HOW DO BELIEFS SIMPLIFY REASONING?

Julia Staffel

University of Colorado at Boulder

(forthcoming in *Noûs*, penultimate version)

Abstract:

According to an increasingly popular epistemological view, people need outright beliefs in addition to credences to simplify their reasoning. Outright beliefs simplify reasoning by allowing thinkers to ignore small error probabilities. What is outright believed can change between contexts. It has been claimed that thinkers manage shifts in their outright beliefs and credences across contexts by an updating procedure resembling conditionalization, which I call pseudo-conditionalization (PC). But conditionalization is notoriously complicated. The claim that thinkers manage their beliefs via PC is thus in tension with the view that the function of beliefs is to simplify our reasoning. I propose to resolve this puzzle by rejecting the view that thinkers employ PC. Based on this solution, I furthermore argue for a descriptive and a normative claim. The descriptive claim is that the available strategies for managing beliefs and credences across contexts that are compatible with the simplifying function of outright beliefs can generate synchronic and diachronic incoherence in a thinker's attitudes. Moreover, I argue that the view of outright belief as a simplifying heuristic is incompatible with the view that there are ideal norms of coherence or consistency governing outright beliefs that are too complicated for human thinkers to comply with.

Introduction

I present a new puzzle for the role of outright beliefs in reasoning. According to an increasingly popular view in epistemology, human thinkers have two kinds of belief-like attitudes: credences, and outright beliefs. Credences are fine-grained belief states that encode varying degrees of uncertainty. Outright beliefs (which are often also just called 'beliefs'), by contrast, are coarse-grained attitudes that don't make reference to how probable or improbable things are. It is commonly assumed that people can outright believe claims that they are not completely certain of, and that this is rationally permissible.

Why do we need outright beliefs in addition to credences? A popular explanation is that outright beliefs simplify our reasoning processes by allowing us to disregard small probabilities of error. For most ordinary empirical claims, even highly probable ones, we're not completely certain that they are true. Forming an outright belief in such a claim lets us reason as if we had full confidence in it. This is useful in many contexts, because we

can thereby dramatically narrow the number of possibilities we need to consider in reasoning and decision-making.

I argue that this claim about the simplifying function of outright beliefs gives rise to a puzzle when combined with other plausible and commonly endorsed claims about outright beliefs. Defenders of such a dual view of belief usually assume that whether we rely on a high credence in some claim p or an outright belief in p in our reasoning can change from context to context. Moreover, in any given context, we can rely on mixtures of credences and outright beliefs, treating some claims as true and others as uncertain. This gives rise to the question of how thinkers manage their credences and outright beliefs across contexts, i.e. how they arrive at a given set of credences and outright beliefs to rely upon in a given context of reasoning. It has been proposed that thinkers do so by employing a strategy I call pseudo-conditionalization (PC). PC appears to be an attractive strategy for managing outright beliefs and credences across contexts, because it ensures that, in any given context of reasoning, the set of attitudes the thinker reasons with is coherent. The problem with PC is that it is psychologically unrealistic that human reasoners use it, since it is very computationally demanding. Hence we arrive at the following puzzle: How could it be true that the purpose of having outright beliefs is to simplify reasoning, when managing one's outright beliefs requires such a complex strategy?

I propose to solve the puzzle by rejecting the view that thinkers manage their beliefs and credences by employing PC. Based on this solution, I will furthermore argue for a descriptive and a normative claim. The descriptive claim is that the available strategies for managing beliefs and credences across contexts that are compatible with the simplifying function of outright beliefs can give rise to synchronic and diachronic incoherence in a thinker's attitudes. By revealing possible tradeoffs between simplicity and coherence in reasoning, we gain a better understanding of why limited human reasoners fail to have ideally coherent doxastic states. Moreover, I argue that the view of outright belief as a simplifying heuristic is incompatible with the view that there are ideal norms of coherence or consistency for outright belief that outstrip human thinkers' reasoning capacities. If the main purpose of outright beliefs is to simplify reasoning, then beliefs can't be governed by norms that are so difficult to comply with that doing so prevents the beliefs from serving this function.

