusion.² I argue for a general constraint on the space of possible
28 algebra-neutral epistemic utility functions. I characterize some possible
29 versions, reflecting different epistemic theories.

31 1 Paul’s challenge epistemologized

33 1.1 Partial credence functions

34 Jackson (1982) argued that before ever seeing the color red, there was
35 something Mary failed to know: what it’s like to see red. Jackson’s thought
36 experiment arguably supports a stronger conclusion: Mary can’t even
37 conceive of what it’s like to see red. Of course, she can entertain the
38 possibility that what it’s like to see red is the same as what it’s like to see
39 dark gray, or the same as what it’s like to taste an oyster. But there are many
40

42 ¹ To highlight the analogy with practical decision theory, I use voluntaristic phrasing (“de-
43 cisions,” “options,” etc.). But epistemic decision theory doesn’t presuppose epistemic
44 voluntarism.

45 ² Parfit 1984.

1 possibilities for what it will be like to see red that Mary can't entertain,
2 including what it's actually like.

3 Like Paul, we begin with the assumption: people don't always have
4 attitudes toward all propositions. Within the partial belief framework:
5 sometimes a person's credence function isn't defined over all propositions.

6 This is already controversial. On the traditional view of the psychol-
7 ogy of credences, both credence and utility functions are understood as
8 abstractions from dispositions to choice behavior. There are representation
9 theorems that show that, if an agent's dispositions satisfy some constraints,
10 the agent is describable by a unique credence function and a utility function
11 unique up to positive affine transformation. The traditionalist may insist
12 that credence functions aren't partial. Or, if she allows that they may be par-
13 tial, the traditionalist will insist may insist that at least credence functions
14 are not in any way that could change over time. Credences are automati-
15 cally defined for all possibilities the agent could possibly encounter, since
16 for each encounterable possibility, the agent will have choice dispositions.³

17 We might respond by simply accepting that the dispositionalist psychol-
18 ogy of credences is wrong. But there are also possible replies to this objec-
19 tion that maintain the spirit of the dispositionalist view. First, we might
20 question whether an agent really has the relevant dispositions to choice
21 behavior even for choices that hinge on propositions she can't entertain. It
22 might be that, instead, agents have dispositions to acquire dispositions to
23 choose once the relevant conceptual resources are active.⁴

24 Second, we might accept that credences are dispositions to choice behav-
25 ior that meet certain necessary conditions. Plausible theories of intentionality
26 impose some conditions on propositional attitudes. These conditions
27 may be externalist: in order to have *de re* thoughts about individuals, it
28 might be that you need to have a particular sort of causal connection with
29 those individuals. (Shakespeare couldn't have had *de re* thoughts about
30 Cher.) They may also be internalist: in order to have thoughts about phe-
31 nomenal redness, it might be that you need to have already experienced
32 phenomenal redness.

33 A different kind of objection: decision theory characterizes agents who
34 are idealized in all kinds of ways. For example, their credences are infinitely
35 non-vague. Traditional decision theory presupposes credences that are
36 infinitely precise. Fans of imprecise credences functions assume credences
37 take sets of reals as values, and so imprecise credences have infinitely precise
38 boundaries. And so on. Ideally rational agents aren't computationally
39 limited. So why allow that they can have merely partial credence functions?

40 Reply: if there are external or experiential conditions on having proposi-
41 tional attitudes, then plausibly ideal rationality is compatible with having
42 partial credence functions. Ideal rationality doesn't require us to be in a
43

44 ³ The traditionalist assumes that agents' dispositions are determinate, though not fixed.

45 ⁴ Thanks to Robbie Williams for this suggestion.

1 certain sort of environment, or to have had certain kinds of perceptual
 2 experiences. Mary is conceptually limited not because of some irrationality
 3 in her belief state, but because a certain kind of experiential state is a
 4 necessary precondition for entertaining what it's like to see red. Rationality
 5 doesn't require one to have already experienced red phenomenology.

6 In short, the presupposition that rational agents have total credence
 7 functions is not an intuitive assumption about ideal rationality. It's an
 8 idealization only in the sense of "simplification." When this simplification is
 9 removed, it opens up substantive questions about rational decision making
 10 and rational update.

11 I'll assume that even partial credence functions are defined over boolean
 12 algebras of propositions, the strongest elements of which form a partition
 13 over the set of worlds. In other words, I assume there's an exhaustive set of
 14 mutually exclusive 'basic possibilities' that are "visible" to an agent; the
 15 agent's credences are defined over all unions of visible basic possibilities
 16 plus the empty set. What justifies this assumption? Plausibly, while ideal
 17 rationality doesn't require being able to see all propositions, it does require
 18 being able to negate and conjoin the propositions you do see.

19 The partition of basic possibilities an agent can see characterizes the
 20 distinctions the agent is able to make. As an intuitive shorthand, I'll talk
 21 about an agent's "concepts" or "conceptual resources." Note: I don't mean
 22 to make any substantive commitments about the psychology of concepts.
 23 Indeed, nothing in the way we're modeling things will play the role of "a
 24 concept," qua subpropositional mental representation. The distinctions in
 25 the space of possibilities that an agent can make are presumably related
 26 to her conceptual resources, but I acknowledge that that relation might
 27 be very messy. If the shorthand seems misleading or distracting, talk of
 28 concepts may be translated into talk of distinctions an agent can make in
 29 logical space.

