

Diachronic and Interpersonal Coherence

Kenny Easwaran and Reuben Stern

Short Abstract:

Bayesians standardly claim that there is rational pressure for agents' credences to cohere across time because they face bad (epistemic or practical) consequences if they fail to diachronically cohere. But as David Christensen has pointed out, groups of individual agents also face bad consequences if they fail to interpersonally cohere, and there is no general rational pressure for one agent's credences to cohere with another's. So it seems that standard Bayesian arguments may prove too much. Here, we agree with Christensen that there is no general rational pressure to diachronically cohere, but we argue that there are particular cases in which there *is* rational pressure to diachronically cohere, as well as particular cases in which *interpersonal* probabilistic coherence is rationally required. More generally, we suggest that Bayesian arguments for coherence apply whenever a collection (of agents or time-slices) have a shared dimension of value and an ability to coordinate their actions in a range of cases relevant to that value. Typically, this shared value and ability to coordinate is very strong across the time-slices of one human being, and very weak across different human beings, but there are special cases where these can switch—i.e., some groups of humans will have as much reason for their beliefs to cohere across a particular range of cases as the time-slices of one human usually do, but some time-slices of a human will have as much freedom to differ in their beliefs from the others as the members of a group usually do.

Long Abstract:

In the diachronic Dutch Book argument for conditionalization, it is argued that if an agent does not update by conditionalization, then she will accept a series of bets at different times that are jointly guaranteed to lose. David Christensen notices that there exists a parallel Dutch Book argument for interpersonal probabilistic coherence which shows that multiple agents will accept a package of bets that are jointly guaranteed to lose if their degrees of belief do not cohere with each other's. Since there is no general rational pressure for one agent's degrees of belief to cohere with another's, Christensen concludes that the diachronic Dutch Book argument cannot be used to justify conditionalization.

Though we agree that the existence of interpersonal Dutch Books tells against using diachronic Dutch Books to argue that conditionalization is a general requirement of rationality, we think that there are particular cases where diachronic Dutch Books reveal that conditionalization *is* rationally required, as well as cases where interpersonal Dutch Books reveal that *interpersonal*

probabilistic coherence is rationally required. More generally, we suggest that a Dutch Book argument applies whenever a collection (of agents or time-slices) have a shared dimension of value and an ability to coordinate their actions in a range of cases relevant to that value. We suggest that similar considerations apply also in the Accuracy framework of James Joyce and Richard Pettigrew.

Typically, this shared value and ability to coordinate is very strong across time-slices of one human being, and very weak across different human beings, but there are special cases where these can switch. We explore analogies from one side to the other - some groups of humans will have as much reason for their beliefs to cohere across a particular range of cases as the time-slices of one human usually do, but some time-slices of a human will have as much freedom to differ in their beliefs from the others as the members of a group usually do. Lastly, we consider an interpretation of these latter cases of believing (and reasoning from) a proposition for a particular fixed amount of time as supposition, rather than thinking of supposition as a distinct mental state from belief.

I. Introduction

According to Bayesian dogma, if you are rational, then your credences should be probabilistic at any given time, and the temporal evolution of your credences (from one probability distribution to the next) should accord with Bayesian conditionalization—i.e., when you learn e , you should adopt the probability distribution that results from conditioning your prior probability distribution on e .

What justifies these Bayesian norms of rationality? There are several ways to answer this question, but two of the most influential ways to motivate Bayesianism are pragmatic *Dutch Book* arguments and epistemic *Accuracy* arguments. Very roughly, Dutch Book arguments are founded on the pragmatic idea that the role of credences (or beliefs) is to promote desirable outcomes in practical decision-making scenarios. These arguments purportedly show that if you fail to conform to Bayesian norms, then you will make decisions that are guaranteed to result in a sure loss of desired ends. Accuracy arguments, by contrast, are founded on the purely alethic idea that credences (or beliefs) aim at the truth. They purportedly show that if you fail to conform to Bayesian norms, then you will judge your own credences to be a bad approximation of the truth. Since these two kinds of arguments are founded on distinct (perhaps incompatible) conceptions of belief (Easwaran, 2017), not every Bayesian must think that both are well-motivated. But it is nevertheless prudent to discuss both kinds of argument in the context of

justifying Bayesian norms since there are plenty of Bayesians who are sympathetic to Dutch Book arguments, as well as plenty who are sympathetic to Accuracy arguments.

In his (1991) paper, “Clever Bookies and Coherent Beliefs,” David Christensen takes no issue with synchronic Dutch Book arguments (that justify probabilism at a given time), but argues that diachronic Dutch Book arguments do not vindicate Bayesian conditionalization. His basic complaint (roughly) is that just as there are diachronic Dutch Books that appear to show that an agent’s credal states should cohere across multiple times, there are interpersonal Dutch Books that appear to show that multiple agents’ credal states should cohere with one another. Since it is clear that there is not, in general, rational pressure for one individual’s credences to cohere with another’s, Christensen concludes that diachronic Dutch Books prove too much.¹ Christensen does not consider Accuracy arguments for the simple reason that they were not prominent when Christensen wrote his paper, but it is easy to see that the same lesson applies. This is because one can transform any diachronic Accuracy argument into an interpersonal Accuracy argument in just the same way that one can transform any diachronic Dutch Book argument into an interpersonal Dutch Book argument—namely, by trading any mention of an agent’s credal state (or probability distribution) at multiple times for the credal states (or probability distributions) of multiple individuals.

In this paper, following, e.g., Gillies (1991), Kopec (2017), and Rowbottom (2013), we argue that there *is* sometimes rational pressure for distinct credal states to cohere (both interpersonally and intrapersonally), but we do not argue that this justifies Bayesian conditionalization as a universal norm of rationality. Instead, we argue that there are times when diachronic Dutch Book arguments and Accuracy arguments succeed at demonstrating that rational agents should conditionalize, but also times when they do not, as well as times when interpersonal Dutch Book and Accuracy arguments succeed, and times when they do not. Our starting point is that a Dutch Book or Accuracy argument is relevant to a *collective*—i.e., collection of individuals or time-slices—when, and only when, the collective shares interest in the value in question.²

Upon extending Dutch Book and Accuracy arguments to the group setting, it is immediately clear that there are cases where our treatment will say that there is rational pressure to cohere, but where there is no available credal state (or strategy) that guarantees coherence because the

¹ The same complaint does not arise for synchronic Dutch Books because these Dutch Books are about the betting behavior implied by having a *single* credal state. The problem with diachronic Dutch Books, according to Christensen, is that they succeed only if *distinct* “time-slices” of the individual accept bets together, and it is not clear that they should. (The point of Christensen’s interpersonal Dutch Book is to illustrate that we cannot generally help ourselves to the assumption that distinct credal states should accept bets together.)

² Strictly speaking, every collective is a collection of time-slices, but the time-slices can be located within one person, or distributed across multiple people.

individuals that constitute the collective cannot share information in a way that enables them to achieve coherence through coordination. Similarly, once it is realized that a given agent at a given time can enter into multiple collectives (e.g., with their other time-slices, and with other individuals), it seems that considerations of coherence can pull a time-slice in multiple directions that aren't jointly satisfiable. For example, even when an individual's time-slices have collective interest in accuracy, there may be reason not to conditionalize, provided that the individual's earlier and later selves sacrifice diachronic coherence in order to promote interpersonal coherence within their social groups.

