Comparative Learning

Abstract

In contemporary epistemology, there are two dominant approaches to explicating the norms of diachronic rationality. The first approach is articulated in terms of qualitative 'all or nothing' beliefs, and is concerned with identifying the diachronic norms of *qualitative belief revision*. The second approach is articulated in terms of numerically graded credences, and is concerned with identifying the diachronic norms of *credal update*. Here, I present and develop a third (hitherto neglected) paradigm for explicating the norms of diachronic rationality – namely, a paradigm articulated in terms of *comparative confidence judgements* of the form 'I am more confident in the truth of p than I am in the truth of q', or 'I am equally confident in the truth of p and q'. Specifically, I address the question 'how should a rational agent revise their comparative confidence judgements over time as they gather new evidence'. I begin by introducing and motivating a novel and intuitively compelling update rule, called 'comparative conditionalisation' (CC), that generalises Bayesian conditionalisation to the comparative setting. I then go on to (i) identify several key characteristic properties of this rule, and (ii) demonstrate that CC can be *partially* justified by a plausible diachronic Dutch book argument, before (iii) showing that, contrary to expectations, CC does *not* minimise expected inaccuracy under the most natural formalisation of epistemic utility for comparative confidence judgements – and hence that *there are learning scenarios in which it is impossible for agents to simultaneously minimise the expected inaccuracy of both their credences and their comparative confidence judgements.*

1 Introduction

We humans are prone to believing things, like when I believe that the moon is made of cheese. We are also prone to lending credence to things, like when I lend a credence of around 0.1 to there being rain in Windhoek tomorrow. Finally, we are also prone to making comparative confidence judgments, like when I judge that it is more likely that the moon is made of cheese than it is that it will rain in Windhoek tomorrow. While epistemic attitudes of the first two kinds (qualitative belief and numerically graded credence) are widely taken to play a crucial role in framing the fundamental norms by which the rationality of an agent's epistemic states are to be assessed, comparative confidence judgements have attracted much less attention in the contemporary philosophical literature. This is somewhat surprising, given that several eminent figures in the history of probability – i.e. Keynes (1921), de Finetti (1937, 1951), Koopman (1940) and Fine (1973) – have contended that comparative confidence judgements are the most fundamental, intuitive and psychologically basic of all our epistemic attitudes.[1]

Over the years, numerous authors have attempted to identify synchronic rationality requirements for comparative confidence orderings (see e.g. Halpern (2003) for a thorough overview). However, the philosophical foundations of this project have, until recently, been largely neglected,[2] and there is still little consensus regarding what kinds of comparative confidence structures are characteristic of rational agents. Happily, this situation is beginning to improve. Icard (2016) has shown that 'money-pump' style arguments can be used to provide a prospective pragmatic justification of the requirement that a rational agent's comparative confidence judgments should always be representable by a probability function. Meanwhile, Fitelson and McCarthy (unpublished) have shown that accuracy dominance arguments can be used to provide epistemic justifications for some significantly weaker synchronic coherence requirements (such as the principle that a rational agent's comparative confidence judgments should always be representable by a Dempster-Schafer belief function).

But despite recent progress in identifying the synchronic coherence constraints that govern the comparative confidence judgments of rational agents at a time, practically nothing has been written on the question of how rational agents should change their comparative confidence judgments over time as they gather new evidence.[3] This is the problem with which I'll be concerned in this paper.

Before moving on, it is worth briefly clarifying an important point. My aim in this paper is *not* to justify the claim that the epistemic states of ideally or boundedly rational agents should be conceived

[1] Thus, we read, for example,

> The fundamental viewpoint of the present work is that the primal intuition of probability expresses itself in a (partial) ordering of eventualities: A certain individual at a certain moment considers the propositions a, b, h, k,...Then the phrase 'a on the presumption that h is true is equally or less probable than b on the presumption that k is true' conveys a precise meaning to his intuition... This is, as we see it, a first essential in the thesis of intuitive probability, and contains the ultimate answer to the question of the meaning of the notion of probability. (Koopman, 1940: 270)

[2] See Fine (1973) for a critical assessment of the philosophical motivations behind several synchronic rationality requirements from the literature.

[3] At this point, the reader might wonder whether authors such as Hawthorne (2016) and Norton (forthcoming) have not already addressed this question with their qualitative axiomatisations of the notion of inductive support. But there is actually a subtle but important distinction between (i) formalising the notion of inductive support in terms of comparative confidence judgements, and (ii) justifying a procedure for revising comparative confidence judgements on the basis of novel evidence. While the first question has already been addressed in the literature, the second question, which is importantly distinct from the first, has been completely neglected. Furthermore, it should be noted that i.e. Hawthorne (2016) assumes a battery of synchronic coherence constraints about whose normative status I remain neutral here.

of in terms comparative confidence judgements rather than qualitative beliefs or (precise or imprecise) numerical credences. Rather, I assume in the background that there are at least some scenarios in which such a conception is principled, and then consider the question of how epistemic states, thus conceived, should evolve over time. After all, it is surely at least possible to conceive of a creature whose epistemic state is characterised purely by comparative confidence judgements, and it is surely philosophically interesting to ask what kinds of epistemic norms would determine the rationality of such a creature's reasoning. As I mentioned above, the comparative conceptualisation of epistemic states has numerous illustrious champions, and I mainly take it for granted that the reader will agree that a proper understanding of the dynamics of comparative confidence judgements is a worthy philosophical goal.

The structure of the paper is as follows. In section 2, I introduce the standard formalism for analysing comparative confidence judgments and provide a concise summary of some of the most important synchronic coherence constraints from the literature, before briefly reviewing some recent work regarding the normative status of these constraints. In section 3 I turn to the central question of the paper: 'how should a rational agent revise their comparative confidence judgements over time as they acquire new evidence?'. I address this question by studying the way in which a Bayesian agent's comparative confidence judgements change when they conditionalize on new evidence. I then show that the resulting revision rule (which I call 'comparative conditionalisation' (CC)) is intuitively compelling, even outside of the context of probabilistic Bayesian epistemology, and establish some basic properties of CC, before going on to illustrate two important senses in which the comparative rule requires less epistemic structure for its application than its numerical counterpart. The aim of the subsequent sections is to investigate whether the normative arguments that are normally given as justifications for Bayesian conditionalisation can be generalised to provide normative justifications for comparative conditionalisation. In section 4 I construct and evaluate a prospective pragmatic diachronic Dutch book argument for CC, and catalogue the decision rules that are needed to get such an argument off the ground. In section 5, I show that the most influential *epistemic* justification of Bayesian conditionalisation, namely the argument from expected inaccuracy, cannot be straightforwardly generalised to the comparative setting, since CC sometimes fails to minimise expected inaccuracy according to the only extant scoring rule for comparative confidence orderings. This implies that *it is sometimes impossible for agents with probabilistic credences to minimise the expected inaccuracy of both their posterior credences and their posterior comparative confidence judgements at the same time.* Section 6 draws some morals, describes the available avenues for escaping this puzzle, and concludes.

2 Coherence Conditions for Confidence Orderings

2.1 Preliminaries

I begin with some technical preliminaries. Firstly, I assume that agents always make comparative confidence judgements about 'propositions' drawn from the Boolean algebra $\mathfrak{B}$ of equivalence classes of logically equivalent sentences of some language $\mathcal{L}$.[4] Intuitively, an agent A can make two kinds of comparative confidence judgement about propositions in $\mathfrak{B}$. Firstly, they can be strictly more confident in the truth of p than they are in the truth of q. I denote this kind of judgement with the notation $[p \succ q]$. Alternatively, A can be equally confident in the truth of p and q. I denote this second kind of judgement with the notation $[p \sim q]$.[5] Together, the set of all A's comparative confidence judgements define a *confidence ordering*, $\succsim$, over some subset of the propositions in $\mathfrak{B}$, defined by $p \succ q$ if and only if $[p \succ q]$, and $p \sim q$ if and only $[p \sim q]$. I write $p \succsim q$ to indicate the disjunction '$p \succ q$ or $p \sim q$'. Following the standard approach in the literature, I make the following basic assumptions about $\succsim$.

Opinionation: For any $p, q \in \mathfrak{B}$, A makes exactly one of the judgements $[p \succ q]$, $[q \succ p]$, $[p \sim q]$, i.e. one of $p \succ q$, $q \succ p$ and $p \sim q$ is true.

Opinionation implies that $\succsim$ is a 'total ordering' of $\mathfrak{B}$.[6] Intuitively, it means that there are 'no gaps' in A's confidence judgements, i.e. A makes a comparative confidence judgement about every pair of propositions in $\mathfrak{B}$. This assumption, though controversial, is standard in the extant literature on comparative confidence orderings.[7] Next, I assume that $\succ$ satisfies the following conditions.

Irreflexivity of $\succ$: For every $p \in \mathfrak{B}$, A does not make the judgement $[p \succ p]$, i.e. $p \not\succ p$.

Transivity of $\succ$: For every $p, q, r \in \mathfrak{B}$, if A makes the judgements $[p \succ q]$ and $[q \succ r]$, then A makes the judgement $[p \succ r]$, i.e. if $p \succ q$ and $q \succ r$, then $p \succ r$.

Finally, I assume that $\sim$ is an equivalence relation, i.e.

---

[4]For simplicity, I assume that $\mathfrak{B}$ and L are always finite. The assumption that the relata of comparative confidence judgements are logical equivalence classes rather than simple sentences can be seen as a logical omniscience assumption, i.e. that the agent is always aware of all logical equivalences.

[5]To be clear, $[p \sim q]$ denotes the judgement that p and q are equally plausible. Depending on one's view of epistemic indifference, this may or may not be distinct from simply being epistemically indifferent between p and q. See Eva (forthcoming) for a discussion of comparative conceptions of epistemic indifference.

[6]Fitelson and McCarthy (unpublished) work in a more general setting than that described here. Specifically, they consider an agent's comparative confidence over arbitrary (possibly proper) subsets of $\mathfrak{B}$, which they call 'agendas'. They then assume only that $\succsim$ is opinionated with respect to the given agenda.

[7]For philosophical critiques of the Opinionation assumption, see e.g. Keynes (1921) and Forrest (1989). One might plausibly contend that one of the primary advantages of conceiving of an agent's epistemic states in terms of comparative confidence judgements rather than numerical credences or qualitative beliefs is that it allows us to study the epistemological consequences of failures of Opinionation. As I note in Section 3, the update rule considered here is easily extended to the non-opinionated setting, but I leave the important problem of systematically *justifying* this extension for another day.

Reflexivity of $\sim$: For every $p \in \mathfrak{B}$, A makes the judgement $[p \sim p]$, i.e. $p \sim p$.

Transivity of $\sim$: For every $p, q, r \in \mathfrak{B}$, if A makes the judgements $[p \sim q]$ and $[q \sim r]$, then A makes the judgement $[p \sim r]$, i.e. if $p \sim q$ and $q \sim r$, then $p \sim r$.

Symmetry of $\sim$: For every $p, q \in \mathfrak{B}$, if A makes the judgements $[p \sim q]$, then A makes the judgement $[q \sim p]$, i.e. if $p \sim q$, then $q \sim p$.