30 Once we notice that rational agents can't always see all propositions,
 31 Paul's challenge arises. How can we make rational choices when we aren't
 32 in a position to entertain the possible outcomes of our actions?⁵

34 1.2 Epistemic decision theory

35 Paul's challenge is a problem for practical decision theory. The episte-
 36 mologized variant we're considering is a problem for epistemic decision
 37 theory.
 38

39 ⁵ Here, it might be that I'm formulating the challenge differently from how Paul (2015)
 40 sees it. On her view, the specification of outcomes won't determine facts about the agent's
 41 phenomenology. So an agent whose credences aren't defined over relevant phenomenological
 42 propositions will nevertheless be able to entertain the possible outcomes of her acts. But
 43 the agent's inability to conceive of the phenomenology of, e.g., having a child will prevent
 44 her from assigning utilities to outcomes in which she has a child. This difference may be
 45 substantive. Indeed, it may be that by shifting the focus away from phenomenology, what I
 46 call "Paul's challenge" would be better called "one of Paul's challenges."

1 Epistemic decision theory is an application of a Savage-style decision
 2 theory, restricting itself to epistemic “acts” and epistemic utilities. Epis-
 3 temic “acts,” typically not construed voluntaristically, involve possessing or
 4 coming to possess a credence function. Epistemic utility functions represent
 5 comparative epistemic goodness. Commonly defended epistemic decision
 6 rules:

7 **Dominance:** if credence function c has higher epistemic utility than
 8 credence function c' at every world, don't adopt c' .

9 **Maximize Expected Epistemic Utility:** adopt the credence function
 10 with the highest expected epistemic utility (i.e. the highest weighted
 11 average of epistemic utilities at all possible worlds, weighted by the
 12 probability of those worlds).
 13

14 As with ordinary decision theory, the relevant sort of goodness is ul-
 15 timate epistemic goodness. Epistemic goodness of this sort is objective,
 16 in the sense that it's non-information-dependent. A common view is that
 17 the relevant sort of ultimate epistemic goodness should be interpreted in
 18 terms of gradational accuracy.⁶ Gradational accuracy is the closeness of
 19 a credence function to the truth, by some measure satisfying a handful of
 20 intuitive constraints. Credence 1 in p is maximally close to the truth iff p
 21 is true; credence 0 is maximally close to the truth if p is false.

22 We can distinguish the epistemic utility of a particular credence an agent
 23 has in an individual proposition from the epistemic utility of the agent's
 24 total belief state. We'll call the former ‘local epistemic utility’ and the latter
 25 ‘global epistemic utility.’ It's usually assumed that global epistemic utility is
 26 determined straightforwardly as a function of local utilities.

27 Turning back to Paul's challenge: can epistemic decision theory be
 28 extended to credence functions defined over different propositions?
 29

30 2 Epistemic decision theory for partial credences

31 2.1 Motivating comparability

32 Epistemic decision theory usually presupposes that the credence functions
 33 it compares are defined over the same algebra of propositions. Once we
 34 abandon this presupposition, new difficulties arise.
 35

36 For example, to narrow the space of epistemic utility functions, con-
 37 straints such as ‘truth-directedness’ and ‘immodesty’ are placed on epistemic
 38 utility functions. Roughly, truth-directedness says credence functions are
 39 epistemically better the closer they are to the truth. Immodesty says that
 40 probabilistic credence functions assign themselves higher expected epistemic
 41 utility than all other credence functions. It's no longer clear what these
 42 constraints amount to, or why they are intuitive, once we compare credence
 43 functions that are defined over different domains. Should a probabilistic
 44

45 ⁶ Joyce 1998, 2009; Leitgeb and Pettigrew 2010a,b.
 46

1 partial credence function assign itself higher expected utility than proba-
 2 bilistic extensions of itself that see more propositions? Is the extension
 3 automatically closer to the truth than the original?

4 It's tempting to prescind from these questions. Epistemic decision theory
 5 wasn't designed to make comparative evaluations of credence functions that
 6 see different propositions. It falls silent about these sorts of comparison.
 7 So perhaps credence functions with different domains are incomparable in
 8 epistemic value.

9 Here are two arguments for why credence functions that can see different
 10 propositions should be comparable.

11 First: decision theory aims to be neutral with respect to substantive nor-
 12 mative questions. It provides only structural constraints on rational choices.
 13 Epistemic decision theory involves placing some substantive constraints
 14 on epistemic utility functions, but these are only meant to delineate the
 15 epistemic subject matter and need to be individually justified. There's no
 16 intuitive basis for assuming that credence functions with different domains
 17 can't be compared. So epistemic utility theory should accommodate the
 18 possibility that conceptual change has epistemic (dis)value.

19 Second: there are intuitive grounds for making at least some credence
 20 functions with different domains comparable. Here's an argument from ev-
 21 identialism: it's irrational to change your credence in a proposition without
 22 new evidence. One way of changing your credence in a proposition is to
 23 abandon it, so that your credence function no longer sees the proposition.
 24 So when evidence is held fixed, it's irrational to abandon your credence
 25 in a proposition. More briefly: it's irrational to undergo an "epistemic
 26 contraction." By contrast, it's plausible that it's at least sometimes rational
 27 to undergo an "epistemic expansion," whereby you retain your previous
 28 credences but come to see new propositions. (Intuitions are murkier about
 29 cases where you come to see new propositions and lose sight of old ones,
 30 and so I focus narrowly on easy cases: pure epistemic expansions and
 31 contractions.)

32 Let's consider stronger and weaker versions of these claims:

- 33 (1) It's always rationally impermissible to undergo an epistemic con-
 34 traction.
- 35 (2) It's sometimes rationally impermissible to undergo an epistemic
 36 contraction.
- 37 (3) It's always rationally permissible to undergo an epistemic expan-
 38 sion.
- 39 (4) It's sometimes rationally permissible to undergo an epistemic ex-
 40 pansion.