In this paper, we address the first problem by arguing that when a collective cannot secure perfect coherence due to information-sharing constraints, the collective should aim to be as coherent as they can be (where, following, e.g., Schervish et al. (2002) and Staffel (2015), degrees of coherence correspond to the distance or divergence from perfect coherence). We then mostly leave the second problem for later since its solution appears to depend on the very difficult problem of how trade-offs should be assessed between competing shared interests (and, relatedly, on what it is for one collective to share interests more than another).

It is worth mentioning at the outset that there may be more to diachronic and interpersonal rationality than coherence. Although formal epistemologists have focused on giving arguments for collective coherence, philosophers of science, political theorists, and other epistemologists have given arguments for *diversity* of opinion within groups. We don't have a clear idea of how to compare these arguments with the ones given for coherence. But since they, like Accuracy arguments, are based on the aim of discovering the truth, there should be some broad perspective from which these competing demands can be considered. As with the problem of considering competing demands that arise from within the coherence framework, we mostly leave this issue for future work. But we do suggest that *supposition* may play two important roles in these settings. First, the agent may temporarily sacrifice coherence with one collective in order to promote coherence with another by temporarily supposing something. Second, if there is ever internal reason to decohere in order to achieve temporal or group diversity (for reasons not related to Dutch Book arguments or Accuracy arguments), then this, too, can be reached via local supposition.

II. Arguments for Synchronic Coherence

Suppose that you and two friends are discussing NASA's Mars 2020 program and the likelihood both that Congress won't cancel the program before the planned launch of the rover in 2020, and that the complicated design specifications will work out so that the rover successfully lands and returns data from Mars. You mention that you are only 40% confident that the rover will actually

launch and 30% confident that it will successfully land. One of your friends replies that she is 90% confident that the rover will actually launch and 60% confident that it will successfully land. Your other friend replies that he partly agrees with each of you—he is 40% confident that the rover will actually launch and 60% confident that it will successfully land.

Even if you think that neither of your friends' credences would be appropriate for you (perhaps because you think that your own credences are uniquely justified by your evidence), there is a sense in which you can be *sure* your second friend errs while your first may not.³ Specifically, your second friend is *incoherent*—viz., no matter what the evidence says, there is something wrong with being both 40% confident that the rover will launch and 60% confident that it will successfully land (since it is common knowledge that the rover can successfully land only if it launches).

It still may not be immediately clear *what* is wrong with your second friend's incoherent credences. After all, each credence had by your irrational friend is shared with someone rational. In order to answer this question, Bayesians standardly provide two types of arguments for the thesis that any rational credal state must be probabilistic—namely, Dutch Book and Accuracy arguments.

Let us first consider Dutch Book arguments. These arguments are premised upon treating credences as guides to action. For example, we might think of credences as betting dispositions—e.g., in the way that our confidence that a particular horse will win the race seems to underwrite dispositions to accept/reject specific bets on that horse. Or we might construe credences as more general guides to actions—e.g., in the way that your willingness to accept a job working for NASA in Pasadena may depend on your confidence that the rover will launch on time (because, e.g., it would be inconvenient to move to Pasadena, but the job would be better than your current job if the rover launched on time). On either interpretation, the credence represents the ratio of certain losses to conditional benefits that would make an action worth taking. In this section, we will focus on monetary gains and losses for ease of exposition.

If your second friend accepts that the rover can't land without having launched, then we can give an argument that he must have a credence at least as high in launching as landing. In the case discussed, he violates this by having a policy of selling bets on “the rover will launch” at over

³ Some philosophers argue that one's evidence always uniquely determines the rational credence for one to have in each proposition, so that you and your first friend could both be rational only if she has different evidence than you. Others argue that evidence is permissive, so that you and your first friend might both be rational with the same evidence. The issues addressed in this paper are independent of the status of this debate, and focus on the explanation of why the *second* friend *couldn't* be completely rational. See Kopec and Titelbaum (2016), Horowitz and Dogramaci (2016), Schoenfield (2014).

40% and of accepting bets on “the rover will successfully land” at up to 60%, and as a result he can be “Dutch-booked”—i.e., he will accept packages of bets that are jointly guaranteed to result in sure loss. For example, someone might offer to pay him \$45 for a bet of \$100 that he will lose if the rover launches, and then offer to charge him \$55 for a bet of \$100 that he will win if the rover lands. Since these prices are favorable compared to his credences, he would accept each. This means that he spends a total of \$10 up front, and has no chance of winning the second \$100 except in circumstances where he loses the first \$100. He can already recognize that this is a bad deal.

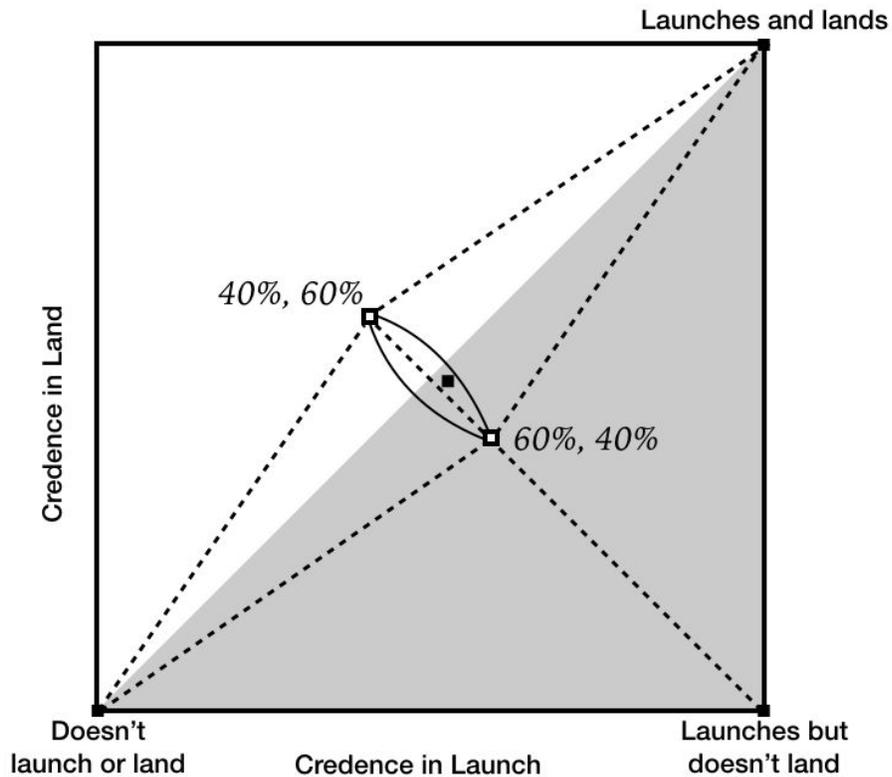
This sort of Dutch Book argument provides one explanation for why his credences should be coherent in this way. The three general axioms of probability theory can be derived from coherence arguments like these—i.e., if an agent (i) has negative credences, (ii) doesn’t assign credence 1 to certainties, or (iii) fails to have additive credences for the disjunctions of incompatible events, then a similar set of bets can be made that he recognizes to guarantee bad results by his own lights. Among the consequences of these axioms are that an agent’s credences in complements must add up to 1, that entailment results in increased probability (which is relevant to this case), and all the more specific implications of mathematical probability theory. Thus, if it’s reasonable to assimilate credences with betting dispositions (especially allowing that the costs and benefits might occur in non-monetary form), then we can use Dutch Books to argue that credences should be probabilistic.