When all of these assumptions are satisfied, I say that the ordering $\succsim$ is a 'total preorder' over $\mathfrak{B}$. For the remainder of the article, I will assume that the confidence orderings being considered are total preorders over $\mathfrak{B}$, unless otherwise stated. The following additional constraint is also sometimes assumed (see i.e. Fitelson and Mccarthy (unpublished)):

Regularity of $\succsim$: For any contingent $p \in \mathfrak{B}$, A makes the judgements $[\top \succ p]$ and $[p \succ \bot]$, i.e. $\top \succ p \succ \bot$.

Regularity requires that A is always strictly more confident in the tautology than they are in any contingent proposition, and that they are always strictly less confident in the contradiction than they are in any contingent proposition. This is a generalisation of the controversial Regularity condition from Bayesian epistemology.[8] Now, the purpose of this paper is to analyse how a rational agent should revise their comparative confidence judgements upon learning, *with certainty*, the truth of some evidential proposition e. Of course, the very possibility of this kind of learning is ruled out a-priori by the regularity assumption, so I reject regularity as a universal coherence constraint.[9]

2.2 Synchronic Norms

Before attempting to explicate the diachronic rationality norms that govern the way in which an agent A should revise their comparative confidence judgements over time, it behooves us to review some of the most influential synchronic rationality constraints that have been forwarded in the literature.[10]

The first and most fundamental synchronic rationality constraints for comparative confidence orderings can be stated both qualitatively and in terms of possible numeric representations. I present the qualitative versions first.

[8]See e.g. Lewis (1980) and Skyrms (1980) for philosophical justifications of the regularity condition.

[9]Naturally, the familiar debate surrounding whether one can ever learn a proposition with certainty can be rerun in the comparative setting, but it seems to me that there probably isn't much to gain by doing so. Here, I am of course assuming that it is possible to learn a proposition with certainty, and hence that the kind of learning experience described by standard Bayesian conditionalisation can actually occur.

[10]Note that the literature is replete with possible synchronic coherence constraints for confidence orderings, and it would be impossible to provide an exhaustive survey here (the interested reader should consult e.g. Halpern (2003), Wong *et al* 1991). I review only those synchronic constraints that play a crucial role in what follows.

$$(A1)\ \top \succ \bot.$$

$$(A2)\ \text{For any } p, q \in \mathfrak{B}, \text{ if } p \vdash q \text{ then } q \succsim p.$$

$A1$ requires that rational agents always be strictly more confident in the tautology than they are in the contradiction, and $A2$ is a general monotonicity requirement, which stipulates that agents should never be strictly more confident in p than they are in the logical consequences of p. As well as being intuitively compelling, these rationality constraints have been given a range of pragmatic justifications (see e.g. Fishburn (1986), Halpern (2003)).

Given a comparative confidence ordering $\succsim$ over $\mathfrak{B}$ and a function $\mu : \mathfrak{B} \to [0, 1]$, say that $\succsim$ is 'fully represnted' by μ if and only if for every $p, q \in \mathfrak{B}$, (i) $p \succ q \Leftrightarrow \mu(p) > \mu(q)$, and (ii) $p \sim q \Leftrightarrow \mu(p) = \mu(q)$. Call a function $\mu : \mathfrak{B} \to [0, 1]$ a 'plausibility function' if it satisfies the following two conditions,

$$(PL1)\ \mu(\top) = 1 \text{ and } \mu(\bot) = 0.$$

$$(PL2)\ \text{For any } p, q \in \mathfrak{B}, \text{ if } p \vdash q \text{ then } \mu(p) \leq \mu(q).$$

It is well known that if $\succsim$ is a total preorder over $\mathfrak{B}$, then $\succsim$ will satisfy $A1$ and $A2$ if and only if $\succsim$ is fully representable by a plausibility function on $\mathfrak{B}$. Thus, our first prospective synchronic coherence requirement for comparative confidence judgements is

($\mathfrak{C}\mathfrak{1}$) $\succsim$ should be fully representable by a plausibility function, or equivalently, $\succsim$ should satisfy $A1$ and $A2$.

Another prospective qualitative rationality constraint on $\succsim$ is

$(A3)$ For any $p, q, r \in \mathfrak{B}$, if $p \vdash q$ and $< q, r >$ are logically incompatible,[11] then

$$q \succ p \Rightarrow q \vee r \succ p \vee r$$

A3 can be thought of as a weak additivity condition. It is best understood in terms of its implications for the representability of $\succsim$ by numerical functions on $\mathfrak{B}$. Call a function $\mu : \mathfrak{B} \to [0, 1]$ a 'mass function' if

$$(M1)\ \mu(\bot) = 0$$

$$(M2)\ \sum_{p \in \mathfrak{B}} \mu(p) = 1$$

[11] i.e. $p \wedge q \equiv \bot$.

Any mass function $\mu : \mathfrak{B} \to [0,1]$ defines a corresponding 'Dempster-Schafer belief function' $b_\mu : \mathfrak{B} \to [0,1]$ defined by $b_\mu(q) = \sum_{\{p \in \mathfrak{B} | p \vdash q\}} \mu(p)$. It is not hard to see that the set of all belief functions on $\mathfrak{B}$ is generally a proper subset of the set of all plausibility functions on $\mathfrak{B}$ and a proper superset of the set of all probability functions on $\mathfrak{B}$, where probability functions are defined as belief functions whose corresponding mass functions assign all of their mass to the possible worlds in $\mathfrak{B}$.[12] It turns out (see Wong *et al.* (1991)) that $\succsim$ satisfies all of $A1$, $A2$ and $A3$ if and only if $\succsim$ is fully representable by a Dempster Schafer belief function. Thus, the second prospective synchronic rationality requirement for comparative confidence judgements is

($\mathfrak{C}2$) $\succsim$ should be fully representable by a Dempster-Schafer belief function, or equivalently, $\succsim$ should satisfy $A1$, $A2$ and $A3$.

The following qualitative constraint is strictly stronger than $A3$, and will play a significant role in what follows.

($A4$) For any $p, q \in \mathfrak{B}$, $p \succ q \Leftrightarrow (p \wedge \neg q) \succ (\neg p \wedge q)$

Finally, the last and strictly strongest prospective synchronic rationality constraint that I will consider here is

($\mathfrak{C}3$) $\succsim$ should be fully representable by a probability function.[13]

It is easy to show that a confidence ordering which satisfies $\mathfrak{C}3$ automatically satisfies all the other synchronic constraints listed here.

We are now ready to ask whether and how these three synchronic rationality constraints can be justified. More specifically, we can ask 'what goes wrong when an agent's comparative confidence judgements fail to satisfy these representability requirements?'. Generally, there are two ways to justify prospective epistemic norms such as these. Firstly, one can provide a pragmatic justification, i.e. one can show that agents who violate the norm will be *instrumentally irrational* in the sense that they will act in ways which fail to produce the best outcomes by their own lights. An example of such a justification is the synchronic Dutch Book argument for probabilism. Alternatively, one could provide an epistemic justification, i.e. one could attempt to show that agents who violate the norm will end up with epistemic attitudes that are in some way defective. Of course, this typically involves appealing to

[12] And 'possible worlds' are just the atoms of the Boolean algebra $\mathfrak{B}$, i.e. the maximal consistent conjunctions of sentences of L.

[13] The qualitative statement of $\mathfrak{C}3$ (standardly referred to as the 'cancellation axiom') is rather technical, so I omit it here (but see i.e. Scott (1964), Konek (2019)).

some more fundamental epistemic norms, since the question of whether an agent's epistemic attitudes are defective is itself inherently normative. An example of such an epistemic justification is the accuracy dominance argument for probabilism (see e.g. Joyce (1998), Pettigrew (2016)), which itself appeals to the fundamental norm that an agent should aim to have 'accurate' numerical credences.

For current purposes, I will not take a substantive stand regarding the synchronic norms of comparative confidence judgements. My aim will rather be to systematically catalogue which synchronic norms one needs to assume in order to obtain normative justifications of the diachronic norm given by comparative conditionalisation. Towards this end, I will briefly pause now to recall the two most prominent extant normative justifications of synchronic norms from the literature.

Firstly, Icard (2016) has given a pragmatic 'money-pump' style argument in favour of $\mathfrak{C}3$. Since $\mathfrak{C}3$ is strictly stronger than all the other synchronic norms presented above, this argument, if successful, can be seen as providing a pragmatic justification of all the synchronic rationality constraints considered here. Secondly, Fitelson and McCarthy (unpublished) have developed a formal framework for assessing the accuracy of comparative confidence orderings.[14] This framework yields epistemic accuracy dominance arguments in support of $\mathfrak{C}1$ and $\mathfrak{C}2$. Importantly, these arguments do not extend to justifications of either $A4$ or $\mathfrak{C}3$, which suggests that $\mathfrak{C}2$ is probably the strongest interesting norm that can be justified by accuracy dominance arguments in Fitelson and McCarthy's framework.

3 Comparative Conditionalisation

We are now ready to address the central question of this article: how should a rational agent revise their comparative confidence judgements after learning the truth of some evidential proposition e? Before going further, it is worth spelling out a couple of important background assumptions.

Firstly, I assume here that the evidential proposition e is learned *with certainty.* Thus, the kind of learning I am interested in is the same as that described by standard Bayesian conditionalisation, where the agent assigns a posterior probability of 1 to the learned proposition. In the context of comparative confidence judgements, the analogous requirement is that after the learning experience, the agent makes the judgement $[e \sim \top]$, i.e. that they become equally confident in the truth of the learned proposition and the tautology.

Secondly, I assume that upon learning e, the agent needs to reorganise their comparative confidence judgements in a way that (i) ensures that they become certain in the truth of e, and (ii) defines

[14] This framework is introduced and explicated in section 5, where it is used to evaluate the possibility of an expected inaccuracy argument for comparative conditionalisation.

a confidence ordering that preserves all the relevant synchronic rationality requirements that were satisfied by their initial ordering. So, for example, if we assume $\mathfrak{C}1$, $\mathfrak{C}2$ and $\mathfrak{C}3$ as synchronic rationality requirements and the agent's initial ordering satisfies all these requirements, then their ordering should still satisfy those requirements after they have revised their comparative confidence judgements to accommodate the new evidence. Whatever the synchronic rationality norms are, learning new evidence should not lead one to violate them.