41 Of these four claims, all that is needed to establish the possibility of cross-
 42 algebra comparison is for one of the weaker claims [claim 2](#) and [claim 4](#) to
 43 be true. So here I rely on intuition. There are at least some circumstances
 44 where losing a credence in a proposition is irrational. In order for it to be
 45

possible to model the comparison between epistemic states before and after an epistemic expansion or contraction, credence functions with different domains must at least sometimes be comparable.

2.2 Strong and weak conceptual conservatism

For the purposes of this paper, I'll commit to [claim 2](#) and [claim 4](#). But in fact, all four have some intuitive plausibility, at least when the space of options is unconstrained. (When the space of options is constrained, there can be cases of forced choices between irrational credences and epistemic contractions, or cases where all the only optional expansions are irrational. In such cases, contractions may be permissible and expansions impermissible, contra [claim 1](#) and [claim 3](#). Note also that my discussion is confined to ideal rationality, where clutter avoidance and other resource constraints are non-issues.)

Consider [claim 1](#), which we can call 'strong conceptual conservatism'. A brief defense of strong conceptual conservatism:

Strong evidentialism: It is irrational to change your credences without acquiring new evidence.

No evidence against concepts: There can be no evidence that justifies losing conceptual resources.

Together these entail strong conceptual conservatism.

One might object: aren't there good reasons to abandon some concepts? For example, concepts that have false presuppositions? For example, once I realize that the concept *slut* has false misogynist presuppositions, isn't it best to abandon that concept altogether? In fact, though, it's better to retain the concept.

One might argue that the real objection to having the concept *slut* is a practical, ethical objection, not an epistemic one. Practical norms have no bearing on epistemic norms. (We are not epistemically required to have nonprobabilistic credences even when we know that doing so will magically save the life of a child.)

A better argument for retaining even problematic concepts, on my view, is that doing so is preferable both epistemically and practically. What's really wrong with objectionable concepts is not possessing them, but rather (in some sense) applying them. Suppose you know that Joe believes that Mary is a slut. If you abandon the concept of a *slut*, then you will no longer know exactly what Joe believes about Mary. If you want to talk him out of this misogynist belief, better to keep the concept *slut* and explain to him what's wrong with applying it.

Nothing in the present framework models a fully general distinction between possessing and applying a concept, but the idea should be roughly clear. Refraining from applying a concept might mean rejecting (having

credence 0) in all propositions that apply the concept. (One might wonder: how is that consistent with Bayesianism? Wouldn't you have to have credence 0 in both a proposition and its negation, thereby violating probabilism? Response: to make sense of concepts with false presuppositions in this sense, we'd have to move to a nonclassical setting.)⁷

So: problematic concepts don't generate a problem for strong conceptual conservatism. But there are other possible different angles of attack.

First, it might be that in some cases, if you face a choice between having credences that employ a problematic concept in an irrational or otherwise problematic way, or else losing the concept altogether, then perhaps you should choose the latter. For example, suppose it's psychologically impossible for you to retain the concept *slut* without lending positive credence to propositions that apply it. Then perhaps it's better just to lose the concept.

Here I think, the fan of strong conceptual conservatism will reply: if your credences are already irrational, then rationality doesn't recommend retaining them. Contracting your conceptual resources is still irrational, but it might be the lesser of two epistemic evils. What this objection is really an objection to is the stronger claim that in all cases, it's better to retain one's own credence function than to adopt any contraction of it. That is plausibly false, but not something that the strong conceptual conservative is committed to.⁸

A different type of objection to strong conceptual conservatism suggests that sometimes evidential considerations favor loss of concepts. For example, it might be that as our knowledge of the world develops, our concepts change; a naive concept is replaced by a more sophisticated one, or two separate concepts merge when they're revealed to be extensionally equivalent or analytically identical.

But these sorts of objections rely on a more inflated theory of conceptual resources than capacities to draw distinctions in logical space. Extensionally equivalent and analytically identical concepts don't draw different distinctions in logical space. Concept change over time may involve getting rid of concepts in some psychological sense, but it's not clear that it rationally requires abandoning distinctions in logical space.

⁷ Robbie Williams has pointed out to me that, in the case of concepts with false presuppositions, the problem may be more complex than I make out. On some views, each concept with a false presupposition determines a concept with the same application conditions that doesn't trigger the false presupposition; and so one can abandon the false presupposition while retaining the same distinctions in logical space by switching to the equivalent, presuppositionally innocuous concept. This suffices for conceptual conservation in the sense of "conceptual" I use throughout. But in the case of thick concepts like *slut*, this form of disentanglement may not be possible, and so other resources (e.g. rational but nonprobabilistic credences) may be necessary. With still other forms of problematic concept, we may have to tell a different story. Still, for the purposes of representing others' beliefs, or merely possible beliefs, or counterfactual or counterpossible scenarios, I suggest, strong conceptual conservatism is compelling.

⁸ Thanks to L. A. Paul, Julia Staffel, and Robbie Williams for discussion.

1 So these observations are in fact compatible with conceptual conser-
 2 vatism, in my (loose, perhaps unfortunate) sense of the word “conceptual.”
 3 There are, of course, interesting questions about the epistemic value of
 4 concepts in a more inflationary sense; but that’s not under discussion here.

5 A final objection to strong conceptual conservatism: it prohibits ever
 6 losing conceptual resources. This is a diachronic epistemic constraint. It’s
 7 controversial whether there are diachronic constraints on rationality.⁹

8 Opponents of diachronic constraints on rationality reject strong eviden-
 9 tialism, in favor of a weaker variant:

10 **Weak evidentialism:** It’s irrational to change your credences without
 11 some change in evidence.