In section one of the paper, I present the ingredients of the puzzle, and explain in what sense they are puzzling when taken together. In section two of the paper, I examine which of the claims comprising the puzzle ought to be rejected. In section three of the paper, I explain how we can think of the simplifying role of outright beliefs in a way that avoids the puzzle, and I draw lessons from the proposed solution to the puzzle.

1. The Puzzle

In the current literature on belief, it is increasingly popular to assume that human thinkers have two kinds of belief-like attitudes: credences, and outright beliefs (which are

often just called 'beliefs').[1] Credences are fine-grained belief states that encode varying degrees of uncertainty.[2] We ascribe credences to people when we say things like "Jane is highly confident that she will be promoted" or "James thinks it's unlikely that his flight is on time." Outright beliefs, by contrast, are coarse-grained attitudes that don't make reference to how probable or improbable things are.[3] We ascribe outright beliefs to people by saying things like "Jane believes the store is still open", or "James thinks that Paris is the capital of France." One of the most important differences between outright beliefs and credences is how they behave in reasoning. If someone relies on an outright belief in p in reasoning, the person takes p for granted, or treats p as true. The possibility that ~p is ruled out. By contrast, if someone reasons with a high credence in p, they don't take p for granted. The possibility that p might be false is not ruled out. To see the difference, consider someone who is selecting what to bring to a vegan brunch. Her reasoning and decision-making will be substantially different depending on whether she reasons with an outright belief that the cookies contain no animal products, or whether she reasons with a high credence that the cookies don't contain animal products.[4]

If we accept that people in fact possess those two types of belief-like attitudes, we may then wonder: How are beliefs and credences in fact related in human thinkers? Which combinations of credences and beliefs are rational for people to have? The answers to these questions can of course come apart, since the first question is a descriptive question, and the second one a normative question. Most normative theories of belief deem an outright belief in some claim p to be permissible when a thinker is highly confident in p, but not necessarily certain of it (with some exceptions such as lottery propositions). The combinations of beliefs and credences we actually find in people seem to be roughly in line with these normative constraints. People seem to readily form beliefs in propositions that they are not completely certain of, and rely on them in reasoning and communication, but refrain from forming beliefs when their confidence is too low (Weisberg forthcoming).

---

[1] E.g. Buchak 2014, Clarke 2013, Easwaran & Fitelson 2015, Greco 2015, Leitgeb 2016, Lin 2013, Lin & Kelly 2012, Ross & Schroeder 2014, Staffel 2016, Sturgeon 2015, Tang 2015, Weatherson 2016, Wedgwood 2012, Weisberg 2013, Weisberg forthcoming. For earlier discussions of this type of view, see e.g. Harsanyi 1985, Levi 1964.

[2] Some philosophers think that credences should be construed as complex, graded attitudes held towards simple contents. Others think that they are best thought of a simple attitudes held toward complex contents, where the contents are sets of probability spaces (Carr forthcoming, Moss 2018). For our purposes in this paper, it doesn't matter which position we adopt.

[3] The terminology varies a bit in the literature: some people call any ungraded belief an outright belief, regardless of whether its content is probabilistic. On this view, both a belief that it's hot outside and a belief that it's probably hot outside would count as outright beliefs. I use the term 'outright belief' in a more limited sense here, which excludes ungraded beliefs towards probabilistic contents. Only the former example is an outright belief in my sense. When I use the term 'belief' in this paper, it is interchangeable with this understanding of 'outright belief.'

[4]This way of distinguishing between credences and beliefs according to their roles in reasoning might not be attractive to adherents of some types of interpretivist views of propositional attitudes. I leave it as a question for future research whether a version of the puzzle I present here arises for such interpretivist views.