It is clear that there are generally many ways that an agent can revise their confidence orderings whilst satisfying these basic requirements (for any fixed specification of the synchronic norms). How to choose between them? It is instructive here to take inspiration from a key structural property of Bayesian conditionalisation. Specifically, given a probability distribution P, let $\succsim_P$ be the confidence ordering defined by $q \succ_P p$ if and only if $P(q) > P(p)$ and $p \sim_P q$ if and only if $P(p) = P(q)$. By definition, P fully represents $\succsim_P$, and we can think of $\succsim_P$ as encoding the comparative confidence judgements of a Bayesian agent whose credal state is given by the probability function P. Now, we can ask 'what is the relationship between $\succsim_P$ and $\succsim_{P(-|e)}$?', where $P(-|e)$ is the probability function obtained by conditionalising P on e. Less formally: 'how does conditionalising on e change the comparative confidence judgments implicit in P?'. Happily, this question has a simple answer:

$$\begin{aligned} q \succ_{P(-|e)} p &\Leftrightarrow P(q|e) > P(p|e) \\ &\Leftrightarrow P(q|e)P(e) > P(p|e)P(e) \\ &\Leftrightarrow P(e \wedge q) > P(e \wedge p) \\ &\Leftrightarrow e \wedge q \succ_P e \wedge p \end{aligned}$$

Thus, if we let $\succsim_e$ denote the ordering that results from revising the initial ordering $\succsim$ after learning e, a Bayesian agent will always revise their confidence orderings according to the rule

(CC:) $q \succ_e p \Leftrightarrow e \wedge q \succ e \wedge p$, and $q \sim_e p \Leftrightarrow e \wedge q \sim e \wedge p$

Where 'CC' stands for 'comparative conditionalisation'. The question now is whether there is anything special about CC as opposed to other revision rules for comparative confidence judgements. One might be tempted here to simply invoke the observation that there are numerous philosophical justifications for viewing Bayesian conditionalisation as the uniquely rational rule for updating numerical credences, and to conclude that the revision rule defined by Bayesian conditionalisation must therefore be the correct one. However, this kind of justification is clearly flawed. For, it assumes at the outset that an agent's comparative confidence judgements are defined by a specific credal state, and that the way in which an agent revises those judgements will be entirely determined by the rule they use to update that credal state. But, as I noted in the introduction, there is a significant minority of authors who contend that comparative confidence judgements are philosophically and psychologically more fundamental than assignments of numerical credence, and so will reject the implicit assumption that an agent's comparative confidence judgments are always determined by some specific credal state. It may be that the content of an agent's epistemic state is exhausted by their confidence ordering, and that they simply have no well defined credal state.[15] Again, it's at least coherent to conceive of such an agent. And in this context, rejecting CC in favour of another revision rule does not bring one into conflict with Bayesian conditionalisation. For, the way in which an agent revises their comparative confidence judgments will have *no* implications regarding the way in which they update their credences if they have no well defined credences in the first place.[16]

If one hopes to justify CC, then they need to do so within the context of the epistemology of comparative confidence judgements. My aim in the rest of this paper is to explore the possibility of systematically justifying CC within the context of a comparativist epistemology. But before doing so, it is worth emphasising the intuitive rationality of CC as opposed to alternative revision rules. Towards this end, consider the following example:

> Mufasa is sitting in a soundproof room with no windows, and he has no idea what the weather is like outside. The room is equipped with a speaker which will occasionally announce some partial information about the weather outside. Based on past experience, he judges that it is more likely to be raining and thundering outside than it is to be sunny and thundering outside, i.e. he makes the judgements $[r \wedge t \succ s \wedge t]$. The speaker then announces that it's thundering outside. Mufasa subsequently revises his comparative confidence judgements in a way that leads him to judge that it is more likely to be sunny outside than it is to be rainy outside, i.e. he makes the judgements $[s \succ_t r]$.

I take it that there is something intuitively bizarre about the dynamics of Mufasa's confidence judgements here. The question is whether this bizarreness is indicative of diachronic rationality. As a first step, I will now survey some basic formal properties of the CC updating rule.

[15]Note that this is true even if their confidence ordering is fully representable by a probability function, since there will generally be infinitely many different probability functions that can be used to represent the ordering.

[16]It is worth stressing again here that my aim in this paper is not to defend the position that we should conceive of an agent's epistemic state purely in terms of comparative confidence judgements. Rather, I start from the assumption that there are at least some situations in which such a conception is desirable, and then address the question of how agents in situations of this sort should revise their epistemic states over time in light of this assumption.

3.1 Formal Properties of CC

I require that any revision procedure for comparative confidence judgements should satisfy two key properties: (i) after learning e, the agent should make the judgment $[e \sim_e \top]$, and (ii) the posterior ordering $\succsim_e$ should preserve all relevant synchronic rationality constraints (for some fixed specification of the synchronic rationality constraints). I begin by noting that CC trivially satisfies (i), since

$e \sim_e \top \Leftrightarrow (e \wedge e) \sim (e \wedge \top) \Leftrightarrow e \sim e$, and all total preorders on $\mathfrak{B}$ satisfy $e \sim e$ for all $e \in \mathfrak{B}$.[17]

Next, we need to establish (ii). A first step is achieved with the following proposition (all proofs in appendix).

Proposition 1 *Let $\succsim$ satisfy $\mathfrak{C}1$ Then $\succsim_e$ satisfies $\mathfrak{C}1$ if and only if $e > \bot$.*

In what follows, I always assume that the agent is initially strictly more confident in the learned evidential proposition e than they are in the contradiction, i.e. $e \succ \bot$.[18] Given this assumption, we can show that CC preserves all the synchronic rationality constraints described in section 2.[19]

Proposition 2 *If $\succsim$ satisfies $\mathfrak{C}2$, then $\succsim_e$ satisfies $\mathfrak{C}2$.*

Proposition 3 *If $\succsim$ satisfies A4, then $\succsim_e$ satisfies A4.*

Proposition 4 *If $\succsim$ satisfies $\mathfrak{C}3$, then $\succsim_e$ satisfies $\mathfrak{C}3$.*

Thus, we know that regardless of which synchronic rationality constraints one imposes on the prior confidence ordering, revising by CC will never lead an agent to replace a coherent confidence ordering with an incoherent one. So CC satisfies the second condition on revision procedures for comparative confidence judgements, regardless of which of these synchronic coherence constraints one adopts as epistemic norms.[20]

---

[17] Similarly, $\neg e \sim_e \bot \Leftrightarrow (e \wedge \neg e) \sim (e \wedge \bot) \Leftrightarrow (e \wedge \neg e) \sim \bot$, which again is satisfied by all total preorders on $\mathfrak{B}$.

[18] This assumption is of course reminiscent of the fact that a Bayesian agent can never condition on a probability 0 event. Critics of Bayesian epistemology typically take this feature to be problematic and unmotivated. I don't address this issue here, but it is certainly worth noting that this aspect of Bayesian inference generalises so naturally to the comparative setting (and so can't be straightforwardly attributed to the ratio definition of conditional probabilities, as is often suggested).

[19] In fact, it is also straightforward to show that CC preserves another synchronic rationality constraint from the literature that is strictly stronger than $\mathfrak{C}2$ but strictly weaker than $\mathfrak{C}3$, namely that $\succsim$ should be a 'comparative probability relation' (see e.g. Kraft *et al.* (1959)).

[20] It is also worth noting that, perhaps unsurprisingly, CC shares many of the key structural properties of Bayesian conditionalisation. For example, CC defines a commutative revision procedure, i.e. the order in which the agent receives novel evidence makes no difference to the comparative confidence judgements that they end up with at the end of the learning process. To see this, let $\succsim_{e_1,e_2}$ be the result of revising $\succsim$ sequentially by e_1 and then e_2. Then

$$p \succ_{e_1,e_2} q \Leftrightarrow e_2 \wedge p \succ_{e_1} e_2 \wedge q \Leftrightarrow e_1 \wedge e_2 \wedge p \succ e_1 \wedge e_2 \wedge q \Leftrightarrow p \succ_{e_1 \wedge e_2} q.$$

The commutativity of CC is of course of fundamental importance, since it ensures that there is always a well defined and

3.2 Opinionation Failures and Conditional Judgement

I turn now to briefly describing two important points regarding the scope of CC's applicability. Firstly, until now, I've assumed that the prior ordering to which CC applies always satisfies Opinionation, a condition whose status as an epistemic norm is contested by many authors (see i.e. Forest (1989), Keynes (1921), Eva (forthcoming)). Happily, the definition of CC given here is straightforwardly generalised to the non-opinionated setting. To see this, let $p \odot q$ denote the case in which the agent makes no comparative confidence judgement regarding the pair (p, q). Then we can extend the original definition of CC to the non-opinionated case as follows (where $p \odot_e q$ denotes the case in which the agent makes no judgement regarding the pair (p, q) after learning e),

(CC*:) $q \succ_e p \Leftrightarrow e \wedge q \succ e \wedge p$, $q \sim_e p \Leftrightarrow e \wedge q \sim e \wedge p$, and $q \odot_e p \Leftrightarrow e \wedge q \odot e \wedge p$

The intuitive motivation for CC* is directly analogous to that of the standard CC rule in the opinionated setting. This analogy can be made formally precise in the context of the following synchronic coherence constraint for non-opinionated orderings.[21]

($A5$) there should exist a non-empty set $\mathcal{P}$ of probability measures on $\mathfrak{B}$ such that for any $p, q \in \mathfrak{B}$:

$$p \succsim q \Leftrightarrow (\forall P \in \mathcal{P})(P(p) \geq P(q))$$

$$p \odot q \Leftrightarrow (\exists P, P' \in \mathcal{P})(P(p) > P(q)) \wedge (P'(p) < P(q))$$

In this case, we say that $\mathcal{P}$ 'fully represents' $\succsim$

$A5$ is familiar from the theory of imprecise credences, where it is often assumed that an agent's credences are represented by a *set* of precise probabilistic credence functions, referred to as the agent's 'representor' (see i.e. Joyce (2010), Weatherson (2007)). On this view, the agent's comparative confidence ordering can be derived through the following supervaluationist semantics. Firstly, the agent makes the judgement $[p \succsim q]$ if and only every function in their representor assigns p a credence which is at least as high as what it assigns to q. Secondly, if there are two functions P_1, P_2 in the agent's representor such that $P_1(p) > P_1(q)$ and $P_2(q) > P_2(p)$, then the agent makes no comparative confidence judgement regarding p and q, i.e. their confidence ordering satisfies $p \odot q$. By definition, the ordering identified by this semantics always satisfies $A5$ (since it is fully represnted by the agent's representor), so just as $\mathfrak{C}_3$ is presupposed by precise Bayesianism, $A5$ is presupposed by influential models of imprecise credence. Typically, these imprecise models assume that upon learning a proposition e, a rational agent

intuitively rational way to iterate the revision procedure in sequential learning scenarios.

[21]For discussions of this constraint, and its qualitative axiomatisation, see e.g. Alon and Lehrer (2014), Forrest (1989), Harrison-Tailor *et al* (2016).

will replace their prior representor $\mathcal{P}$ by the set $\mathcal{P}(-|e) = \{P(-|e)|P \in \mathcal{P}\}$, i.e. that they will simply condition every function in their prior representor on e and take the set of updated functions as their new representor. Now, it's easy to see that if $\succsim$ is fully represented by the agent's representor, then the posterior ordering obtained by applying CC* will always be fully represented by the agent's posterior representor. So just as CC coheres perfectly with standard conditionalisation, CC* coheres perfectly with its imprecise counterpart. Thus, CC can be straightforwardly and naturally generalised to the non-opinionated setting. For now, I am content to stay within the opinionated setting and investigate the normative status of CC, but the task of exploring prospective normative justifications of CC* is a worthy and pressing one that I intend to return to in future work.