12 Weak evidentialism allows changes in your credences if you either gain new
 13 evidence or lose old evidence, e.g. by forgetting information. On its own,
 14 weak evidentialism entails [claim 2](#), which we can call ‘weak conceptual
 15 conservatism.’ Indeed, it entails something stronger:

16 2’ ‘Medium-strength conceptual conservatism’: it’s rationally imper-
 17 missible to undergo an epistemic contraction without some change
 18 in evidence.
 19

20 We can equally provide arguments from [claim 4](#) to comparability. For
 21 example, suppose c is a partial credence function, and c^+ is an extension of
 22 c such that for all propositions A visible to c^+ and not c , $c^+(A)$ is maximally
 23 accurate (or otherwise has maximal epistemic value). Intuitively, c^+ must be
 24 as accurate as (or as valuable as) c . And so c and c^+ must be comparable.¹⁰

25 The upshot: evidentialism supports the conclusion that at least some
 26 credence functions with different domains can be compared.
 27

28 2.3 Impact on epistemic decision theory

29 What this means for epistemic decision theory: we need epistemic utilities,
 30 pr epistemic decision rules, that make at least some credence functions
 31 defined over different domains comparable.
 32

33 One natural temptation would be to be as neutral as possible with
 34 respect to the epistemic utilities of partial credence functions: for example,
 35 to assign partial credence functions only imprecise epistemic utilities. A
 36 neutral imprecise utility assignment for a partial credence function c would
 37 equal the set of global utilities of all total extensions of c .

38 How do we compare imprecise utilities? Perhaps the imprecise utility
 39 of c will be greater than the imprecise utility of c' iff the precise utilities of
 40 all total extensions of c were greater than the precise utilities of all total
 41

42
 43 ⁹ Against diachronic rationality, see [Talbot 1991](#), [Christensen 2000](#), [Williamson 2000](#),
 44 [Meacham 2010](#), [Moss Unpublished](#), [Hedden 2013](#).

45 ¹⁰ Thanks to an anonymous referee for suggesting this compelling intuition pump.
 46

1 extensions of c' . The problem is that this generates widespread incompara-
 2 bility. In particular, partial credences will always be incomparable with all
 3 of their extensions. The imprecise utility of a credence function will be a
 4 superset of the imprecise utility of its extensions.¹¹

5 We may be happy not to compare the utilities of credence functions
 6 defined over disjoint sets of propositions, or overlapping sets of proposi-
 7 tions where neither includes the other. But comparing credence functions
 8 and their extensions was supposed to be the easy case. An algebra-neutral
 9 epistemic utility function should be able to compare at least some par-
 10 tial credence functions and their extensions. Otherwise we can't predict
 11 conceptual conservatism.

12 It would be hasty to rule out imprecise utilities altogether. But I've
 13 argued that epistemic decision theory should allow at least some credence
 14 functions over different domains to be comparable, including at least some
 15 credence functions and their extensions. So either way, we can't avoid
 16 substantive epistemological questions about what constraints there are on
 17 epistemic utility functions that range over credence functions with different
 18 domains.

20 3 Epistemic utility functions for partial credences

21
 22 Instead of retaining a utility function that ranges only over total credence
 23 functions, and assigning partial credence functions imprecise utilities, I
 24 suggest we consider epistemic utility functions that range over (at least
 25 some) credence functions with different domains.

26 Which utility functions? Instead of positing unique epistemic utility func-
 27 tions, epistemic utility theory generally proceeds by imposing constraints on
 28 the space of candidate epistemic utility functions, constraints which suffice
 29 represent epistemic norms of various sorts. We can think of principles like
 30 strong and weak conceptual conservatism as constraints on the space of
 31 epistemic credence functions, constraints that encode the epistemic value of
 32 conceptual resources. There may be other intuitive principles linking the
 33 utilities of partial credence functions and their extensions.

34 Other constraints on algebra-neutral epistemic utility functions may be
 35 natural generalizations of accepted constraints on algebra-specific functions.
 36 For example, in place of strict propriety (the constraint that algebra-specific
 37 utility functions should make probabilistic credence functions assign them-
 38 selves higher expected utility than all alternatives), an algebra-neutral
 39

40 ¹¹ Obviously, we might use a different rule for comparing imprecise credences: for example,
 41 perhaps c is strictly preferable to c' iff the maximal utility in c is greater than the maximal
 42 utility in c' and the minimal utility in c is greater than the minimal utility in c' . Then we
 43 might generalize: c is weakly preferable to c' iff the maximum utility of c is greater than or
 44 equal to the maximum utility of c' and the minimum utility of c is greater than or equal to
 45 the minimum utility of c' . Even then, we can never predict that an extension of c is strictly
 46 preferable to c , and so we can't predict weak conceptual conservatism.

1 epistemic utility function should perhaps satisfy a generalization of strict
 2 propriety that makes probabilistic partial credence functions assign them-
 3 selves higher expected utility than all alternatives defined over the same
 4 partition.

6 3.1 Local and global utilities

7
 8 In order to extend epistemic decision theory to partial credence functions,
 9 we need to look at how adding or subtracting individual credences affects
 10 the epistemic utility of an agent's total epistemic state. In other words, we
 11 need to look at the relationship between the local utilities of credences in
 12 individual propositions and the global utilities of overall credence functions.