As we mentioned in the introduction, it is also possible to argue that credences should be probabilistic even if one rejects the idea that credences are evaluated pragmatically on the types of actions they license. Accuracy arguments are a bit more complex than Dutch Book arguments—and we will refrain from delving into the full mathematical details here—but they start with the idea that credal states are evaluated in terms of how well they approximate the truth. Like William James (1896), then, the authors of Accuracy arguments start with the thought that a belief is good to the extent that it approximates the truth, and bad to the extent that it approximates the false, though individuals might differ in exactly how much they care about these two aspects of accuracy. The particular profile is expressed by a “scoring rule”, which is a mathematical function that has some particular features. We refrain from discussing the details here for ease of exposition, but they (along with the more formal versions of these arguments) can be found in the literature on this topic. (Joyce, 1998; Pettigrew, 2016))

The Accuracy argument for coherence goes like this. If your second friend is 40% confident that the rover will launch and 60% confident that the rover will successfully land, he doesn’t know exactly how accurate his credences are overall. But as we will show, he *does* know that there are other credal states he could have adopted that *weakly accuracy-dominate* his own credal state—i.e., that are at least as accurate in every possible state of the world, and that are more

accurate in some possible states of the world. As long as the different propositions contribute equally to his overall accuracy, then a credence of 60% that the rover would *launch* and 40% that it would *land*, would be equally accurate as his actual credences in the cases where both propositions are true (i.e., the rover both launches and lands) or where both are false (i.e., the rover neither launches nor lands) and strictly *more* accurate if the rover launches without landing.

If the metaphor of “closeness to truth” is appropriate for understanding accuracy, as it is to some extent under all scoring rules, then the following diagram illustrates the situation:



The square indicates the set of possible pairs of credences in the two propositions, with the perfectly accurate credences for the three possible situations indicated in three of the corners. The arcs indicate the sets of credences in these two propositions that are equally accurate as the actual credences in two of the three possibilities. Any pair of credences in the region bounded by those arcs is strictly more accurate than the actual credences of (40%, 60%) in all three possibilities. For credences in the white region in the upper left, a similar pair of curves gives a set of alternative credences that are guaranteed to be more accurate in all possibilities. But for credences in the shaded region, to get closer to one possibility requires getting farther from one

or both of the others. This diagram indicates the Accuracy argument for the rule that one shouldn't have higher credence in one proposition than in another proposition entailed by it. Parallel arguments again justify all the rules of probabilism.

III. Arguments for Diachronic Coherence

There also versions of Dutch Book and Accuracy arguments that apply to an agent's credences at different times.

Suppose once more that you are currently 40% confident that the rover will launch and 30% confident that it will successfully land. Because you are 40% confident that it will launch, you will be willing to sell a bet that costs \$30 and returns \$75 if the rover launches. (This is just a scaled-down version of the bet that costs \$40 and returns \$100.) Because you are 30% confident that it will land, you will be willing to buy a bet that costs \$30 and returns \$100 if the rover successfully lands. With both bets together, you will lose a total of \$75 if the rover launches and doesn't land, will win a total of \$25 if the rover launches and lands, and will come out even if the rover doesn't launch. This is equivalent to a *conditional* bet, costing \$75, for potential winnings of \$100 if the rover lands, where the bet is called off (and your initial payment returned) if the rover fails to launch. This sort of conditional bet is sometimes taken to indicate a kind of "conditional credence" in one proposition conditional on another.

Now suppose that you learn of the rover's successful launch and then come to assign some credence other than 75% to the rover landing—i.e., you fail to update your credences by Bayesian conditionalization. It can be shown that you will now accept bets that, when considered in a package with the bets that you accepted before, are guaranteed to result in a net loss. For example, if you become 50% confident that the rover will not land, then you will accept \$50 for a bet that costs you \$100 if the rover lands. If you already made the earlier conditional bet, then you will now lose a total of \$25, no matter whether the rover lands or not. By adding a small bet where you win a little if the rover launches and lose a little if it doesn't, a clever bookie could put together a strategy for betting with you, from before you learn whether the rover launches until afterwards, that guarantees a net loss for you no matter what happens.

Similar versions of this argument exist for accuracy. We spare the reader the details, (see Greaves and Wallace (2006) and Easwaran (2013) for arguments about conditional credence, and Briggs and Pettigrew (2016) for a true diachronic Accuracy argument) but their upshot is that if an agent plans to update in a way other than by conditionalization for some duration of time, then there is some alternative plan that she could have adopted that she would recognize to be more accurate across that time than her actual plan.

It thus seems that agents not only have strong reason to be coherent *at* a given time, but also to be coherent *across* time. While coherence at a given time requires satisfying the probability axioms, coherence across time requires updating via conditionalization (if changes are driven by coming to be certain of some proposition).

IV. Christensen's Double Agent Dutch Book

Christensen (1991) is not impressed with these arguments for diachronic coherence. He argues against Dutch Book arguments for diachronic norms of coherence on the grounds that (i) it is possible to develop parallel interpersonal Dutch Book arguments, and (ii) there is clearly not rational pressure for multiple individuals' credal states to cohere with one another. In order to argue that there is no such rational pressure, Christensen introduces the case of *The Double Agent Dutch Book*:

“The Double Agent Dutch Book: Suppose that I am shopping with my wife. My credence in rain today is 25%. My wife, who is somewhat more pessimistic than I, sets the probability of rain at 50%. I am approached by a bookie, who offers to bet me \$1 to my \$3 that it will rain (that is, he will win if it rains). Given my credence, I regard this bet as fair, accept it, and go back to weighing out wax beans. The bookie then approaches my wife, offering her a bet at \$2 to \$2, which he will win if it doesn't rain. Given her credence, she regards this bet as fair, and accepts it. The bookie has now assured himself of a \$1 profit: if it rains, he gets my \$3 and pays my wife \$2; if it doesn't, he pays me \$1 and gets \$2 from my wife.” (Christensen, 1991, pp. 239-240)

The Double Agent Dutch Book resembles a diachronic Dutch Book in every way but one—namely, while diachronic Dutch Books are about the credal states and betting profiles of one individual at two distinct times, the Double Agent Dutch Book is about the credal states and betting profiles of two distinct individuals. Thus, just as diachronic Dutch Books show that the time-slices of an agent will collectively accept bets resulting in sure loss when their degrees of belief do not cohere, the Double Agent Dutch Book shows that the distinct individuals that comprise a group will collectively accept bets resulting in sure loss when their degrees of belief do not cohere.

The normative upshot of the Double Agent Dutch Book is controversial. Christensen uses it to *modus tollens* diachronic Dutch Books—viz., he thinks that there is obviously no rational pressure to cohere with his wife (for the simple reason that it seems entirely reasonable for their opinions to diverge about the weather, despite the shared bank account), and that

collective (diachronic or interpersonal) Dutch Books therefore fail to establish norms of rationality. But following Gillies (1991), Rowbottom (2013) draws a different lesson from Christensen's example:

“Surely Christensen cannot think that it *is* rational – in an ordinary language sense of ‘rational’ – for a husband and wife to make individual bets with a bookie such that they lose joint funds whatever happens, and then just continue to make such bets, come what may, despite having communication channels by which they could easily prevent this occurring (and joint interest in not losing money for no utility)? But it would appear that this is a consequence of denying the notion that avoiding Dutch Books at the group level is ever a (practical) rationality requirement. Before I ‘bet’ with a significant sum of money shared with my wife, e.g. in making an investment, I discuss this with her. And the same is true, *mutatis mutandis*, when she wishes to ‘bet’. I trust we are not the only working couple with a joint bank account to take this precaution!” (Rowbottom, 2013, p. 195)

Rowbottom's point is that when we take stock of the couple's shared interest, and when the amounts of money at play are reasonable stand-ins for utility, it seems like there may be rational pressure for Christensen and his wife to coordinate their bets (and credences) after all. Thus Rowbottom opts for *modus ponens* where Christensen opts for *modus tollens*—viz., he maintains that when some individuals have shared interest and are in a position to coordinate their credences or bets (as both Christensen and his wife are, and as the time-slices of an individual plausibly are), there *is* rational pressure for the collective to cohere.