The second important point to note regarding the scope of CC's applicability concerns the rule's relation to supposition and conditional judgement. Here, it is significant that the definition of standard Bayesian conditionalisation relies on the availability of *conditional degrees* of belief. In order to calculate my new credence in q after conditionalising on p, I need to know my prior conditional degree of credence in q given p, $P(q|p)$, which is standardly interpreted as representing my credence in q *under the (indicative) supposition* that p is true. In the comparative context, Koopman (1940) forwarded a set of axioms whose satisfaction allowed for the definition of an analogous notion of *comparative conditional confidence*. It is significant that the definition of CC does not involve reference to any such notion. The rule can be straightforwardly and intuitively applied without any appeal to representations of conditional or suppositional judgement. This suggests that the close relationship between learning, supposition and conditional judgement that is familiar from Bayesian epistemology is likely to be fundamentally different in the comparative setting.

Enough with the preliminaries. As I noted above, CC looks like an intuitively compelling approach to revising comparative confidence judgements on the basis of novel evidence. Indeed, there is something intuitively inconsistent about agents who revise their judgements in a way that conflicts with CC. But I have not yet given any definite argument in support of the claim that agents who revise their confidence orderings by a rule other than CC are irrational. In what follows, I explore the most salient avenues for the pursuit of such an argument.

4 A Comparative Diachronic Dutch Book Argument

One of the most influential normative justifications for standard Bayesian conditionalisation is the diachronic Dutch book argument, first reported in Teller (1973). The diachronic Dutch book argument purports to provide a purely pragmatic justification of conditionalisation, in the sense that it (allegedly) demonstrates that all violations of conditionalisation are liable to yield behaviours which are instrumentally irrational. Specifically, diachronic Dutch book arguments aim to show that agents who update their credences via a strategy other than Bayesian conditionalisation can be 'Dutch booked' by bookies who exploit the agent's credences to trade them a series of bets which generate a sure loss. In this section, I'll present a natural generalisation of the diachronic Dutch book argument to the comparative setting (§4.1), catalogue which assumptions are needed to get the argument going, and evaluate the extent to which this generalised argument constitutes a genuine normative vindication of CC (§4.2). Note that in the remainder of the paper, $\succsim_e$ will always denote the posterior ordering produced by CC, and alternative posterior ordering will always be denoted by $\succsim^*$.

4.1 Comparative Betting Principles

The first step in constructing the generalised argument is to identify principles linking an agent's comparative confidence judgements to their betting behaviour. In most presentations of the standard diachronic Dutch book argument for Bayesian conditionalisation, it is simply assumed that agents accept all and only bets wth positive expected utility (see i.e. Pettigrew (forthcoming), Lewis (1999)). In the absence of numerical credences, we need more general principles that do not rely on probabilistic expectations for their articulation. In the current setting, where Opinionation is assumed, we will need one betting principle for equal confidence judgements of the form $[p \sim q]$, and one for judgements of unequal confidence, of the form $[p \succ q]$. Before presenting the specific principles to be used in the Dutch book argument for CC, consider first the following intuitively compelling principle (which concerns judgements of unequal confidence).

Comparative Betting Principle 1 (CBP1): Let A be an agent who makes the comparative confidence judgement $[p \succ q]$. Then for any price α, it is rational for A to pay α to buy a bet with the following payoff matrix.

Outcome	Payoff
$p \wedge \neg q$	2α
$\neg p \wedge q$	0
$p \equiv q$	α

To get a feeling for the content of CBP1, suppose first that A's comparative confidence judgements are fully represented by some probability function P, and that A makes the judgement $[p \succ q]$. Then it's easy to see that the bet outlined above will always have positive expected utility according to P. So if we assume that agents should always accept bets which have positive expected utility according to all probabilistic representations of their comparative confidence judgements, then CBP1 is entailed as a practical norm whenever $\succsim$ is probabilistically representable (whenever $\mathfrak{C}3$ is assumed as a synchronic norm). In fact, CBP1 can be more generally justified on the basis of an even more intuitive betting principle, namely that whenever an agent makes the judgement $[p \succ \neg p]$, they should always be willing to buy an even odds bet on the truth of p. The justification runs as follows (where it's assumed that A makes the judgement $[p \succ q]$).

1: In the event that A makes the judgement $[p \succ \neg p]$, they should always be willing to buy an even odds bet on the truth of p.

2: In the event that A makes the judgement $[p \succ q]$, they should also make the judgement $[p \wedge \neg q \succ \neg p \wedge q]$.

3: The bet in CBP1 is effectively called off in the event that $p \equiv q$ is true. So when assessing the bet, the agent can safely assume that $p \equiv q$ is false.

4: In the event that $p \equiv q$ is false, p is equivalent to $p \wedge \neg q$ and $\neg p$ is equivalent to $\neg p \wedge q$.

5: If A makes the judgement $[p \wedge \neg q \succ \neg p \wedge q]$, they should continue to do so under the assumption that $p \equiv q$ is false.

6: So A should make the judgement $[p \succ \neg p]$ under the assumption that $p \equiv q$ is false. (From 2, 4, 5)

7: Assuming that $p \equiv q$ is false, the bet in CBP1 is effectively an even odds bet on the truth of p.

8: Since A makes the judgement $[p \succ \neg p]$ under the assumption that $p \equiv q$ is false, they should be willing to accept an even odds bet on the truth of p under that assumption. (From 1, 3, 6)

9: So A should accept the bet in CBP1. (From 7, 8)

I take it that steps 3, 4 and 7 are uncontroversial. Step 1 is the even odds betting principle from which we are trying to derive a justification of CBP1. Steps 6, 8 and 9 all follow logically from some combination of the preceding steps. So the only steps that require explicit consideration are 2 and 5. I address each in turn.

Firstly, note that step 2 follows directly from the synchronic norm $A4$. So in the event that we assume $A4$ as a synchronic rationality constraint, step 2 requires no further justification. Importantly, $A4$ is significantly weaker than the full probabilistic representability constraint $\mathfrak{C}3$, so we do not need full probabilistic representability to get this justification of CBP1 off the ground (although some authors think we always have good practical reason to assume probabilistic representability anyway, i.e. Icard (2016)). However, it is worth recalling that $A4$ is still strictly stronger than $\mathfrak{C}2$ and is not justifiable by an accuracy dominance argument in Fitelson and McCarthy's epistemic utility framework for comparative confidence judgements. So the reliance on $A4$ at this stage is by no means uncontroversial.

Next, note that step 5 is also a substantive assumption. It can be interpreted as a constraint on the kind of learning rule used by A – namely that if one initially makes the judgement $[p \succ q]$, they should continue to make that judgement after learning (or assuming) any proposition r which is a logical consequence of both p and q. This constraint is trivially satisfied by CC, and also by infinitely many other updating rules. Like $A4$, it is intuitively compelling. If I'm more confident that there are three people in the room than I am that there are two in the room, and then I learn (only) that there are less than five people in the room, that should never lead me to change my initial judgement since it does nothing to discriminate between those two possibilities. It doesn't rule out any worlds in which either of them are true, and it also doesn't tell us anything about the relative plausibilities of any of those worlds.

So, if one is happy to accept steps 2 and 5 above, along with the even odds betting principle expressed by 1, then one has good reason to accept CBP1. But CBP1 on its own isn't quite enough to get the Dutch book argument going. We require the following, slightly more controversial, principle.

Comparative Betting Principle 2 (CBP2): Let A be a rational agent who makes the comparative confidence judgment $p \succ q$, and let $\alpha > 0$ be any price. Then there will exist β_1, β_2 such that (i) $\alpha > \beta_1, \beta_2$, (ii) $\beta_1, \beta_2 > 0$, and (iii) A will pay α to buy a bet with the following structure.

outcome	payoff
$p \wedge \neg q$	$2\alpha - \beta_1$
$\neg p \wedge q$	0
$p \equiv q$	$\alpha - \beta_2$

CBP2 is strictly stronger than CBP1 in so far as it positively identifies a broader range of bets as rationally acceptable. Specifically, if requires that when A makes the judgement $[p \succ q]$, there should be *some* bet which is acceptable to the agent, and which has worse expected utility than the CBP1 bet in the sense that the payoffs in the $p \wedge \neg q$ and $p \equiv q$ cells are smaller than the corresponding payoffs in the CBP1 bet (by a degree of β_1 and β_2, respectively). The difference between this bet and the bet in CBP1 can be arbitrarily small. The idea is that if one is willing to accept the CBP1 bet, they should also be willing to accept *some* bet with an infinitesimally less appealing payoff structure. Since the bet in CBP1 is equivalent to the bet in CBP2 in the limit where β_1, β_2 approach zero, and β_1, β_2 can be arbitrarily small, there will always be some values of β_1, β_2 for which the bet in CBP2 becomes acceptable - or so the story goes. Again, it is easy to see that in the special case where $\succsim$ is representable by some probability function P, there will always exist $\beta_1, \beta_2 > 0$ for which the bet in CBP2 has positive expected utility according to P. So in the special case of probabilistic representability, CBP2 coheres with the requirement that agents accept all and only bets with positive expected utility.

Before presenting the argument, I will also need a principle that connects judgements of equal confidence to rational betting behaviour, namely

Comparative Betting Principle 3 (CBP3): Let A be a rational agent who makes the comparative confidence judgment $p \sim q$, and let $\alpha > 0$ be any price. Then for any β_1, β_2 satisfying $\beta_1 + \beta_2 > 2\alpha$, A will buy bets of the following form.

outcome	**payoff**
$p \wedge \neg q$	β_1
$\neg p \wedge q$	β_2
$p \equiv q$	α

As with CBP1/2, it is worth observing that whenever $\succsim$ is probabilistically representable, all representations of $\succsim$ will expect the bet in CBP3 to have positive utility whenever $p \sim q$ obtains. More generally, CBP3 can be given a justification that is directly analogous to the justification given for CBP1. To see this, consider the simple betting principle which states that whenever an agent A is equally confident in p and $\neg p$, they should be willing to buy a bet depending only on the truth value of p whenever the combined payoffs for the possibilities p and $\neg p$ are greater than twice the price α.[22] This principle can be used to justify CBP3 in the following manner (assuming that A makes the judgement $[p \sim q]$).

1: In the event that A makes the judgement $[p \sim \neg p]$, they should always be willing to buy a bet depending only on the truth of p where the combined payoffs for the possibilities p and $\neg p$ are greater than twice the price α.

2: In the event that A makes the judgement $[p \sim q]$, they should also make the judgement $[p \wedge \neg q \sim \neg p \wedge q]$.

3: The bet in CBP3 is effectively called off in the event that $p \equiv q$ is true. So when assessing the bet, the agent can safely assume that $p \equiv q$ is false.

4: In the event that $p \equiv q$ is false, p is equivalent to $p \wedge \neg q$ and $\neg p$ is equivalent to $\neg p \wedge q$.

5: If A makes the judgement $[p \wedge \neg q \sim \neg p \wedge q]$, they should continue to do so under the assumption that $p \equiv q$ is false.

6: So A should make the judgement $[p \sim \neg p]$ under the assumption that $p \equiv q$ is false. (From 2, 4, 5)

7: Assuming that $p \equiv q$ is false, the bet in CBP3 is effectively a bet depending only on the truth of p where the combined payoffs for the possibilities p and $\neg p$ are greater than twice the price α.

[22]Where a bet depending only on the truth value of p is a bet whose payoff one can perfectly predict as soon as they have knowledge of the truth value of p.