13 Global epistemic utility is usually interpreted (noncommittally) as a sum
 14 or average of local epistemic utilities. Where g and l are global and local
 15 utility functions, respectively, and \mathcal{A} is the algebra of propositions c is
 16 defined over,

17 **Summing proposal:**

$$18 \quad g(c, \omega) = \sum_{A \in \mathcal{A}} l(c, A, \omega)$$

19
 20 **Averaging proposal:**

$$21 \quad g(c, \omega) = \frac{1}{|\mathcal{A}|} \sum_{A \in \mathcal{A}} l(c, A, \omega)$$

22
 23 Now, epistemic utility is standardly understood as gradational accuracy.
 24 Accuracy-first epistemic decision theory assigns credences value in terms
 25 of their distance from the truth, by some measure satisfying a handful of
 26 intuitively plausible constraints.

27
 28 So we can ask: is an accuracy-first epistemology compatible with treating
 29 conceptual resources as epistemically valuable? The answer can be *yes* only
 30 if the epistemic value of conceptual resources is reducible to the epistemic
 31 value of accuracy. As it turns out, both the summing hypothesis and the
 32 averaging hypothesis for global (in)accuracy introduce commitments about
 33 the value of conceptual resources. These commitments potentially are
 34 unattractive.

35 We could, of course, switch to epistemic utility functions that aren't
 36 accuracy-directed. Indeed, we might be forced to do so in order to avoid
 37 implausible conclusions about the value of conceptual resources. But it's
 38 worth exploring whether we can represent the value of conceptual resources
 39 without departing to far from accuracy-first epistemology, especially be-
 40 cause the research project has otherwise proven fruitful.

41 42 3.2 Summing positive disutilities

43
 44 In accuracy-first epistemic decision theory, local inaccuracy is interpreted as
 45 a positive penalty for distance from the truth. There is a maximum degree
 46

1 of accuracy, which is distance 0 from the truth. Any credence other than 1
 2 in truths or 0 in falsehoods incurs a positive disutility.

3 And so summing local disutilities generates an immediate consequence: if
 4 c^+ is an extension of c and any of the new credences it brings are uncertain,
 5 c^+ will incur whatever disutilities c has plus disutilities for its new credences.
 6 So c^+ will automatically have greater global disutility c . In other words,
 7 c dominates c^+ , merely because c^+ can make new distinctions but isn't
 8 omniscient about them.

9 So on this proposal, we find ourselves with the result that the fewer im-
 10 perfect credences you have, the better you are epistemically. This amounts
 11 to treating nonattitudes toward propositions as epistemically *perfect*: maxi-
 12 mally accurate. They are matched only by the epistemic utility of credence
 13 1 in truths and credence 0 in falsehoods.

14 This proposal, paired with either weak dominance avoidance or expected
 15 utility maximization, yields the following verdicts: rational agents only
 16 have credences in propositions such that their credence matches the truth
 17 value at every world (i.e. \top and \perp and any other known necessities).
 18 Specifically, rational agents will have credence 1 in all tautologies, credence
 19 0 in all contradictions, and no other defined credences. And of course,
 20 similar problems afflict any other, non-accuracy-based interpretations of
 21 global epistemic disutility that treat global disutilities as sums of positive
 22 disutilities.

23 I take it as a datum that epistemic utility functions shouldn't universally
 24 prohibit attitudes toward contingent propositions, or propositions the agent
 25 can't be certain of.

27 3.3 Averaging

28 What if instead of treating the global (dis)utility of a credence function as
 29 the sum of its local (dis)utilities, we treated it as the average?

30 Then there is no automatic dominance relation between credence func-
 31 tions and all of their extensions. Sometimes seeing new propositions will
 32 increase global utility, sometimes decrease it. At some worlds, new un-
 33 certain credences may increase your average accuracy; at others decrease
 34 it.
 35

36 The averaging proposal seems like an obvious move. But it brings
 37 with it new problems. Suppose you have credence function c , which sees
 38 propositions A and $\neg A$ and assigns .8 in A , which is true. Then suppose
 39 you have the option of extending your credence function to c^+ , which also
 40 sees new propositions B and $\neg B$, and has credence .6 in the true one.¹²

41 Two consequences. First, adopting c^+ entails a reduction in your global
 42 accuracy at the actual world. Your old credences in A and $\neg A$ were pretty
 43 accurate. Your new credences in B and $\neg B$ are on the right track, but

44 ¹² To keep the example as simple as possible, I don't stick to my assumption that credences
 45 must defined be over the boolean closure of a partition.
 46

1 they're still not as close to the truth as your old credences were. So they
2 drag the average down.

3 Second, adopting c^+ will typically entail a reduction in your expected
4 global accuracy. After all, you are more confident than not that you're at a
5 world where your average accuracy is dragged down.¹³

6 Two objections. First, it's not clear that this is an intuitive way of
7 characterizing the overall accuracy of a credence function. Second, if we
8 accept this as a characterization of the accuracy of a credence function,
9 then accuracy is a bad measure of epistemic value.

10 On the first point: it's clear that your average local accuracy is reduced
11 when you move from c to c^+ . But it's not obvious that your global accuracy
12 should be reduced. There is some intuitive sense in which c^+ is doing
13 better at the actual world, accuracywise, than c . For one thing, c couldn't
14 distinguish B from $\neg B$. c^+ not only distinguishes them, but is closer to truth
15 than to falsehood. So this example may motivate rejecting the equivalence
16 of global accuracy with average local accuracy.

17 On the second, less controversial point: if global accuracy is average
18 local accuracy, then it's implausible that global accuracy is a good charac-
19 terization of epistemic value.

20 In this example, suppose our agent is an expected utility maximizer and
21 epistemic utility is simply average accuracy, measured by the most familiar
22 accuracy measure, the Brier score.¹⁴ Then if she has c^+ as her credence,
23 she'll prefer to abandon her attitudes toward B and its negation. c has
24 higher expected accuracy from c^+ 's perspective than c^+ itself.