This way of thinking about outright beliefs and credences gives rise to a natural way of incorporating outright beliefs into formal models of doxastic attitudes. In slogan format, the idea is that "belief is credence 1 in context." Clarke (2013) and Greco (2015, 2017) argue that we should identify an outright belief with a credence of 1 in a proposition, but allow that whether a proposition is assigned credence 1 can vary between contexts (see also Harsanyi 1985). Similarly, Wedgwood (2012) argues that reasoners like us have theoretical credences (the credences that we adopt purely in light of our evidence), and practical credences. The latter are simplified versions of our theoretical credences. An outright belief is identified with a practical credence of 1. Tang (2015) also defends the view that outright beliefs can be identified with high probability estimates that get rounded up to 1. Hence, on this view, a credence of one in a claim p no longer exclusively represents a thinker's (stable) certainty that p is true, it can also represent the attitude of treating p as true in a context of reasoning. I will adopt this modeling convention throughout the paper, but the substance of my argument does not depend on it.

While the view that human thinkers have both outright beliefs and credences has struck many philosophers as extremely plausible, it also raises an important question, which has sometimes been called the 'Bayesian challenge.' Having both outright beliefs and credences creates a kind of redundancy in our attitudes that might seem unnecessary. Why have outright beliefs at all, when it seems like credences by themselves can do the job? This explanatory challenge is thought to have been first expressed by Richard Jeffrey (1970), who said:

> I am inclined to think that Ramsey sucked the marrow out of the ordinary notion [of belief] and used it to nourish a more adequate view. But maybe there is more there, of value. I hope so. Show me; I have not seen it at all clearly, but it may be there for all that.

There is a small literature on answers to the Bayesian challenge. A highly popular answer to the challenge is that human thinkers need outright beliefs in addition to credences in their inventory of doxastic attitudes, because outright beliefs help simplify reasoning (see e.g. Harsanyi 1985, Lance 1995, Leitgeb 2016, Lin 2013, Lin & Kelly 2012, Ross & Schroeder 2014, Tang 2015). Human reasoners' cognitive limitations make it infeasible for them to use only credences in their reasoning, because it requires keeping track of many different possibilities, even if some of those possibilities are very improbable and could safely be ignored. This is where outright beliefs help. By letting reasoners treat highly probable claims as true, the number of possibilities that need to be considered is greatly reduced, thereby making reasoning problems more tractable.[5] We can illustrate

---

[5] One might wonder why I have classified the "treating as true" attitude as being an outright belief, rather than the attitude of acceptance. Acceptances also allow reasoners to treat claims as true, and can thus help

this with a simple example. Suppose you are wondering how likely it is to rain during an upcoming tennis match. The problem is that you don't remember where the tennis match will take place. You think it might be in New York or Boston or LA. Your credences are as follows:

Cr (NY) = 0.48
Cr (Boston) = 0.48
Cr (LA) = 0.04

Of course, how likely it is to rain during the match depends on where it will take place. You have the following conditional credences reflecting this:

Cr (rain|NY) = 0.7
Cr (rain|Boston) = 0.9
Cr (rain|LA) = 0.1

In order to correctly compute Cr(rain), you need to plug your conditional credences of the form (rain|place of match) and your unconditional credences about where the match happens into the total probability theorem:

Cr (rain) = Cr(rain|NY) Cr(NY) + Cr(rain|Boston) Cr(Boston) + Cr(rain|LA) Cr(LA)
Cr (rain) = 0.772

This computation could be simplified if you disregarded the possibility that the match might be in LA, which you consider to be very improbable. If you reasoned with an outright belief that the match is in Boston or New York, which you consider to be equally likely, you could simply take the average between the rain probabilities in Boston and New York, which comes out to Cr(rain) = 0.8. Of course this is not quite correct, but it's very close to the correct answer. Formally, we might represent your reasoning as follows then, representing outright belief as credence 1 in context:

---

simplify reasoning processes. In the philosophical literature, acceptances and outright beliefs are commonly distinguished because they differ in how they are formed. Acceptance is usually taken to be under a thinker's voluntary control. I can decide to treat any claim as true in reasoning, no matter how low my confidence in the claim is, while being fully aware that the claim might be false. By contrast, outright beliefs cannot be voluntarily adopted in the same way (although some philosophers think we have some amount of control over what we believe). Rather, whether or not we outright believe a claim is usually regulated by cognitive processes that are automatic. We seem to have some deliberative control over switching from relying on an outright belief to relying on a credence, for example by directing our attention to ways in which we might be mistaken. But we usually can't employ deliberative control over which claims we take for granted in framing a reasoning problem, this is done automatically and without our conscious awareness. The Bayesian challenge, as I understand it, is the question of why our minds are equipped to employ outright beliefs, so construed, and this paper discusses one possible answer and its implications.