8: Since A makes the judgement $[p \sim \neg p]$ under the assumption that $p \equiv q$ is false, they should (under that assumption) buy any bet depending only on the truth of p where the combined payoffs for the possibilities p and $\neg p$ are greater than twice the price α. (From 1, 3, 6)

9: So A should accept the bet in CBP3. (From 7, 8)

Again, this justification relies crucially on the synchronic norm $A4$ and the constraint that learning a proposition r which is a logical consequence of both p and q should not lead on to revise their comparative confidence judgements regarding the pair (p, q). The only substantive difference with the justification for CBP2 is in the starting point from which the principle is eventually derived.

In the following, I will assume CBP2 and CBP3 as constraints governing the betting behaviour of rational agents (recall that CBP2 is strictly stronger than CBP1, and hence that CBP1 is also implicitly assumed).

4.2 The Argument

I am now ready to present the main assumptions which, combined with CBP2 and CBP3, yield a diachronic Dutch book argument for CC. First, note that, as in the standard Dutch book arguments for Bayesian conditionalisation, I assume that (i) the agent A has a fixed updating strategy, i.e. a plan for how they will revise their comparative confidence judgements given any possible piece of evidence e, and (ii) that this updating strategy is known to both A and the bookie.[23] Secondly, I assume that the updating strategy is such that for some proposition e, learning e would lead A to revise their comparative confidence judgements in a way that conflicts with CC. This implies that there exist propositions $x, y \in \mathfrak{B}$ such that one of the following four cases obtains, where $\succsim$ denotes the agent's prior ordering and $\succsim^*$ denotes their posterior ordering after learning e

(i) $(e \wedge p) \succ (e \wedge q)$, $q \succ^* p$.

(ii) $(e \wedge p) \succ (e \wedge q)$, $q \sim^* p$.

(iii) $(e \sim p) \succ (e \wedge q)$, $q \succ^* p$.

(iv) $(e \sim p) \succ (e \wedge q)$, $q \succ^* p$.

These cases enumerate the four possible types of CC violation. The argument aims to demonstrate that violations of type (i) and (ii) are liable to yield the acceptance of series of bets which generate a sure loss.

[23]Both of these assumptions have been the targets for recent critiques of Dutch book arguments for diachronic rationality norms (see i.e. Pettigrew (forthcoming)).

Thirdly, I also assume (importantly) that the bookie knows nothing that the agent A doesn't know, and hence that the truth values of the propositions e, x and y are unknown to the bookie. Finally, I assume (as in standard in Dutch book arguments) that there exists some future time t_e at which the truth value of e will be simultaneously revealed to both the agent and the bookie.[24]

The argument runs as follows, where we are concerned with showing that CC violations of types (i) and (ii) can lead to the acceptance of sequences of bets that generate a sure loss.

1: Suppose that A makes the judgement $[(e \wedge x) \succ (e \wedge y)]$ at time t_0. Let $\alpha > 0$ be any price. By CBP2, there will exist numbers $\beta_1, \beta_2 > 0$ for which A is willing to pay α for the following bet.

outcome	**payoff**
$e \wedge x \wedge \neg y$	$2\alpha - \beta_1$
$e \wedge \neg x \wedge y$	0
$(e \wedge x) \equiv (e \wedge y)$	$\alpha - \beta_2$

So at time t_0, the bookie sells this bet to A for α.

2: The bookie now waits for time t_e, at which point the truth value of e is revealed to A and the bookie. If e is false, then A makes a loss of $\beta_2 > 0$ on the first bet, and the bookie can quit while she's ahead. If e is true, then A will follow their updating strategy to revise their prior confidence judgements. By assumption, they will now make one of the two judgements $y \succ^* x$ or $x \sim^* y$. Suppose first that they make the judgement $y \succ^* x$. By CBP2 there will exist β_3, β_4 for which they are now willing to buy the following bet for price α.

[24]Where it is also assumed that A does not learn anything else in the intervening time before t_e.

outcome	**payoff**
$\neg x \wedge y$	$2\alpha - \beta_3$
$x \wedge \neg y$	0
$x \equiv y$	$\alpha - \beta_4$

So at time t_e, the bookie sells this bet to A for α. There are now three possibilities. If $x \wedge \neg y$ is true, then the agent will have won the first bet and lost the second, yielding an overall loss of $\beta 1$. If $\neg x \wedge y$ is true, the agent will have lost the first bet and won the second, yielding an overall loss of β_3. If $x \equiv y$ is true, the agent will accrue an overall loss of $\beta_2 + \beta_4$. So in all eventualities, A will lose money.

It remains to address the case in which A makes the judgement $x \sim^* y$ at t_e. In this case, CBP3 entails that A will be willing to buy the following bet for the same price α they paid for their first bet, where β_3, β_4 are such that $\beta_3 + \beta_4 > 2\alpha$, $\beta_4 < 2\alpha$ and $\beta_1 > \beta_3$.[25]

outcome	**payoff**
$x \wedge \neg y$	β_3
$\neg x \wedge y$	β_4
$x \equiv y$	α

So at time t_e, the bookie will sell this bet to the agent for the price α.There are now three possibilities. If $x \wedge \neg y$ is true, then the agent will have won the first bet and 'lost' the second, yielding an overall loss of $\beta_1 - \beta_3 > 0$. If $\neg x \wedge y$ is true, the agent will have lost the first bet and 'won' the second, yielding an overall loss of $2\alpha - \beta_4 > 0$. If $x \equiv y$ is true, the agent will accrue an overall loss of β_2. So in all eventualities, A will lose money.

In sum then, we have considered two of the four ways that the agent can diverge from CC at time

[25]It's easy to check that there will always exist suitable β_3, β_4. For example, set $\beta_4 = 2\alpha - \frac{\beta_1}{2}$, $\beta_3 = \frac{3\beta_1}{4}$.

t_e. Either they make the judgement $[x \sim^* y]$ or they make the judgement $[y \succ^* x]$. Either deviation will leave A susceptible to a sequence of bets which will yield a sure loss.

4.3 Replies

To what extent does the preceding argument constitute a genuine pragmatic vindication of CC?[26] The first thing to note is that the argument only counts against two specific types of CC violation, namely the types that occur when an agent makes the prior judgement $[e \wedge x \succ e \wedge y]$ but fails to make the posterior judgement $x \succ^* y$. It does not apply to the kinds of violations that occur when the agent makes the prior judgement $[e \wedge x \sim e \wedge y]$ but fails to make the posterior judgement $x \sim^* y$. In fact, it is easy to see that the betting principles forwarded here are not sufficiently strong to yield a Dutch book for this kind of violation. So at best, the preceding argument only rules out a subset of possible CC violations as instrumentally irrational.

To make this precise, say that a confidence ordering $\succsim_2$ is an 'extension' of an ordering $\succsim_1$ if and only if $x \succ_1 y$ implies $x \succ_2 y$. Note that $\succsim_2$'s being an extension of $\succsim_1$ leaves open the possibility that $\succsim_1$ and $\succsim_2$ can differ in the sense that there can exist x, y such that $x \sim_1 y$ and $x \succ_2 y$. The comparative diachronic Dutch book argument purports to show that an agent who learns e and adopts a posterior ordering that is not an extension of the posterior ordering $\succsim_e$ produced by CC will be disposed to accept a series of bets that generate a sure loss. Thus, the argument is not a fully fledged justification of CC as a practical norm, but rather a justification of the rule that one should always update in a way which produces confidence orderings that are extensions of those produced by CC. Informally, the argument shows that while it may be permissible to update in a way that deviates slightly from CC, any 'large violations' of CC lead one away from the path of instrumental rationality.

Now, one could well search for alternative betting principles in order to reinforce the argument and achieve a full justification of CC.[27] But the argument presented above suffices for current purporses. Although it doesn't fully justify the normative status of CC, it strongly constrains the range of CC violations that are permissible from the perspective of instrumental rationality, and does so by assuming a very weak decision theory for comparative confidence judgements. The conclusion of the argument is also sufficiently strong to draw out the conflict between the pragmatic norm of sure loss avoidance

[26] I stress that my focus here is on those aspects of the argument that are peculiar to the current comparative formulation. For present purposes, I am happy to ignore those aspects of the argument which are shared with standard probabilistic formulations (i.e. that the bookie knows the agent's updating strategy, that the strategy is deterministic etc). In that sense, I am interested only in evaluating whether the diachronic Dutch book for CC is *as good* as the corresponding arguments for Bayesian conditionalisation, rather than on whether any of those arguments work at all.

[27] In ongoing work, I conduct a more thorough survey of the betting principles that can be used to ground Dutch book arguments for CC and related rules.

and the epistemic norm of expected inaccuracy minimisation for comparative updating, on which we will focus in subsequent sections.

Secondly, it should be stressed that much of the heavy lifting in the argument is done by the comparative betting principles CBP2 and CBP3. Once these principles are accepted, the argument follows quite straightforwardly from the same assumptions that are at play in the standard diachronic Dutch book argument for Bayesian conditionalisation.

Thus, the argument is only as compelling CBP2 and CBP3. And at first blush, these are both plausible principles. Both cohere naturally with the idea that agents with probabilisitic credences should accept all and only bets with positive expected utility. But of course, this observation alone is not a dialectically acceptable justification, since we are concerned with investigating the prospects of normatively justifying CC in a purely comparative setting, where it is not assumed that agents come equipped with probabilistic credences – hence the alternative justifications for CBP2/CBP3 presented in section 4.1.

The first thing to note about these justifications is that they rely crucially on the synchronic norm $A4$, which is strictly much weaker than full probabilistic representability, but still stronger than what can be justified by Fitelson and McCarthy's accuracy dominance arguments, for example. So the present Dutch book argument for CC is contingent on a controversial but plausible and widely accepted synchronic rationality constraint.

Secondly, the justifications for CBP1/CBP3 themselves start from weaker betting principles that one may also take issue with. For instance, the justification for CBP1 starts from the premise that when one makes the judgement $[p \succ \neg p]$, it is always rational to buy an even odds bet on the truth of p. While I take this to be an eminently plausible condition, I readily acknowledge that it is also substantive and could well be resisted by those that are skeptical about any tight connection between epistemic attitudes and betting behaviour. However, if one hopes to take a broadly functionalist approach to comparative confidence judgments, on which judgements are intimately connected to i.e. betting dispositions, then it's hard to think of a more plausible principle than this one (and the same goes for the principle from which CBP3 is derived).

Thirdly, the justifications for CBP1 and CBP3 both relied on the constraint that learning a proposition r which is a logical consequence of both p and q should never affect the comparative confidence judgments one makes regarding the pair (p, q). In this case, it's hard to conceive of a situation in which the norm could lose its force (if one is happy with the idea that there exist any diachronic norms governing the relationship between one's prior and posterior epistemic judgements), and I am happy to take this assumption at face value. One might still worry that it is circular to make an assumption about the nature of rational learning to justify betting principles that are subsequently used in a pragmatic justification of a particular learning rule, but this worry is easily dispelled. For, the constraint that learning a proposition r which is a logical consequence of both p and q should never affect the comparative confidence judgments one makes regarding the pair (p, q) is satisfied by infinitely many updating rules, of which CC is only one. While assuming this constraint significantly narrows down the range of candidate updating rules, it still allows for the possibility of radical CC violations, and is therefore a legitimate assumption in the current context.