25 The problem doesn't just afflict expected utility maximizers. Even re-
26 stricting ourselves to strong accuracy dominance avoidance, we end up at
27 the bottom of the same slippery slope. If the agent has attitudes towards
28 propositions that she's uncertain of, then her average accuracy is imperfect.
29 But if she only has attitudes toward propositions that she's certain of—
30 again, tautologies and contradictions—then her attitudes will have perfect
31 average accuracy, and hence will be more accurate at every possible world.
32 So in order to avoid having credences that are strongly dominated, she
33 must restrict her credences to propositions that she's certain of. In order to
34 avoid having credences that are weakly dominated, she must restrict her
35 credences to propositions that are necessary across all possibilities that she
36 can entertain.

40 ¹³ This depends on the choice of local utility function or accuracy measure, since the loss of
41 average accuracy at some worlds may be offset by a more dramatic gain in average accuracy
42 at other worlds. But for any continuous, truth-directed accuracy measure, we can generate an
43 analogous case where the expected average accuracy will decrease with the addition of new
44 credences.

45 ¹⁴ The Brier score of a probability is the squared Euclidean distance between the probability
46 assigned to a proposition and its truth value (1 = true and 0 = false).

1 And so, like the summing disutilities proposal, I think we have good
 2 reason to reject an algebra-neutral epistemic utility function that derives
 3 global utilities from averages of local utilities.

4 5 **3.4 Summing positive utilities**

6 Let's return to the summing proposal. What if instead of penalizing distance
 7 from truth, we reward distance from falsehood? That is, what if we derive
 8 global utility by summing positive local utilities? Then, of course, the
 9 situation is reversed with respect to summing positive disutilities. Each
 10 new proposition added to the domain of a credence function increases
 11 the credence function's epistemic value, as long as the credence it assigns
 12 isn't maximally inaccurate. This means treating nonattitudes towards
 13 propositions as maximally inaccurate—just as bad, from an epistemic point
 14 of view, as certainty of falsehood.

15 This is not counterintuitive in the way that summing disutilities is.
 16 There's a perfectly reasonable position in logical space according to which
 17 any increase in conceptual resources contributes positively to epistemic
 18 utility. It says that conceptual resources are a pure epistemic good: they
 19 trump any epistemic badness that might be required to achieve them. So it
 20 automatically gains us strong conceptual conservatism.

21 Whereas summing disutilities generated a dominance argument against
 22 having credences in any contingent propositions, summing positive utilities
 23 generates a dominance argument for having credences in every proposition.
 24 One might object: doesn't this contradict the premise of this paper: that
 25 there can be rational agents with partial credence functions? Doesn't this
 26 entail that before she sees red, Mary is irrational?

27 Answer: having dominated credences is only irrational if there are any
 28 non-dominated credences that are among one's epistemic options. In Mary's
 29 case, though, adopting credence functions that can conceive of what it's
 30 like to see red is not an option for her. She would need to experience
 31 phenomenal redness before she'd be in a position to adopt that credence
 32 function. But she can't: she's locked in a black room. So she's not irrational
 33 for having a credence function that's dominated only by a non-option, any
 34 more than you are irrational for not spontaneously acquiring a billion
 35 dollars.

36 It does mean, though, that if your credence function is needlessly partial,
 37 then you're irrational. If you simply fail to have credences in some possi-
 38 bilities, even though you're perfectly well in a position to do so, then this
 39 view judges you irrational. Two reasons to think this is not a problem: first,
 40 it's presumably controversial whether this is even psychologically possible.
 41 Second, if it is, it's not obvious that it should be rational.

42 A different kind of objection: suppose there were a way to make up new
 43 concepts. For example, you might give your friend Rachel a separate name,
 44 "Srachel," just for Tuesdays. Then you might conceptually distinguish
 45

1 Srachel (who went to karaoke on Tuesday) from Rachel (who gave a talk
2 on Wednesday). Even though you know they're the same person, perhaps
3 you can now conceive of the possibility that they're two different people.

4 The details of the example don't really matter; we simply need some way
5 for a person to be able to generate new concepts. Then, the proposal seems
6 to suggest that doing so is epistemically mandatory. That's implausible.

7 First, it's not clear that the act of voluntarily inventing a new concept
8 is an epistemic act, as opposed to a practical one. Second, it's not clear
9 that in introducing new concepts of this sort, you're really exposing any
10 distinctions in logical space. (After all, you plausibly already have the
11 concept of *rachel on tuesdays*; possessing that concept was instrumental
12 in your invention of the *scrachel* concept. Third, if somehow you are
13 exposing real distinctions in logical space, then it's perhaps doing so is
14 epistemically preferable after all. The only obvious reason against it is
15 clutter avoidance; but for ideally rational agents I take it that epistemic
16 clutter is unproblematic.

17 Summing positive utilities does have some questionable consequences,
18 though. For example, suppose a partial credence function c has an exten-
19 sion, c^+ . Then c^+ dominates c , regardless of what credences it assigns to
20 new propositions. c^+ could be irrational in lots of ways: it might assign
21 credence 1 to both a proposition and its negation. Still, c^+ would dominate
22 c , even if c were probabilistically unimpeachable.

23 This may not be a deep problem. If c^+ isn't probabilistic, then it will
24 itself be dominated by some probabilistic credence function defined over the
25 same algebra.¹⁵ And so as long as an agent's epistemic options include all
26 those credence functions defined over subsets of the propositions potentially
27 visible to the agent, the summing utilities proposal will never require one
28 to adopt nonprobabilistic credences.

29 Still: it's not obvious that seeing new propositions in an irrational way is
30 necessarily epistemically better than not seeing them at all, even when both
31 are suboptimal.