Who is right? Our own take is closer to Rowbottom's. We agree with Rowbottom that there are circumstances in which there is rational pressure for a collective to cohere (e.g., when he and his wife coordinate their investments). But we also believe that Christensen's argument does reveal that there are no universal norms of diachronic coherence—including even Bayesian conditionalization! There are collections of time-slices that either do not share interests, or are in no position to cohere because there is limited information flow between the time-slices.

More generally, our view is that norms of rational belief (or credence) do not derive from some metaphysically primitive notion of agential identity, but rather derive from the functional role that credences play in practical or epistemic agency. Each type of agency involves control over a set of behaviors or beliefs, and is directed towards some goal (whether pragmatic, as in Dutch Book arguments, or alethic, as in Accuracy arguments). When there is a unified goal to a set of beliefs, and an ability to structure them, we think that considerations of rationality mandate that these beliefs to work together in a way that properly aims at this goal. But when

a set of beliefs are not unified in their pursuit of a goal, or are not causally connected in a way that allows for structural organization, then rational standards do not apply.

In order to motivate our view, it is helpful to consider some further specifications of Christensen's case in which the standards of rationality clearly do apply, and some where they clearly do not.

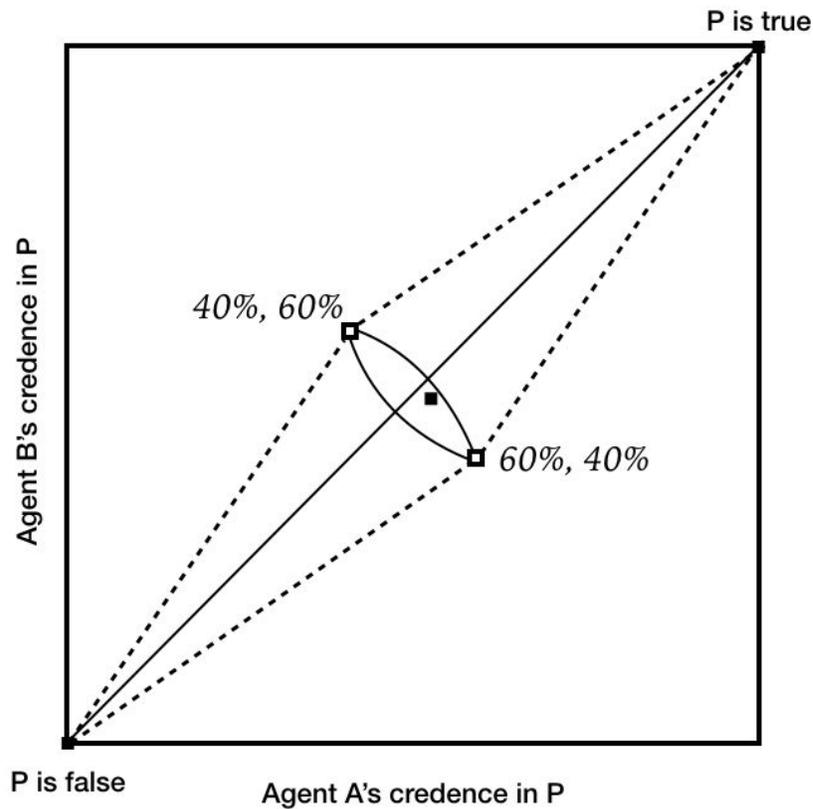
Let us suppose that Christensen was describing a gambling incident as they stood in the grocery store parking lot (before entering the store). In this case, it seems that Christensen is right—i.e., since the couple's collective fate did not seem to depend on their beliefs about rain for this duration of time, it seems entirely appropriate for their opinions to differ. But things plausibly change when the couple is inside the grocery store, shopping for tonight's dinner. Suppose, for example, that the couple is shopping the day before leaving for vacation, for a meal that they plan to eat that evening. If she has high credence in sun, then she might get a picnic basket with strawberries and cheeses to eat outdoors. If he has high credence in rain, then he might get ingredients for a warm roast to have by the fire. Each action corresponds to a gamble that is reasonable under one credence function but not the other. But if both are taken, they spell disaster by ensuring that the couple is out some money and only eat one of the meals (or, worse, that they have a strawberry cheese roast).

Since the purpose of the shopping trip was to organize their collective resources in order to promote their shared desire for a pleasant meal together, it seems that the couple has erred by failing to cohere. That is, since the couple *could* have better realized their shared goals by communicating with each other and coordinating their credences, it seems that they *should* have done so. (Perhaps this would have ended up with them purchasing ingredients for a pasta that would be fine to have regardless of weather.) If credences about rain are relevant to the goals of their shopping trip, then bad decisions will be made if they don't cohere. But if credences about rain are not relevant to the shopping trip, then from the perspective of this collective activity, there is no rational reason to cohere. There might be unfortunate consequences if the parking lot happens to contain a clever bookie, but this possibility is usually not one that couples need to consider in order to be rational.

On the Accuracy conception there can also be a kind of pressure for interpersonal coherence if several agents see each other as aiming for a sort of "collective accuracy", and can coordinate their beliefs. We need to assume that this notion of collective accuracy satisfies some particular formal constraints (in particular, that it is convex, so that the total accuracy of two people with different credences in the same proposition is less than the total accuracy they *would* have if they both adopted some particular intermediate credence). But with that

assumption (which is satisfied by all standard scoring rules), we can show that a pair of people who can coordinate their credences, and aim at collective accuracy, have rational pressure to have the same credences.

Consider the following diagram. The two axes represent the credences of the two agents in a single proposition P . The upper-right corner represents the situation where the proposition is true, and the lower-left represents the situation where the proposition is false. The total accuracy of the two agents in either situation is represented by something like the distance from the point representing the pair of credences to the point representing the situation that is actual.



The curved lines indicate the sets of points with the same total accuracy for the pair as (40%, 60%) in each of the two situations. If the scoring rule is the Brier score,⁴ then these arcs will be perfect circular arcs centered at the corners, but for any convex scoring rule the arcs will curve

⁴ The Brier score is probably the most frequently deployed scoring rule in Accuracy arguments. As Pettigrew (2015) reports, “the Brier score measures the inaccuracy of a credence function at a world as follows: it takes each proposition to which the credence function assigns credences; it takes the difference between the credence that the credence function assigns to that proposition and the ideal or vindicated credence in that proposition at that world; it squares this difference; and it sums up the results.”

in this sort of way. Any pair of credences represented by a point in the region bounded by these two arcs will be strictly better from the perspective of total accuracy than the credences of (40%, 60%). The only way to avoid this possibility is for both to have the same credence, so the point representing them lies on the diagonal line.

One might worry that there is no reason for interpersonal collective to care about their *total* accuracy. However, the Accuracy justification of a truly diachronic norm given by Briggs and Pettigrew (2016) (unlike the “pseudo-diachronic” norms justified by Greaves and Wallace (2006) and Easwaran (2013)) requires consideration of total accuracy across time-slices, and we think that whatever motivates this concern for total accuracy across time can also motivate concern for total accuracy of a group.