Overall then, the justifications for CBP1/2/3 presented above rest on some substantive, but broadly plausible assumptions regarding the synchronic norms of comparative confidence, the relation between rational betting behaviour and comparative confidence, and the nature of rational learning. Dialectically, one could either use these assumptions to derive CBP1/2/3 , or simply take CBP1/2/3 as axiomatic. Both options looks broadly plausible to me, but I consider CBP1/2/3 to be just as intuitively compelling as the conditions from which they can be derived. Again, I should stress that my aim here is to survey the extent to which CC can be normatively justified in the same way that its Bayesian counterpart typically is, and to catalogue the assumptions that are needed to get these justifications off the ground, rather than to pin my colours to any particular justifying argument.

5 Comparative Expected Inaccuracy

In the previous section, I presented a purely pragmatic argument for the normative status of CC. Although the argument didn't achieve a full justification of CC, it did very strongly constrain the set of rationally permissible violations of CC, and in doing so identified CC as normatively significant. I turn now to exploring the possibility of constructing a purely *epistemic* justification of CC. Again, the strategy will be to identify an extant justification for Bayesian conditionalisation, and investigate whether it can be transferred to the comparative setting to ground a normative justification for CC. The most influential epistemic justification for Bayesian conditionalisation in the current literature is the argument from expected inaccuracy, first articulated by Graves and Wallace (2006). The basic idea behind the argument is the following.

Let P be the probability distribution that encodes the prior credences of a Bayesian agent A. Upon learning the evidential proposition e, A needs to adopt a new probability function P^* such that $P^*(e) = 1$. But how should they choose amongst the infinitely many functions which satisfy this condition? Greaves and Wallace propose that A should choose the function P^* that satisfies $P^*(e) = 1$ and minimises the quantity

$$Exp(P^*|P) = \sum_{w \in W} P(w)\mathfrak{I}(P^*, w)$$

where W is the set of possible worlds and $\mathfrak{I}$ is the strictly proper scoring that is assumed to encode A's conception of the inaccuracy of a credal state. Intuitively, this quantity is supposed to encode the *expected inaccuracy* of the credal state represented by the probability function P^* according to A's prior credal state P. By minimising this quantity, A will find the function P^* which, by the lights of A's own initial credences, is expected to be the most accurate function that satisfies the new evidential constraint. Greaves and Wallace show that the expected inaccuracy of P^* is uniquely minimised (as a function of P^*) when $P^* = P(-|e)$, i.e. updating by conditionalisation will always produce the posterior probability function with the lowest possible expected inaccuracy. Assuming that rational agents should always attempt to minimise the inaccuracy of their posterior credences, we then reach the conclusion that conditionalisation is the unique rational updating rule for agents with probabilistic credences. To date, this is the most influential epistemic justification of Bayesian conditionalisation as a principle of ideal rationality.

The idea now is to study whether a similar argument can be made for CC, i.e. to check whether CC always leads to those comparative confidence judgements that an agent initially expected to be the most accurate in the comparative context. In order to do this, we need access to a suitable scoring rule for quantifying the inaccuracy of an agent's comparative confidence judgements. Happily, the problem of identifying such a scoring rule has already been addressed by Fitelson and McCarthy (unpublished), whose framework I will now briefly introduce.

5.1 Accuracy and Confidence Orderings

Just as there is an intuitive sense in which an agent's numerical credences can be more or less accurate, there is likewise an intuitive sense in which an agent's comparative confidence judgements can be more less accurate. One might hope that, by making this intuitive notion formally precise, it will be possible to obtain epistemic justifications for norms governing the comparative confidence judgements of rational agents. This is precisely the project taken up by Fitelson and McCarthy (unpublished), who base their formalisation of accuracy for confidence orderings on the following premise,

> [A] confidence ordering is (qualitatively) inaccurate (at w) if and only if it fails to rank all the

truths strictly above all the falsehoods (at w). (Fitelson and McCarthy (unpublished): 6)[28]

The idea is that an agent A makes an epistemic mistake (at a world w) when they fail to be strictly more confident in a proposition that is true at w than they are in a proposition that is false at w. There are two ways in which this can happen:

(Case 1) A makes the judgement $[q \succ p]$, but $w \models p \wedge \neg q$.

(Case 2) A makes the judgement $[q \sim p]$, but $w \models \neg(p \equiv q)$.

Intuitively, A's epistemic mistake in Case 1 is worse than their epistemic mistake in Case 2. For, in Case 2, A merely fails to be more confident in a truth than they are in a falsehood, while in Case 1, A is actually more confident in the falsehood than they are in the truth. These observations suggest that the overall inaccuracy of A's comparative confidence judgements at a world w should be a weighted sum

$$\mathfrak{I}_{\geq}(\succsim, w) = \alpha_1 \cdot M_1(\succsim, w) + \alpha_2 \cdot M_2(\succsim, w)$$

where $\succsim$ is A's confidence ordering, $M_1(\succsim, w)$ and $M_2(\succsim, w)$ are the numbers of Case 1 mistakes and Case 2 mistakes that A makes at w, respectively, and $\alpha_1 > \alpha_2$.[29] At this point, we need to specify the exact strength of this inequality, i.e. we need to specify how much worse a Case 1 mistake is than a case 2 mistake. Fitelson and McCarthy show that there is actually a good argument for requiring that $\alpha_1 = 2\alpha_2$, i.e. that Case 1 mistakes are exactly twice as bad as Case 2 mistakes. Specifically, they show that setting $\alpha_1 = 2\alpha_2$ is the only way to ensure that the scoring rule $\mathfrak{I}_{\geq}$ is *evidentially proper*, where an evidentially proper scoring rule is one that guarantees that whenever $\succsim$ is fully representable by a probability function P, $\succsim$ is the unique confidence ordering that minimises expected inaccuracy according to P.[30] If our scoring rule were not evidentially proper, then there would be orderings that were fully represented by probability functions, but which didn't minimise expected inaccuracy according to those functions.[31] Such orderings would appear to be 'self-undermining' in the sense that their probabilistic representations would expect other orderings to be more accurate.[32] For example,

[28] A brief clarificatory note is in order here. At the time of the writing of this article, Fitelson and McCarthy's manuscript on scoring rules for comparative confidence judgements remains unpublished. However, all of the material referenced herein is publicly available in draft form at the url https://davidmccarthy.org/wp-content/uploads/2018/01/fitelson-mccarthy2015.pdf.

[29] i.e. $M1 = |\{(p, q) \in \mathfrak{B} \times \mathfrak{B} | (q \succ p) \wedge (w \models p \wedge \neg q)\}|$, $M2 = |\{(p, q) \in \mathfrak{B} \times \mathfrak{B} | (q \sim p) \wedge (w \models p \wedge \neg q)\}|$.

[30] Note that their proof of this claim does not assume Regularity, and so is relevant in the current generalised setting.

[31] Where the expected inaccuracy of an ordering $\succsim$ according to a probability function P is defined in the obvious way, i.e. $Exp_{\geq}(\succsim |P) = \sum_{w \in W} P(w) \cdot \mathfrak{I}_{\geq}(\succsim, w)$.

[32] This is of course similar to the sense in which immodest credences are said to be 'self-undermining' (see e.g. Joyce (1998)).

we can imagine a case in which an agent is informed of the relevant chance distribution determining the objective probabilities of the propositions under consideration. If they apply the Principal Principle, they will then adopt credences that match this distribution. These credences will define a corresponding confidence ordering, and it would be peculiar if the agent expected (relevant to their actual credences) another ordering to be more accurate than their own. It should be noted that this justification does not assume that agents are always equipped with probabilistic credences, but rather than *whenever an agent has such credences*, they should expect their own confidence ordering to be maximally accurate.

So the requirement that one's scoring rule be evidentially proper is, at first blush, a natural and persuasive one, and it motivates the following explicit definition,

$$\mathfrak{I}_{\geq}(\succsim, w) = 2 \cdot M1(\succsim, w) + M2(\succsim, w)$$

This is the final form of the scoring rule proposed by Fitelson and Mccarthy. It is important to stress that this scoring rule follows directly from two compelling axioms: evidential propriety and the requirement that inaccuracy should be a linear sum of the inaccuracies of the individual judgements encoded by an agent's ordering. As I've already stressed, violations of evidential propriety would seem to commit agents to a problematic form of epistemic modesty, on which they expect their own confidence orderings to be less accurate than some other possible orderings. The linearity assumption is also familiar from standard epistemic utility theory for probabilistic credences, where it is standardly assumed that global inaccuracy should be a weighted sum of the 'local inaccuracies' generated by individual credal judgements (see i.e. Pettgirew (2016), pp39–40 for a discussion and motivation of this assumption). It is of course possible to reject the linearity assumption, and argue that the global inaccuracy of a confidence ordering can not be calculated by simply adding up all of its individual local inaccuracies. However, it is hard to conceive of a non-additive rule that is anything like as natural as the additive rule, and there is (as of yet) no obvious theoretical ground for making such a choice. At any rate, I am happy to temporarily accept Fitelson and McCarthy's linearity assumption for current purposes (but see section 6 for more discussion).

Before moving on, it is also worth preempting one possible criticism of Fitelson and McCarthy's scoring rule, namely that it ignores a third kind of epistemic mistake that should also be taken into account when assessing the inaccuracy of confidence orderings. Specifically, one might think that it is epistemically non-optimal for an agent to make the judgement $[p \succ q]$ when $p \equiv q$ is actually the case. Certainly, an omniscient being would never make such a judgement in this case. And the current scoring rule doesn't take these mistakes into account. Again though, it turns out that evidential propriety is the reason for ignoring this third class of epistemic mistake. If we add a third term to the scoring rule corresponding to the number $M3$ of these mistakes, and weight that number by some factor α_3, then the resulting rule will be evidentially proper if and only if $\alpha_3 = 0$. Evidential propriety implies that these 'mistakes' aren't really mistakes at all. And that's why they are omitted from the scoring rule.

With their scoring rule in hand, Fitelson and McCarthy employ the following accuracy dominance avoidance norm (familiar from the epistemic utility theoretic justification of probabilism) to provide epistemic justifications for some of the synchronic rationality requirements described in section 2.

Weak Accuracy Dominance Avoidance (WADA): $\succsim$ should not be weakly dominated in accuracy, i.e. there should not exist any $\succsim^*$ such that

$$\text{(i) } (\forall w)[\mathfrak{I}_{\geq}(\succsim^*, w) \leq \mathfrak{I}_{\geq}(\succsim, w)]$$

$$\text{(ii) } (\exists w)[\mathfrak{I}_{\geq}(\succsim^*, w) < \mathfrak{I}_{\geq}(\succsim, w)]$$

Intuitively, an ordering $\succsim$ that is weakly dominated by another ordering $\succsim^*$ is guaranteed to be alethically sub-optimal in the sense that it can never be more accurate than $\succsim^*$, but it can be less accurate. Assuming that rational agents aim to have accurate epistemic states, it is clear that they should strive to ensure that their comparative confidence judgements can not be weakly accuracy dominated. Fitelson and McCarthy prove the following results.[33]

Theorem 5 *(1) If $\succsim$ violates $\mathfrak{C}1$, then $\succsim$ is weakly accuracy dominable. (2) If $\succsim$ violates $\mathfrak{C}2$, then $\succsim$ is weakly accuracy dominable. (3) $\succsim$ can violate $A4$ and $\mathfrak{C}3$ without being weakly accuracy dominable.*

Thus, the norm **WADA** alone is sufficient to ground a purely epistemic justification for the requirement that $\succsim$ be fully representable by a Dempster-Schafer belief function. However, it is not sufficient to justify the stronger requirement that $\succsim$ be fully representable by a probability function, or that $\succsim$ satisfy the qualitative constraint given by $A4$.