32 There are other reasons why one might not be satisfied with this proposal.
33 Because it treats nonattitudes toward propositions as having the lowest
34 possible epistemic utility, expanding one's epistemic vision is an absolute
35 epistemic good. One might doubt this: for example, one might think that
36 having a very inaccurate credence in a proposition is worse, at a world,
37 than not having any credence in that proposition.

38 I don't want to rule out the summing positive utilities proposal. It has
39 many intuitively appealing features, including its preservation of conceptual
40 conservatism. But these considerations suggest that it's worth exploring
41 other options. There's room in first-order epistemology for controversy
42

43 ¹⁵ Assuming our algebra-neutral utility functions preserve the probabilism-entailing properties
44 of algebra-specific epistemic utility functions when comparing credence functions that share
45 an algebra.
46

1 about whether new conceptual resources are always an absolute epistemic
 2 good, or whether their utility can be outweighed.

3.5 “Better than chance”?

5 So far, a pattern has emerged. When we sum positive disutilities, in effect
 6 we treat nonattitudes toward propositions as if they had maximal epistemic
 7 utility. When we sum positive utilities, in effect we treat nonattitudes
 8 toward propositions as if they had maximal epistemic disutility. When we
 9 average local (dis)utilities, in effect we treat nonattitudes as if they had the
 10 same utility as the average utility of attitudes.

11 Instead of treating nonattitudes as though they all having maximal utility
 12 or minimal utility, it seems plausible that we should treat nonattitudes as
 13 having middling utility. The averaging proposal accomplishes that, but not
 14 in the right way.

15 One might be tempted to say: nonattitudes toward propositions are no
 16 closer to truth than to falsehood. And so they should have the same utility
 17 as credences that are no closer to truth than falsehood, i.e. credence .5.

18 Something like this might be on the right track. But suppose Mary has
 19 no attitude toward the propositions that seeing red is like experiencing
 20 phenomenal redness, or like experiencing phenomenal blueness, or like
 21 experiencing phenomenal azureness, or Suppose Mary’s partial
 22 credence function has the same utility as an extension of itself that assigns
 23 credence .5 in *all* of these propositions. Then it has the same utility as a
 24 wildly non-probabilistic credence function. This will generate unpredictable
 25 consequences for when a very inaccurate credence function has higher
 26 utility than locally very accurate but partial credence function. In any
 27 case, this certainly makes partial credence functions dominated (since all
 28 non-probabilistic credence functions are dominated, assuming familiar
 29 constraints on epistemic utility functions). This is a substantive, non-
 30 obvious epistemological hypothesis.¹⁶

31 To avoid this problem, the local utility of nonattitudes need not be
 32 the same for both truth and falsehood. That does satisfy a plausible
 33 desideratum: that the global utility of a partial credence function is uniform
 34 across worlds it doesn’t distinguish. But the proposal is stronger than
 35 necessary: one might satisfy the same desideratum without assuming that a
 36 nonattitude’s local utility is the same at worlds where it’s true and worlds
 37 where it’s false. Local utilities may differ as long as it all evens out at
 38 the global level. Let’s consider another sort of proposal that generates
 39 uniformity in global utilities across indistinguishable worlds.

40 Consider again the example in the last subsection. A fan of accuracy-first
 41 epistemology might reason: even though c^+ brings down average local
 42 accuracy, it seems to be doing pretty well, accuracywise. After all, it assigns
 43 .6 credence in B , a truth; so it’s closer to truth than falsehood. One might

44
 45 ¹⁶ Thanks to Kenny Easwaran for prompting me to consider this proposal and for discussion.

1 say, intuitively, c^+ is doing “better than chance”. Maybe that’s the sense in
 2 which it seems to be an improvement in accuracy over having a nonattitude.

3 And so, one might suppose nonattitudes toward propositions are worse
 4 than credences that are more accurate than chance, but better than cre-
 5 dences that are less accurate than chance. The credences that are neither
 6 worse nor better than chance are the maximally unopinionated attitudes. So,
 7 on this view, nonattitudes have the same utility as maximally unopinionated
 8 attitudes.

9 Spelling this out: suppose each partial credence function c has a unique
 10 maximally unopinionated total extension: call this credence function $c^{\cdot|}$.
 11 On the hypothesis we’re considering, the utility of c ’s nonattitude toward
 12 A is equal to the utility of $c^{\cdot|}$ ’s attitude toward A .

13 This hypothesis has a number of attractive features. Unlike summing
 14 disutilities or averaging, there is no automatic epistemic gain in seeing fewer
 15 propositions. There is also no automatic epistemic gain in seeing more
 16 propositions: this view allows that whatever epistemic good there is in
 17 seeing new propositions, it can be outweighed.

18 Another feature of this view is that it ensures weak conceptual conser-
 19 vatism, and comes close to ensuring strong conceptual conservatism. If c
 20 is probabilistic, then so is $c^{\cdot|}$. And so if c ’s utility matches $c^{\cdot|}$ ’s, then with
 21 standard (strictly proper) scoring rules, c will be non-dominated and will
 22 maximize expected utility from its own perspective. So it will never be rati-
 23 onally required to give up having any credence at all in some propositions.
 24 And, if the agent is an expected utility maximizer, she’ll prefer her own
 25 partial credence function to almost any alternative the domain of which is
 26 a subset of hers.¹⁷

27 But this hypothesis does face standard symmetry worries. It depends
 28 on the existence of a unique, maximally unopinionated extension of c .
 29 But problems for the principle of indifference suggest that there isn’t any
 30 objective basis for determining a unique maximally unopinionated $c^{\cdot|}$.¹⁸ In
 31 other words, there is no objective way of isolating a uniform distribution of
 32 probabilities c_0 over \mathcal{W} such that $c^{\cdot|}$ is (roughly) c_0 updated on c ’s evidence.