Sarah Moss (2011) uses concern for total accuracy in a group to motivate a version of the “equal weight view” for disagreement. She argues that if two agents use the same scoring rule, then they ought to compromise by “splitting the difference” and jointly adopting the credence that is the average of their credences, because this one has the best total expected accuracy. We don’t necessarily think that total *expected* accuracy is the method they should use to negotiate their compromise. But if they care about total *actual* accuracy, and have some means of negotiating a compromise, then they can recognize that *any* pair of credences inside this region would be *strictly* better than their current situation of different credences. This region might have an odd shape if they use different scoring rules for the credences of the different individuals, or give different weights to them, but as long as the scoring rule is convex, there are many compromises that would be an improvement if they don’t share the same credence.

While there is no *specific* compromise among different credences that would be required, as long as a group does *in fact* share an interest in total accuracy, and has some means of coming to a compromise, they are collectively incoherent if they fail to do so.

One might wonder about which conditions are such that individuals care about the total accuracy of the collective, and not just about their own individual accuracies. These conditions are ones that motivate tradeoffs among the accuracy of different individuals on the same proposition. Regardless of whether the proposition in question is true or false, one person’s credence will become more accurate while the other’s becomes less. One might worry that taking stock of these tradeoffs violates an intuitive notion of “separateness of persons” for epistemology—i.e., that the epistemic value of one individual person’s credal state cannot be traded off for the epistemic value of another’s. But as Selim Berker (2013a,b) argues, even intrapersonal Accuracy arguments (or Dutch Book arguments, for that matter) require the

violation of an analogous “separateness of propositions” insofar as they are premised upon the existence of tradeoffs between the epistemic value of credal attitudes towards distinct propositions. According to Berker, this gives us reason to abandon any sort of consequentialist or teleological argument for coherence. But we think that some such teleological conception of the requirements of rationality (e.g., that they are derived either from the goal of using credences to guide actions, or from the goal of having accurate credences) is the most promising way of understanding coherence requirements. And if tradeoffs among different propositions can be motivated for a single person that sees her credences in different propositions as part of a single epistemic project aimed at overall accuracy, then tradeoffs among different persons can be motivated for a group that sees their credences as part of a single epistemic project aimed at overall accuracy.

Rather than arguing for some particular conception of “sharing values” or “ability to coordinate” as uniquely relevant to an overarching notion of rationality, we think that *any* meaningful precisification of these two concepts gives rise to *some* conception of rationality that can apply either diachronically or interpersonally. In the next section, we present some examples that should give the reader a better idea of the sorts of conceptions we have in mind.

V. Conditions for Collective Coherence

The first case that draws our attention involves Paul and Mitch, two congressmen from the same political party (albeit in different houses of congress) who both want to do what they can to promote their party’s agenda in their respective branches of congress. Now suppose that Paul has credence $\frac{2}{3}$ that the stock markets will go up in March while Mitch has only credence $\frac{1}{3}$ that the stock markets will go up in March. Given Paul’s credence, he might take some action that ties the political party’s fortune to the markets so that the party gains 1 point in the polls if the markets go up and loses 2 points if they go down. Given Mitch’s credence, he might take some other action that results in a gain of 1 point if the market goes down and a loss of 2 points if it goes up. No matter whether the market goes up or down, the result of Paul and Mitch’s collective actions will be a loss of 1 point. Are Paul and Mitch collectively irrational?

Just as it seemed clear that Christensen and his wife should have coordinated in order to avoid the strawberry cheese roast, it seems clear that Paul and Mitch should coordinate their

credences in various propositions relevant to their foreseeable political actions.⁵ This is something that politicians can (and often do) manage, and if they fail to do this, their political party will suffer. So this case squares with the view that if a collective is fully able to share information, and likewise shares interest, then there is rational pressure for them to cohere. But is the converse true? Must collectives share interest and be fully able to share information in order for there to be rational pressure to cohere?

Consider Paul's relationship with Nancy, who is a member of the opposing party. Even though Paul and Nancy's actions may also collectively result in a net loss for Paul's party, no matter how the market performs, this would reflect no failure of rationality. This is because Nancy has no reason to care whether they collectively do well for Paul's interests. She may in fact *prefer* this sort of outcome!

Similarly, we can consider Sheldon, who is a billionaire in charge of a Political Action Committee, who shares Paul's concern for the party, but is legally barred from coordinating political strategy with him. If Sheldon and Paul fail to align their credences, then they may behave in ways that are collectively unfortunate for the party. But it wouldn't be a sign of irrationality, given the legal barrier to coordination.

The political arena is useful for demonstrating when and whether there is rational pressure for collectives to cohere that are jointly interested in some *practical* good, but since the goals of politics often stray far from the truth, there are better places to gauge when and whether there is *epistemic* reason for a collective to cohere.

Suppose that Rosalind, James, and Francis are working together on a shared epistemic project of trying to understand the structure of some biological molecules, and that they can coordinate their beliefs, because they work down the hall from each other. Collective Accuracy arguments seem to show that they will do a better job of promoting their shared goal of accuracy on this subject matter if their credences are collectively coherent than if they are not. This sort of collective coherence may be particularly important if all of them know that Rosalind's expertise is very different from that of James and Francis.

However, if the project is not shared, or if there is no possibility of coordination, then there may be no failure of rationality on the part of the project if the beliefs don't cohere.

⁵ It is important to keep in mind that we claim that Paul and Mitch should cohere *given* their shared desire to promote their party's agenda. We certainly do not mean to claim that Paul and Mitch *should* desire to promote their party's agenda.

Consider Edward and Othniel, two paleontologists who are both interested in knowing more about the Cretaceous dinosaurs of western North America, but who view each other as competitors in the so-called “Bone Wars” (rather than as collaborators in the pursuit of knowledge). Suppose that Edward is 80% confident that the head of *Elasmosaurus* belongs on the *short* sequence of vertebrae, while Othniel is 80% confident that the head of *Elasmosaurus* belongs on the *long* sequence of vertebrae. The two researchers would be sure to have greater *total* accuracy if they were both 50% confident of each end, but if they don’t care about each other’s accuracy, then they can each reasonably take their separate gambles on the truth.⁶

Similarly, suppose that it’s 1965 and that Julia and Yuri are both interested in Hilbert’s tenth problem, on the solutions of systems of diophantine equations. Julia is quite confident that if the Fibonacci numbers are diophantine then Hilbert’s tenth problem has no solution, and Yuri is quite confident that the Fibonacci numbers are diophantine, but neither is very confident about Hilbert’s tenth problem. If Julia is working in Berkeley and Yuri is working in Leningrad, and the Cold War prevents them from communicating, then there is no irrationality here, but rather just tragedy.⁷

It thus seems that there is rational pressure for an interpersonal collective to cohere when the members of the collective can share information and share interest in some value (whether pragmatic or alethic), and that there is no rational pressure for an interpersonal collective to cohere when the members of the collective either cannot share information or lack shared interest. We do not attempt to give a detailed characterization of what it means for Rosalind, James, and Francis to share an epistemic value, such that Julia and Yuri count as well, but Edward and Othniel do not. One approach to this question is suggested by Kopec (2015), but we think more research on this question would be fruitful.

This covers the interpersonal case, but there is reason to wonder whether these considerations carry over to the diachronic case. After all, one could attempt to vindicate Bayesian norms of diachronic coherence as universal norms of rationality by arguing that Christensen’s analogy is

⁶ There may be separate irrationality if, say, the competition between the two drives them both to bankruptcy, and results in the destruction of many fossils. But if each individual were entirely motivated by his epistemic goals, then no *individual* may have been irrational here, and the lack of shared interest means there is no clear *group* to have been irrational. For more details on this period in the history of paleontology, see Shor (1974) or Jaffe (2000).