5.2 CC Doesn't Live Up to Expectations

Armed with a framework for quantifying the inaccuracy of comparative confidence orderings, the advocate of CC can now attempt to produce an expected inaccuracy argument for the rule, i.e. they

[33] As I noted in the previous section, Fitelson and McCarthy assume the Regularity condition, which I reject. Regularity plays no role in their proof that their scoring rule is the only evidentially proper linear rule, but some of the results stated in Theorem 5 may depend on the condition. This is not important here, since the current analysis does not rely on these results, which are only stated for the sake of completeness.

can attempt to show that CC always leads to those comparative confidence judgements that the agent initially expected to be the most accurate. If we assume that an agent A's initial ordering $\succsim$ is representable by a probability function, then for any P that represents $\succsim$ and any potential posterior ordering $\succsim^*$, we can define the expected inaccuracy of $\succsim^*$ according to P in the obvious way, i.e.

$$Exp(\succsim^* |P) = \sum_{w \in W} P(w)\mathfrak{I}_{\geq}(\succsim^*, w)$$

where $\mathfrak{I}_{\geq}$ is Fitelson and Mccarthy's scoring rule for confidence orderings. The most natural way to formulate the claim that CC minimises expected inaccuracy is as follows (assuming that the initial ordering $\succsim$ satisfies $\mathfrak{C}3$ and letting $R(\succsim)$ denote the set of probability functions that fully represent $\succsim$):

Expected Inaccuracy Minimisation for $\succsim$ (EIM$_{\succsim}$): For any proposition e, and any confidence ordering $\succsim^* \neq \succsim_e$ satisfying $e \sim^* \top$:

$$(\forall P \in R(\succsim))(Exp(\succsim^* |P) \geq Exp(\succsim_e |P))$$

EIM$_{\succsim}$ states that whenever an ordering $\succsim$ is probabilistically representable, every probability function that represents $\succsim$ will expect the posterior ordering produced by CC to be maximally accurate (amongst the set of orderings that accommodate the new evidence). If EIM is true, then we can at least be sure that whenever an agent's confidence ordering is representable by a probabilistic credence function, that function will expect the posterior ordering produced by CC to be maximally accurate. As it turns out though, *EIM$_{\succsim}$ is false.*

Theorem 6 *It is possible for there to exist confidence orderings $\succsim$, $\succsim^*$, a probability function P and a proposition e such that (i) P fully represents $\succsim$, (ii) $\succsim^*$ satisfies $e \sim^* \top$, and (iii) $Exp(\succsim^* |P) < Exp(\succsim_e |P)$.*

Theorem 6 establishes that it is in fact possible for CC to produce posterior orderings whose expected inaccuracy is non-minimal (amongst the class of posterior orderings that accommodate the evidence) according to some probabilistic representations of the prior ordering. Before considering the philosophical implications of this observation, it will be helpful to briefly sketch a case which illustrate's the theorem's truth.

Consider a situation in which $\mathfrak{B}$ is just the powerset of a set $W = \{w_0, w_1, w_2, w_3\}$ of four possible worlds. Let P be the probability distribution over $\mathfrak{B}$ generated by the following assignment: $P(w_0) = 0.65, P(w_1) = 0.17, P(w_2) = 0.18, P(w_3) = 0$, and let $\succsim$ and $\succsim_{w_1 \vee w_2}$ be the orderings that are fully represented by P and $P(-|w_1 \vee w_2)$, respectively (i.e. $\succsim_{w_1 \vee w_2}$ is the posterior ordering produced by CC upon learning $w_1 \vee w_2$). Finally, let $\succsim^*$ be the ordering that is identical to $\succsim_{w_1 \vee w_2}$ except that instead of ranking $\neg w_2$ directly above $\bot$, it ranks it directly below $\top$ (see below).

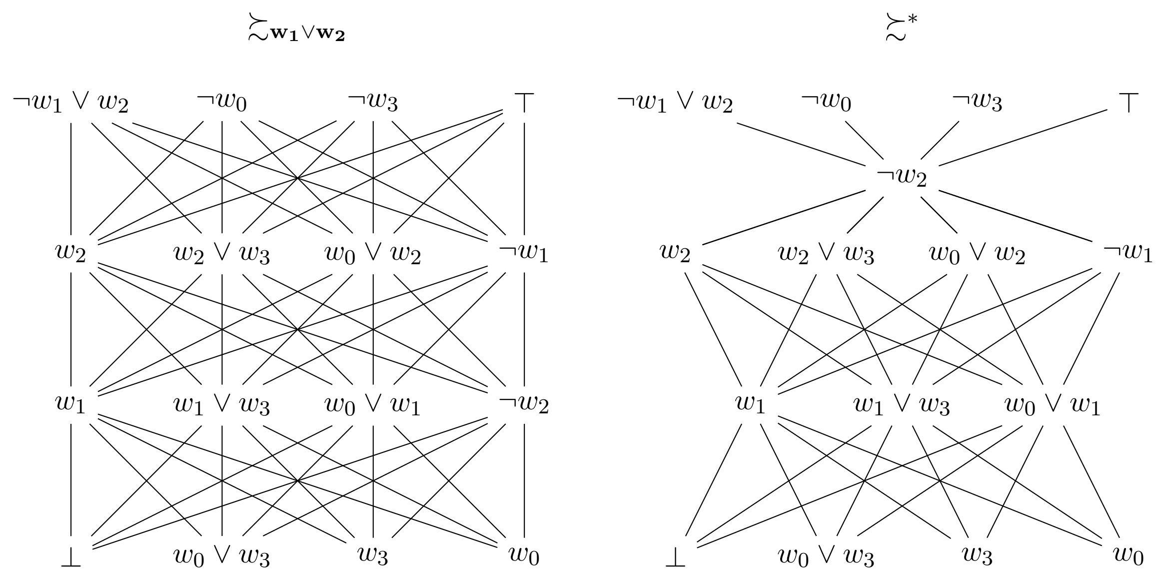

Then it is simple to verify that $Exp(\succsim^* |P) = 43.22$ and $Exp(\succsim_{w_1 \vee w_2} |P) = 47.04$, and hence that $Exp(\succsim^* |P) < Exp(\succsim_{w_1 \vee w_2} |P)$. This means that an agent who starts with the prior ordering $\succsim$ and learns the truth of $w_1 \vee w_2$ would expect $\succsim^*$ to be strictly more accurate than the posterior ordering produced by CC if they used the probabilistic representation P to calculate the expectation values.[34]

An important observation to make here is that the ordering $\succsim^*$ is not an extension of $\succsim_{w_1 \vee w_2}$ (in the sense defined in section 4) since $w_2 \succ_{w_1 \vee w_2} \neg w_2$ and $\neg w_2 \succ^* w_2$ both obtain. So it is not the case that the expected inaccuracy minimisation norm allows only for 'small' violations of CC, of the type which escape the Dutch book argument presented in section 4. If one aims to minimise expected inaccuracy whenever the prior ordering is represented by a probability distribution, then it is sometimes advisable, from the perspective of some probabilistic representations of one's prior confidence ordering, to adopt posterior orderings that strongly differ from those produced by CC, in the sense that they invert some of the 'more confident than' judgements yielded by CC updating. Thus, there seems to be a deep conflict between the following two prospective norms.

[34]The attentive reader might note that the $\succsim^*$ ordering depicted here is not probabilistically representable, although it is representable by a Dempster Schafer belief function. Thus, if one adopts the synchronic norm $\mathfrak{C}_3$, then adopting $\succsim^*$ as one's posterior ordering violates the requirement that updating on evidence should never lead one from a coherent prior ordering to an incoherent posterior ordering. This could provide an escape route for the CC advocate, as long as they were able to demonstrate that CC minimises expected inaccuracy amongst the set of posterior orderings that satisfy $\mathfrak{C}_3$ and accommodate the evidence. However, Fitelson and McCarthy showed that $\mathfrak{C}_3$ can *not* be justified by accuracy dominance arguments. The strongest synchronic norm that is justified by such arguments is $\mathfrak{C}_2$, which *is* satisfied by $\succsim^*$. Thus, if we are really interested in providing a full epistemic justification of CC, then there's no legitimate reason to assume $\mathfrak{C}_3$ at the outset.

(Sure Loss Avoidance (SLA):) An agent should never be willing to accept a series of bets which generate a sure loss.

(Expected Inaccuracy Minimisation (MIN):) Whenever an agent's prior confidence ordering is fully represented by a probability function P, they should update their ordering in a way that minimises expected inaccuracy according to P upon learning new evidence.

We've seen that, assuming the betting principles CBP2 and CBP3, coherence with SLA requires that agents always update their confidence judgements in a way that yields confidence orderings which extend those given by CC, i.e. their posterior confidence orderings should not invert any of the 'more confident than' judgements made by the CC ordering. In contrast, coherence with the MIN norm can sometimes *require* that agents adopt posterior orderings that are *not* extensions of the ordering generated by CC. So while the purely pragmatic SLA norm attributes *some* normative significance to CC, the epistemic MIN norm does not identify anything normatively privileged about the CC updating rule. In fact, it sometimes recommends that the rule be broken.

At this stage, the CC advocate is likely to ask for a justification of the normative status of MIN, i.e. they will ask 'why should we care about whether probabilistic representations of an agent's prior ordering always expects their posterior orderings to be maximally accurate?'. There are two natural ways to construe this question. Firstly, one could ask why an agent should care about what their prior epistemic state says about the expected accuracy of their posterior epistemic state *in general*. For, when the agent acquires new evidence, they are by definition forced to abandon their prior epistemic state and adopt a new one that accommodates the evidence. It's not obvious why one should generally care about what the old state, which had to be abandoned in the face of the new evidence, has to say about the new state. Variants of this reply have been made by i.e. Briggs and Pettigrew (2018) in response to Grieves and Wallace's expected inaccuracy justification for Bayesian conditionalisation. It's clear that if one accepts this response, then any prospective evaluation of CC that relies on a norm of expected inaccuracy minimisation can be summarily rejected.

However, there is another way for the CC advocate to resist the normative significance of EIM without rejecting the justificatory force of inaccuracy minimisation norms in general. Specifically, they could ask 'why should I care about what is expected by probabilistic representations of an agent's prior confidence ordering, when it may well be that the agent has no determinate credences, and hence that these probabilistic representations have no epistemic significance?'. At first pass, this looks like a plausible critique of the relevance of MIN for the epistemology of comparative confidence. We are working in a setting where it is not generally assumed that agents possess epistemic states with any structure over and above their comparative confidence judgements. So why should we care about the expectations of probabilistic representations of those judgements, when they have no a-priori claim to represent any structures that are actually present in an agent's epistemic state?