33 This is obviously a legitimate worry. And it isn’t easy to avoid: for this
 34 style of proposal, we need maximal unopinionatedness to ensure that the
 35 global utility of c is uniform over worlds that c can’t distinguish.¹⁹ It is
 36

37 ¹⁷ It will, however, be epistemically permissible for an agent to give up credence in a proposition
 38 if her credence is already maximally unopinionated (and if giving up that credence doesn’t
 39 violate the requirement that credences be defined over a boolean algebra).

40 ¹⁸ See van Fraassen 1989.

41 ¹⁹ This rules out a natural generalization. Objective Bayesianism—the view that there’s a
 42 unique rational prior credence function—was originally envisioned as recommending absolute
 43 unopinionatedness, i.e. the uniform distribution. In the face of symmetry worries, contempo-
 44 rary objective Bayesians typically think that the rationally privileged prior need not be the
 45 uniform distribution over worlds. Suppose we said that the utility of a nonattitude toward
 46 a proposition at a world is equal to the utility of the credence assigned by the rationally
 privileged prior, updated on the partial credence holder’s evidence. If the prior is not uniform,

1 plausibly a desideratum of an algebra-neutral epistemic utility functions
 2 that if a partial credence function doesn't distinguish two worlds, then the
 3 credence function should have the same epistemic utility at both worlds.

4 If we like this proposal, we might have to accept that there's some
 5 arbitrariness or assessment-sensitivity in the assignment of epistemic utilities.
 6 The motivating thought behind this proposal: having a credence in a
 7 proposition is better (at a world) than having no attitude toward it if the
 8 credence assigned is "better than chance": if one's credence is closer to the
 9 truth than if one were to withhold judgment as much as possible. What
 10 amounts to doing "better than chance" depends here on a conception
 11 of chance as a uniform distribution over indistinguishable worlds, and
 12 so presupposes a specific space of worlds. The relevant space of worlds
 13 depends on the perspective of the theorist.

14 4 Conclusion

15
 16
 17 We've seen a variety of accuracy-based proposals for algebra-neutral epis-
 18 temic utility functions. The most natural extensions of traditional epis-
 19 temic utility functions—summing local disutilities and averaging local
 20 (dis)utilities—both generate terrible consequences. They require agents
 21 to give up any credence in propositions that the agent can't be certain of.
 22 And so they violate both strong and weak conceptual conservatism with
 23 respect to epistemically contingent propositions.

24 The other two proposals we've seen don't face this problem. Summing
 25 positive local utilities entails strong conceptual conservatism. Any loss of
 26 concepts will generate a loss of epistemic utility and so will be dominated.
 27 And the final proposal—where partial credence functions' utilities match
 28 that of their maximally unopinionated extensions—entails at least weak
 29 conceptual conservatism.

30 It's obvious that my discussion has been far from exhaustive. There
 31 may be other plausible ways of generating algebra-neutral epistemic utility
 32 functions that still make epistemic utility a function of accuracy. There
 33 are certainly other plausible epistemic utility functions that don't rely on
 34 accuracy alone. It is not the ambition of this paper to narrow down the
 35 space of epistemic utility functions to one or the other of these proposals.

36 What I hope to have made clear, however, is the need for certain con-
 37 straints on algebra-neutral epistemic utility functions. First, they must
 38 permit rational agents to have some uncertainty, and to have attitudes
 39 toward contingent propositions. Second, they must permit rational agents
 40 to have attitudes toward contingent propositions. Neither of the traditional
 41

44 though, then it might assign different credences to worlds that the partial credence function
 45 under evaluation can't distinguish. And so its utility at those worlds might differ.
 46

ways of aggregating local inaccuracies have satisfied these constraints. Furthermore, there is some intuitive support for stronger constraints: perhaps weak conceptual conservatism; perhaps strong conceptual conservatism.

Jennifer Carr

References:

- Christensen, David. 2000. "Diachronic Coherence Versus Epistemic Impartiality." *Philosophical Review* 109 (3): 349–371.
- Hedden, Brian. 2013. "Options and Diachronic Tragedy." *Philosophy and Phenomenological Research* 87 (1).
- Jackson, Frank. 1982. "Epiphenomenal Qualia." *The Philosophical Quarterly* 32: 127–136. <http://dx.doi.org/10.2307/2960077>.
- Joyce, James M. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65 (4): 575–603.
- Joyce, James M. 2009. "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief." In *Degrees of Belief*, edited by Franz Huber and Christoph Schmidt-Petri, Vol. 342, 263–297. Dordrecht: Springer.
- Leitgeb, Hannes and Richard Pettigrew. 2010a. "An Objective Justification of Bayesianism I: Measuring Inaccuracy." *Philosophy of Science* 77 (2): 201–235.
- Leitgeb, Hannes and Richard Pettigrew. 2010b. "An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy." *Philosophy of Science* 77 (2): 236–272.
- Meacham, Christopher J. G. 2010. "Unravelling the Tangled Web: Continuity, Internalism, Non-Uniqueness and Self-Locating Beliefs." *Oxford Studies in Epistemology* 3: 86–125.
- Moss, Sarah. Unpublished. "Credal Dilemmas."
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Paul, L. A. 2015. "What You Can't Expect When You're Expecting." *Res Philosophica* 92 (2).
- Talbott, W. J. 1991. "Two Principles of Bayesian Epistemology." *Philosophical Studies* 62 (2): 135–150.
- van Fraassen, Bas. 1989. *Laws and Symmetry*. Oxford: Oxford University Press.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.