⁷ There may have been a separate irrationality in the political situation that led to the ban on communication between American and Soviet mathematicians, but it is not an irrationality of the mathematicians themselves.

otiose because (i) the temporal parts that comprise a person always share practical and/or epistemic interest and (ii) are in a position to share information in a way that make coordination possible.

Though the time-slices that comprise a person for some duration of time very often have these features since (i) people tend to see their lives as unified practical and epistemic projects and (ii) memory and intention serve to provide strong coordination among temporal parts, there are cases where collections of time-slices lack one or both of these aspects.⁸ Consider a lawyer who at different times is tasked with arguing that distinct individuals are guilty of one and the same murder.⁹ Or consider Saul, who sets out for Damascus with one set of (pragmatic and alethic) goals, but who undergoes a drastic shift in values along the road. He may reasonably see his earlier self as a competitor rather than a collaborator, and might find that no particular purpose is served by lining his beliefs up with those of his earlier self. (“Transformative experiences” of the sort discussed by Paul (2014) make this phenomenon especially sharp.)

It is easier to identify collections of time-slices that lack the ability to coordinate. For example, when a person suffers from deficit or disorder in memory or planning, their time-slices will be less able to coordinate. Thus there are individual people who are more like Julia and Yuri than Rosalind, James, and Francis, and our account says that there is no rational pressure for them to diachronically cohere.¹⁰ Similarly, even if someone suffers from no disorder, they may have trouble remembering what they did or learned last night, e.g., because they had a bit too much to drink. In such circumstances, it might have been irrational for the agent to drink because of its effect on the agent’s ability to promote their desires over the time in question, but it is not obvious that the entire collection of time-slices rationally err insofar as their credences fail to cohere (since the earlier time-slices made coordination impossible).¹¹

⁸ See Bratman (1987), Korsgaard (2009), and Velleman (2000) for discussion of the first point and Bratman (1987) and Holton (2009) for discussion of how intentions play this binding role.

⁹ Armstrong (2017) reports this phenomenon.

¹⁰ This aspect of our account is seemingly relevant to the literature in formal epistemology on how agents should update their beliefs in cases involving memory loss. See Talbott (1991), Elga (2000), and Arntzenius (2003) for examples of the complexities of dealing with such cases.

¹¹ Cases like these may point towards one of the ambiguities of our account. We say that there is rational pressure for a collective to cohere when they could have coordinated, or when the causal structure of the world permits coordination. But what do we mean by “could” and “permits” here? *Could* the drinker’s credences have been coordinated in the relevant sense? By our lights, this is one of the places where our view yields distinct conceptions of rationality as special cases. On the one hand, it is possible to argue that there is rational pressure for the members of a collective to cohere only if *each individual member* could have cohered, where an individual member counts as incapable if one of the other members does or believes something that makes it hard for them to cohere. According to this conception, the later time-slices of the drunk could not have cohered, and there is thus not pressure to cohere. On the other

We take these examples to reveal that diachronic coherence and interpersonal coherence can productively be thought of as two sides of the same coin, where there is rational pressure for a collection of time-slices to cohere (regardless of whether they are in one person or multiple people) whenever they collectively share interest in some value (e.g., money, political fortune, or purely epistemic accuracy) and are causally connected in a way that enables them to share information and coordinate their credences.

Before proceeding to discuss competing coherence requirements and degrees of coherence, it is worth briefly mentioning the possibility that an individual may fail to satisfy the requirements for rational assessment even *within* one time-slice. Suppose, for example, that David has beliefs about local streets at t , has beliefs about local trains at t , and that these packages of beliefs are integral to separate sets of plans (navigating town on the streets, and traveling to distant cities on the trains) and separate projects of aiming at the truth. If there is some difficulty in coordinating all of David's different actions and epistemic projects at t , this may not constitute a failure of rationality when David's beliefs across these different domains fail to cohere. (Lewis, 1982) This sort of thing may happen quite frequently when we compartmentalize various distinct projects that come with their own beliefs and plans for local action—e.g., our academic projects, our social projects, and our political projects. Such personal incoherence is a common phenomenon, and we suggest it is less of a problem than epistemologists might traditionally think. Regardless, in what follows, we talk as though time-slices are the atoms that comprise (diachronic or interpersonal) collectives for ease of exposition, and, in so doing, ignore questions about synchronic intrapersonal coherence.

VI. Degrees of Coherence

At this point, it is natural to wonder whether the demand to cohere can vary in degree from one collective to the next. After all, when multiple people are involved, the extent to which individuals can coordinate their credences can vary quite a bit, and it may seem that the demands of rationality should take this variance into account. For example, the two of us can often easily talk about this paper over the computer, but there are times when we cannot reach other (since we live on different continents). It may seem that the demand to cohere should be

hand, one can argue that an individual member counts as capable of cohering whenever there is something the *entire collective* could have done or believed to achieve coherence. According to this conception, the later time-slices could have cohered since the earlier time-slices could have refrained from drinking so much.

more demanding during times when we're talking than when we can't reach each other. Similarly, even within one person, the ability to coordinate can vary. For example, as time passes, we are more and more likely to forget what we once knew, and it thus seems more difficult for present time-slices to coordinate with distant-future time-slices than with near-future time-slices (even if these time-slices are part of the same plan or epistemic project).

Furthermore, the extent to which values and desires are shared also comes in degrees, particularly given the unclarity around the idea of what it takes to be part of a single epistemic project. But it's also true for successive time-slices of the self—conversion and change of values often happen gradually rather than all at once.

In these cases, we think it is useful to consider some work on *degrees of coherence*. There are several proposals for how we can measure *the extent to which* a credence function is coherent (or approximates coherence), some of which are developed in parallel with the pragmatic and alethic arguments for coherence. For example, Schervish, Seidenfeld, and Kadane (2002) and Staffel (2015) suggest that in addition to caring whether or not a Dutch Book can be made against you, we might care about how *big* the loss is. Along these lines, they propose some particular ways of measuring how vulnerable credence functions are to Dutch Books, and argue that we should aim to *decrease* this vulnerability even if we can't fully eliminate it. Similarly, in the case of Accuracy arguments, de Bona and Staffel (2017) propose an alternative measure of degree of coherence based on the scoring rule that is used to evaluate closeness to the truth—it can also evaluate closeness to *coherence*.

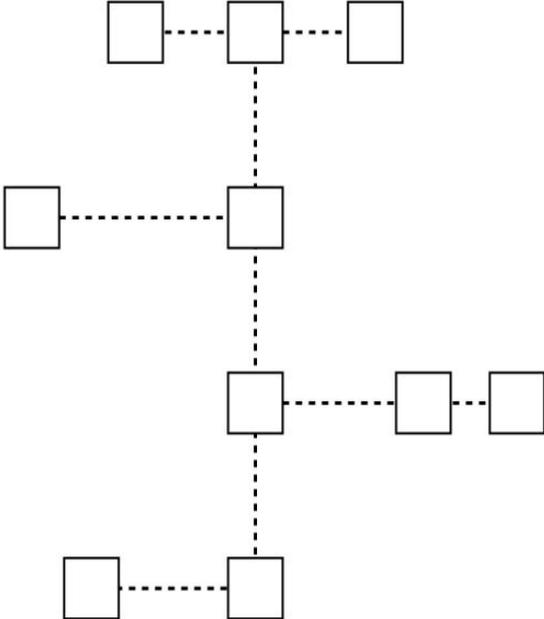
In this paper, we do not wish to commit ourselves to some particular way of measuring degrees of coherence (or degrees of approximation to coherence). We think that any of these methods may be useful when determining *how* coherent a collective should be, so long as the evaluation of any given collective takes into account the degree of shared value and the degree of ability to coordinate. Our claim is simply that collectives that share *any* pragmatic or alethic interest will also share an interest in approximating coherence. Different collectives have different amounts of shared interest, and will thus have different degrees of interest in the degree to which they approximate coherence, as well as different abilities to ensure they do so.