But despite its initial plausibility, this response doesn't pan out. While it is plausibly true that the expectations of probabilistic representations have no normative force in cases where the agent has no determinate credences, it is still prima facie desirable (for those that aren't generally sceptical of expected inaccuracy minimisation norms), *in the special cases where an agent's confidence ordering is set by determinate numerical credences*,[35] for agents to expect (relative to their actual credences) their posterior ordering to be maximally accurate. And agents that always update by CC will sometimes violate this desideratum.

Thus, CC's problem with inaccuracy minimisation is a grave one. It's not only that CC can't be justified by an expected inaccuracy minimisation norm like MIN. In fact, if we care about expected inaccuracy minimisation at all, then there are some cases (where the agent's confidence ordering is defined by a probabilistic credence function) in which updating by CC is normatively prohibited. As it turns out though, this problem is not unique to CC. It's easy to see that if one cares about expected inaccuracy minimisation, then any prospective update rule for comparative confidence orderings runs into similar difficulties. To illustrate, imagine a case in which an agent has determinate probabilistic prior credences, which define a corresponding prior confidence ordering. If the agent then learns some new evidence e, it is natural to think that they should update their epistemic state in a way that ensures that their posterior state has maximal expected accurate, relative to their prior credences. Now, if the agent cares about the expected inaccuracy of their posterior *credences*, then they will update by Bayesian conditionalisation. This will yield the same posterior confidence ordering that they would obtain by simply updating their prior ordering by CC. But in some cases, this will result in a posterior confidence ordering whose expected inaccuracy is non-minimal according to the agent's prior credences (by Theorem 6). Alternatively, in a case where CC fails to minimise expected inaccuracy, the agent could update in a way that does minimise the expected inaccuracy of the posterior ordering, and adopt a posterior credence function that represents this new ordering. But this will then lead them into conflict with Bayesian conditionalisation, and will yield posterior credences with non-minimal expected inaccuracy. In cases where CC fails to minimise expected inaccuracy, there is no way for the agent to update in a way that minimises the expected inaccuracy of both their posterior credences and their posterior confidence ordering.

[35] I take it to be uncontroversial that these cases exist, although one could in principle take the (in my view dialectically unenviable) position that rational agents never possess determinate credences, in which case this response would of course fail.

Thus, it's not simply that the norm of minimising expected inaccuracy conflicts with sure loss avoidance in the context of comparative confidence judgements. Rather, the expected inaccuracy minimisation norm conflicts *with itself* when it is simultaneously applied as a diachronic rationality constraint to both credences and confidence orderings. There are some situations in which the norm simply can't be simultaneously satisfied for both credences and comparative confidence judgements.

In fact, the preceding analysis could equally be construed as an argument against *Bayesian conditionalisation*. Whenever an agent with determinate credences updates those credences by Bayesian conditionalisation, they implicitly update their comparative confidence judgements by CC. And in some cases, doing so will yield posterior confidence orderings with non-minimal expected inaccuracy. So agents that care about the expected inaccuracy of their comparative confidence judgements would sometimes be well advised to update their credences in a manner that diverges from Bayesian conditionalisation.

6 Conclusion

We are left with a dilemma. CC is an intuitively compelling update rule that coheres with Bayesian conditionalisation and is partially justifiable by a Dutch book argument (assuming a weak and plausible decision theory for comparative confidence judgements). However, if we care about expected inaccuracy, then CC is sometimes normatively prohibited. More generally, it is sometimes impossible to simultaneously minimise the expected inaccuracy of both one's posterior credences and one's posterior confidence ordering. This seems to speak against the prima-facie desirable assumption that all learning scenarios allow for a learning strategy that is epistemically justifiable from the perspective of both credences and comparative confidence judgements. What's the solution?

The most obvious option here is to reject Fitelson and McCarhty's formalisation of the accuracy of comparative confidence orderings. However, doing so would incur a significant theoretical cost. For, their scoring rule follows directly from two compelling axioms: evidential propriety and linearity. So if we want to escape the dilemma by advocating an alternative conception of inaccuracy for confidence orderings, we need to reject either evidential propriety or linearity. As I've argued, evidential propriety is a highly desirable property, and it would be peculiar to care about the expected inaccuracy of one's future judgments without caring about the expected inaccuracy of one's present judgements, which is what evidential propriety requires you to do. So the only remaining option is to contend that the global inaccuracy of confidence orderings cannot be a linear function of the local inaccuracies encoded in those orderings. But at first glance, the prospects for identifying a privileged non-linear functional form in a principled way look dim. If the total inaccuracy of one's comparative confidence judgements isn't just the sum of the inaccuracy of the individual judgements, then what else could it be? In sum then, we are left with the following options. One can either

(i) Reject the normative significance of diachronic expected inaccuracy minimisation in general, or

(ii) Adopt *and justify* a non-linear (and evidentially proper) conception of the inaccuracy of confidence orderings on which CC does always minimise expected inaccuracy, or

(iii) Accept that the laws of diachronic rationality are sometimes cruel, and simply cannot be fully satisfied (since it is sometimes impossible to minimise expected inaccuracy for both credences and confidence orderings).

While the second horn of this dilemma will strike many readers as the most appealing, it is far from obvious that there exists a scoring rule for confidence orderings which is (i) philosophically principled, (ii) non-linear, (iii) evidentially proper, and (iv) for which CC always minimises expected inaccuracy. And until such a rule is found, the other horns of the dilemma need to be seriously countenanced.

Appendix

Proofs

Proof of Proposition 1: $\top \succ_e \bot \Leftrightarrow (e \wedge \top) \succ (e \wedge \bot) \Leftrightarrow e \succ \bot$. So $\succsim_e$ satisfies A1 if and only if $e \succ \bot$. To see that $\succsim_e$ satisfies A2 as long as $\succsim$ does, let $p \vdash q$. Then $(e \wedge p) \vdash (e \wedge q)$. So $p \succsim_e q \Leftrightarrow (e \wedge p) \succsim (e \wedge q)$, which is guaranteed by $\succsim$ satisfying A2. ∎

Proof of Proposition 2: Let $\succsim$ satisfy the condition. Proposition 2 guarantees that $\succsim_e$ satisfies A1 and A2. So we just need to show A3. Let $x \vdash y$, $y \succ_e x$ and let $y \vdash \neg z$. This implies that $(e \wedge x) \vdash (e \wedge y)$, $(e \wedge y) \succ (e \wedge x)$ and $(e \wedge z)$ is incompatible with $(e \wedge y)$. So, since $\succsim$ satisfies A3, we get that $(e \wedge y) \vee (e \wedge z) \succ (e \wedge x) \vee (e \wedge z)$, i.e. $e \wedge (y \vee z) > e \wedge (x \vee z)$, i.e. $(y \vee z) >_e (x \vee z)$, which proves that $\succsim_e$ satisfies A3, as desired. ∎

Proof of Proposition 3: Let $\succsim$ satisfy the condition. Then

$$x \succ_e y \Leftrightarrow (e \wedge x) \succ (e \wedge y) \Leftrightarrow (e \wedge x \wedge \neg y) \succ (e \wedge \neg x \wedge y) \Leftrightarrow (x \wedge \neg y) \succ_e (\neg x \wedge y) \quad ∎$$

Proof of Proposition 4: By definition, if $\succsim$ is fully representable by a probability function P, then $\succsim_e$ is fully representable by the function $P(-|e)$, which proves the proposition. ∎

References

Alon, S. and Lehrer, E. (2014). Subjective multi-prior probability: A representation of a partial likelihood relation. *Journal of Economic Theory*, 151: 476-492

de Finetti, B. (1937). La prevision: ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincare*, 7: 1-68.

Briggs, R. and Pettigrew, R. (2018). An Accuracy-Dominance Argument for Conditionalisation. *Nous*, https://doi.org/10.1111/nous.12258.

Eva, B. (Forthcoming). Principles of Indifference. *The Journal of Philosophy.*

Fine, T. L. (1973). *Theories of Probability.* Academic Press.

Fishburn, P. C. (1986). The axioms of subjective probability. *Statistical Science* 1(3), 335-345.

Fitelson, B. and McCarthy, D. (unpublished). *Toward an Epistemic Foundation for Comparative Confidence.* Manuscript.

Forrest, P. (1989). The Problem of Representing Incompletely Ordered Doxastic Systems. *Synthese*, 79:279-303.

Greaves, H. (2006). Justifying Conditionalization: Conditionalization Maximises Expected Utility. *Mind*, 459: 607-632.

Halpern, J. (2003). *Reasoning About Uncertainty.* MIT Press.

Harrison-Trainor, M., Holliday, W. and Icard, P. (2016). A note on cancellation axioms for comparative probability. *Theory and Decision*, 80: 159-166.

Hawthorne, J. (2016). A Logic of Comparative Support: Qualitative Conditional Probability Relations Representable by Popper Functions. In A. Hajek and C. Hitchcock (eds.), *The Oxford Handbook of Probability and Philosophy.* Oxford University Press: 277-295.

Icard, T. (2016). Pragmatic Considerations on Comparative Probability. *Philosophy of Science*, 83: 348-370.

Joyce, J.M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65(4): 575-603.

Joyce, J. (2010). A Defence of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives* 24(1): 281-323

Keynes, J. M. (1921). *A Treatise on Probability.* Macmillan.

Konek, J. (2019). *Comparative Probabilities.* In J Weisberg and R Pettigrew (eds), *The Open Handbook of Formal Epistemology*: 267 – 348.

Koopman, B.O. (1940). The axioms and algebra of intuitive probability. *Annals of Mathematics*, 41(2): 269-292.

Kraft, C. H., J. W. Pratt, and A. Seidenberg (1959). Intuitive probability on finite sets. *The Annals of Mathematical Statistics*, 30(2): 408-419.

Leitgeb, H. and R. Pettigrew (2010). An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philosophy of Science* 77: 236–272.

Lewis, D. (1976). The Probability of Conditionals and Conditional probabilities. *The Philosophical Review*, 85(3): 297-315.

Lewis, D. (1980). A Subjectivist's Guide to Objective Chance. In *Studies in Inductive Logic and Probability*, ed. R. Jeffrey, 263–293. Berkeley: University of California Press.

Lewis, Davi (1999). *Papers in Metaphysics and Epistemology.* Cambridge: Cambridge University Press, 1999.

Norton, J. (Forthcoming). Eternal Inflation: When Probabilities Fail. *Synthese.*

Pettigrew, R. (2016). *Accuracy and the Laws of Credence.* Oxford University Press.

Pettigrew, R. (unpublished). *What is conditionalization and why should we do it?* Manuscript.

Pettigrew, R. (forthcoming). *Dutch Book Arguments.* Cambridge University Press

Skyrms, B. (1980). *Causal Necessity.* New Haven: Yale University Press.

Scott, D. (1964). Measurement structures and linear inequalities. *Journal of Mathematical Psychology* 1(2): 233-247

Weatherson, B. (2007). The Bayesian and the Dogmatist. *Proceedings of the Aristotelian Society* 107: 169-185

Wong, S. M., Y. Yao, P. Bollmann, and H. Burger (1991). Axiomatization of qualitative belief structure. *IEEE Transactions on Systems, Man and Cybernetics*, 21(4): 726-734.