VII. Competing Coherence Requirements

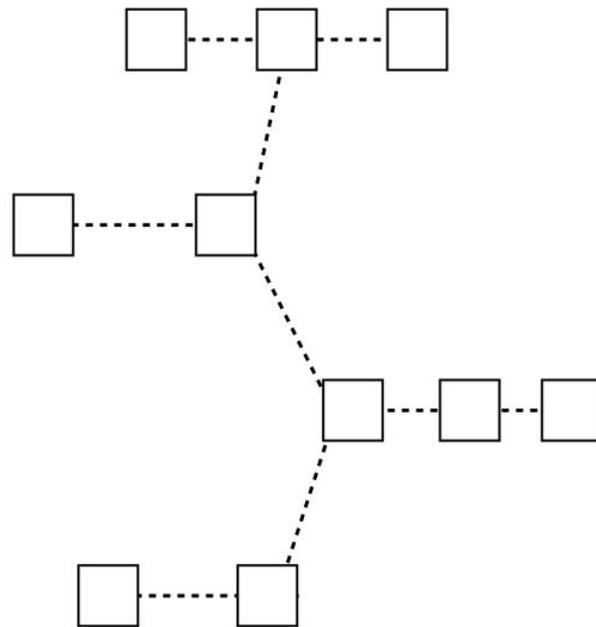
Just as a collective of time-slices may be prevented from achieving perfect coherence because its members cannot readily communicate with each other, an individual time-slice can be

prevented from meeting all of its coherence requirements because it enters into multiple collectives whose corresponding coherence requirements are not jointly satisfiable. This can happen because neither shared interest nor ability to coordinate is transitive—i.e., one time-slice can share some interests with a second time-slice, who shares other interests with a third time-slice, without the first time-slice sharing any interest with the third time-slice. One time-slice can have the ability to coordinate with a second time-slice, who has the ability to coordinate with a third time-slice, without the first time-slice having the ability to coordinate with the third time-slice.

Consider Angela, who is a chemist, and a Lutheran, and the leader of her political party. Some of her time slices are in church, and others are in the lab, and others are in the Bundestag. She may maximize her individual diachronic coherence (and thus avoid pragmatic or alethic dominance) by having credences that obey Bayesian conditionalization between these different time slices. However, if her church, her scientific community, and her political party disagree about various things, then most of these groups will be somewhat less coherent for having her time-slice as a member. She can trade off the benefits of coherence with one group for the benefits of coherence with another, by aiming for perfectly diachronically coherent credences that somehow split the difference between these groups. In these diagrams, each horizontal layer represents one group, while the vertical line down the middle represents Angela.



This first diagram puts Angela's diachronic self above all—she conditionalizes perfectly (the vertical line is perfectly straight), at some expense to how much she coheres with each group. Perhaps, as depicted in the next diagram, she should diverge from the path of perfect personal diachronic coherence, in order to approach each group at the times that it is most relevant. Nothing we have said thus far answers the difficult question of how to prioritize these notions of coherence.



There are many considerations that seem potentially relevant to the question of what we should believe in the great many cases where there are conflicting coherence requirements. Do considerations of ease matter? In the case of Angela, do the requirements of rationality take into account how much time she spends with them? How much she cares about each group? Or is there perhaps some brute fact about how much power each collective has over her that determines how trade-offs should be assessed between the various groups?

Some of this may depend on the details of the measure of degree of coherence under consideration, and also whether we are talking about pragmatic or alethic value. For example, there might there be one way of assessing trade-offs that makes sense in the alethic case, and another that makes sense in the pragmatic case. If this is right, then further questions emerge. Is there some particular way of balancing the demands that identifies what is rationally required? Or is it only possible to say what is rationally required relative to some particular value?

Many questions, no answers. Nevertheless, the mere recognition of these issues poses some interesting possibilities.

When Angela allows her different time slices to diverge from the Bayesian requirements of conditionalization, then she may better balance the demands of the groups of which she is a member. The total incoherence of the four groups of which she is a part (i.e., the church, the lab, the political party, and the diachronic life of Angela) may be less (on some particular total measure) if she only *approximates* conditionalization than if she tries to get her time slices into perfectly coherent order. This means that some of the behaviors that we typically identify paradigmatically irrational or flippant—e.g., letting your opinion be swayed by the group, or “flip-flopping”—may be justified (at least from some perspective) since one can avoid extreme *interpersonal* incoherence by sacrificing some amount of *intrapersonal* coherence.

We might also wonder *for whom* each potential credal state is better. For example, it may be best for Angela herself (considered as a diachronic entity) to have perfectly coherent beliefs across time, but better for each group of which she is a member for her to waver. Is there some grand perspective from which to assess these tradeoffs? If so, is it in Angela’s best interest to act for the greater good? By definition of being in a group, Angela shares the interests of the group, but she likely has other interests as well, both pragmatic and alethic. There may be some way to use Angela’s priorities, and the pragmatic and alethic values she shares with her past and future time slices and the rest of the various groups, to determine whether and when it’s rational to waver. But all that we feel comfortable saying is that Angela may be required to have one credal profile by one group’s standards (e.g., the political party) and another credal profile by another group’s standards (e.g., her own diachronic life). We are silent for the purposes of this paper about whether or how these standards can be integrated into one master standard.

One might think it’s clear that all of the groups of which Angela is a member should cohere with each other. But, first, it’s hard to see why this would be true since they don’t share interests with each other. Second, they generally have no ability to coordinate unless, e.g., Angela is particularly powerful in each of these groups. (Angela might be the rare person who has this power within her social organizations, but most of us lack Angela’s power.) Some who believe in a “uniqueness thesis” for credences might think there are some objective requirements of rationality that could enable coordination across groups, but we submit that just as the truth itself is an objective requirement that everyone aims at, but ultimately misses, these “evidential probabilities” will at best be groped at darkly and thereby will not enable the actual coordination of groups.

VIII. Supposition as Fleeting Belief

The general picture of rationality that we advocate may seem strange since it seems *prima facie* reasonable to criticize others when they've shown some diachronic inconsistency. For example, when someone says P in one breath, and in another says $\sim P$ (or something that implies $\sim P$), this may seem like grounds for criticism. But for everything we've said thus far, these people may be approximating coherence in a justified way—i.e., they may be sacrificing some diachronic coherence in order to achieve coherence with groups of which they are members. One might think, for example, that we erroneously legitimate the practice of a politician who utters P in order to cohere with the constituents in one neighborhood and then utters $\sim P$ in order to cohere with the constituents of another. This may seem all the more troubling since most Bayesian treatments of the subject say that there is rational pressure to cohere across time (e.g., in order to avoid diachronic Dutch Books), and we seemingly take the wind out of the Bayesian sails.

There are at least two things to say in response to this challenge.

First, just like the standard Bayesian, we can say that the politician errs *personally* insofar as they care about the joint outcome of successive actions they take on these credences, or care about personal accuracy across time. This is because we believe that Dutch Book arguments and Accuracy arguments are applicable to a collection of time-slices when they have shared interest and when it's within their power to coordinate, and the time-slices of a given politician typically share interest and are capable of coordinating.¹² Nevertheless, we do not argue that it is *all things considered* irrational to flip-flop in this way, since we countenance the existence of conflicting desires involving the joint outcome of collective actions undertaken with each separate group, or collective accuracy of these different groups.

Though we take no stand on how to generally resolve the conflict between these separate collective goals, the fact that our account does not entail that such flip-flopping is always irrational strikes us a feature rather than a bug. Consider a Miami politician who is ambivalent about the significance of global warming. At one point he might approve a plan for a new community center by the ocean that only makes sense if beachfront property in Miami is likely to be safe for several decades. But he might also campaign for congressional flood insurance legislation that is a pressing national priority only if sea levels are generally rising. There is certainly a kind of incoherence here. But since the two actions take place at different political

¹² We realize that there are some prominent exceptions to this capability claim.

levels, and with different groups of collaborators, they may be only slightly self-defeating, just as David Lewis's contradictory beliefs about roads and rails are only slightly self-defeating. If each project needs broad cooperation from many disagreeing participants, and results in many benefits that aren't entirely dependent on the future of the climate, this implicit contradiction in one person's beliefs may result in the best compromise possible.

Second, though we think that it may sometimes be rationally permissible to temporarily sacrifice diachronic coherence for group coherence, it may be incumbent upon an agent to let her interlocutors know (if asked) when or how much. By our lights, we may do this in natural language by flagging various beliefs as suppositional. For example, if Christensen and his wife temporarily adopted a credence of 50% that it would rain while shopping in order to promote their collective coherence, they would probably assent to characterizing their temporary credences as suppositional. The same is not true of beliefs or credences that are stable across time. Thus, by our lights, flip-flopping politicians may sometimes err insofar as they refuse to accept or neglect to mention that their beliefs (or credences) are suppositional. (Although we don't expect perfect diachronic consistency from any politician, we do often expect their commitments made before one audience to guide at least *some* of their future governmental behavior.)

One might think that we err insofar as we treat supposition as a species of belief, rather than as a kind of mental state that is distinct from belief, complete with its own cognitive function. But we embrace this consequence of our view. When epistemology gives up absolutely general norms of diachronic coherence—e.g., when it is agreed that one can permissibly abstain from conditionalizing—there is less reason to treat supposition as distinct from belief. This is because there is less reason to posit a distinct mental state (complete with its own distinct machinery) that underwrites our practice of temporally accepting certain propositions as true (or as probably true) when beliefs (or credences) are allowed to fleet. We can simplify our view of what goes on in the head by treating our suppositional temporary acceptances as beliefs that conflict with what one accepts most of the time, and whose temporary nature is justified for one reason or another. Many philosophers follow Descartes in sharply distinguishing the acceptance or belief in a proposition from its mere entertainment. But some psychologists (Gilbert et al., 1990) follow Spinoza in suggesting that even supposition partakes in *some* of the force of belief. Our take is that stable belief, fleeting belief, and mere supposition are just members of a single continuum.

IX. Conclusion

We have argued that there is rational pressure for a collection of time-slices to cohere whenever (i) the collection shares interest in some value that they stand to maximize by cohering, and (ii)

the time-slices are causally connected in a way that allows them to coordinate. When a collection of time-slices has *some* ability to coordinate, but not enough to guarantee perfect coherence, and only has *some* shared interests, we say that the collective has interest in coherence to the degree of shared interest and ability to coordinate. When a time-slice enters into multiple collectives and cannot jointly satisfy the coherence requirements that arise from the time-slice's membership in these groups, we are not sure what rationality requires of the time-slice, but we do think that there may be times where it is appropriate to sacrifice some coherence with one collective for greater coherence with another. When an agent temporarily sacrifices some diachronic coherence for greater coherence with a group, we think that the agent's temporary beliefs or credences can be fruitfully regarded as suppositional.

References

- Armstrong, K. (2017). Two Murder Convictions for One Fatal Shot. *New Yorker*, <https://www.newyorker.com/magazine/2017/11/13/two-murder-convictions-for-one-fatal-shot>.
- Arntzenius, F. (2003). Some Problems for Conditionalization and Reflection. *Journal of Philosophy*, 100, 356-370.
- Berker, S. (2013a). Epistemic teleology and the separateness of propositions. *Philosophical Review*, 122, 337-393.
- Berker, S. (2013b). The rejection of epistemic consequentialism. *Philosophical Issues*, 23, 363-387.
- Bratman, M. (1987). *Intentions, Plans, and Practical Reason*. Stanford, CA: CSLI Publications.
- Briggs, R.A. and R. Pettigrew (2016). An Accuracy-Dominance Argument for Conditionalization. Unpublished ms.
- Christensen, D. (1991). Clever Bookies and Coherent Beliefs. *Philosophical Review*, 100, 229-247.
- De Bona, G., and Staffel, J. (2017). Graded Incoherence for Accuracy-Firsters. *Philosophy of Science*, 84, 189-213.
- Easwaran, K. (2013). Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection. *Philosophy of Science*, 80, 119-142.

Easwaran, K. (2017). The tripartite role of belief: evidence, truth, and action. *Res Philosophica*, 94, 189-206.

Elga, A. (2000). Self-Locating Belief and the Sleeping Beauty Problem. *Analysis*, 60, 143-147.

Gilbert, D., Krull, D., and Malone, P. (1990). Unbelieving the Unbelievable: Some problems in the rejection of false information. *Journal of Personality and Social Psychology*, 59, 601-613.

Gillies, D. (1991). Intersubjective Probability and Confirmation Theory. *British Journal for the Philosophy of Science*, 42, 513-533.

Greaves, H., and Wallace, D. (2006) Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115, 607-632.

Holton, R. (2009). *Willing, Wanting, Waiting*. Oxford: Oxford University Press.

Horowitz, S. and S. Dogramaci (2016). Uniqueness: a New Argument. *Philosophical Issues*, 26.

Jaffe, Mark (2000). *The Gilded Dinosaur: The Fossil War Between E. D. Cope and O. C. Marsh and the Rise of American Science*. New York: Crown Publishing Group.

James, W. (1896). The Will to Believe. *The New World*, 5, 327-347.

Joyce, J. (1998). A Non-Pragmatic Vindication of Probabilism. *Philosophy of Science*, 65, 575-603.

Kopec, M. (2017). A New Group Dutch Book Argument. *Ratio*, 30, 122-136.

Kopec, M. and M. Titelbaum (2016). The Uniqueness Thesis. *Philosophy Compass*.

Korsgaard, C. (2009). *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.

Lewis, D. (1982). Logic for Equivocators. *Nous*, 16, 431-441.

Moss, S. (2011). Scoring rules and epistemic compromise. *Mind*, 120(480):1053–1069.

Paul, L.A. (2014). *Transformative Experience*. Oxford University Press.

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

Rowbottom, D. (2013). Group Level Interpretations of Probability: New Directions. *Pacific Philosophical Quarterly*, 94, 188-203.

Schervish, M. J., T. Seidenfeld, and J. B. Kadane (2002). Measuring incoherence. *Sankhya: The Indian Journal of Statistics, Series A*, 561-587.

Schoenfield, M. (2014). Permission to Believe: Why Permissivism is True and What It Tells Us about Irrelevant Influences on Belief. *Nous*, 48, 193-218.

Shor, Elizabeth (1974). *The Fossil Feud Between E. D. Cope and O. C. Marsh*. Detroit, Michigan: Exposition Press.

Staffel, J. (2015). Measuring the Overall Incoherence of Credence Functions. *Synthese*, 192, 1467-1493.

Talbott, W. (1991). Two Principles of Bayesian Epistemology. *Philosophical Studies*, 62, 135-150.

Velleman, D. (2000). *The Possibility of Practical Reason*. Oxford University Press.