Cr (NY) = 0.5
Cr (Boston) = 0.5
~~Cr (LA) = 0.04~~

Cr (rain) = Cr(rain|NY) Cr(NY) + Cr(rain|Boston) Cr(Boston)
Cr (rain) = (0.7 + 0.9) 0.5 = 0.8

This example is a very simple illustration of the general idea that reducing the number of possibilities under consideration simplifies reasoning problems, both of the theoretical and the practical kind. This observation helps explain why it makes sense for limited human reasoners to have outright beliefs in addition to credences: outright beliefs let us eliminate improbable options from consideration in framing reasoning problems, thus making them easier to solve. We will capture this response to the Bayesian challenge in claim (1):

(1) Human thinkers have outright beliefs in addition to credences because outright beliefs help simplify reasoning processes.

In order to show that (1) leads to a puzzle, I need to introduce three further claims. I will state them first, then explain why each of them is plausible, and then argue that they give rise to a puzzle when combined with (1).

(2) Reasoning processes can involve mixtures of credences and outright beliefs, and it is flexible from context to context which outright beliefs are relied upon.

(3) Outright beliefs and credences in a context are determined by pseudo-conditionalizing on a set of background credences.

(4) Pseudo-conditionalizing is difficult to execute for human reasoners, because it is computationally expensive.

Claim (2) consists of two sub-claims, which we might call *Mixing* and *Switching*. According to *Mixing*, any given reasoning task can involve both outright beliefs and credences. The tennis match example illustrates this claim: you rely on an outright belief in the disjunction that the match is in New York or Boston, but you also rely on various conditional and unconditional credences when determining how confident you should be that it will rain during the match. Many, perhaps even most, practical and theoretical reasoning problems have this structure, in which some contingent claims are treated as true, and some claims are treated as uncertain in generating an answer to the problem.

According to *Switching*, whether or not we rely on a credence or on an outright belief in a relevant claim can change from context to context. (In a given context, we

obviously have to rely on one or the other, not both). Such a change can be triggered, for example, by a change in the stakes of a situation. Suppose you are talking to your friend about when you will be able to meet her. You believe that you can take the bus to arrive at her house at 6 pm, and you tell this to your friend. Your friend might then make you aware of the fact that it would be very bad if you came later than 6 pm. Realizing that an error would be costly, you might no longer want to rely on your outright belief that the bus arrives at 6 pm, and instead take into consideration the possibility that the bus might be late and not get there by 6 pm. One plausible way to think about what happened here is that you switched from relying on an outright belief to relying on a high credence in the claim that the bus arrives at your friend's house at 6 pm. Again, it seems common to switch between treating a claim as true and treating it as merely highly likely when the context changes, for example when a new aspect of the situation becomes salient, an error possibility comes into focus, or the stakes change. This type of phenomenon is well documented in the philosophical literature (see, e.g. Hawthorne 2004, Greco 2015, 2017 for discussion, and especially Nagel 2011 for an empirically informed account of switching).

The phenomena of *Mixing* and *Switching* lead to a question, to which claim (3) of the puzzle offers an answer. The question is: How do thinkers (or better, their cognitive systems that execute this task) manage the shifts between different contexts of reasoning, and determine which credences and outright beliefs to rely upon in a given context? Managing these attitude shifts requires solving two problems. The first problem that needs solving is to figure out which claims to treat as true and which claims to treat as uncertain in a given context. In other words, the thinker (or her cognitive system) must figure out which probabilities of error can safely be ignored in a given context of reasoning. Obviously, a thinker can't run an expected utility calculation to decide which error possibilities can safely be ignored, since this would defy the purpose of simplifying the reasoning problem. Hence, a different mechanism must be at work. This problem has received some attention in the philosophy literature, see for example Lin (2014).[6]

But even if the problem of choosing which outright beliefs to rely upon in a given context has been solved, there's a second problem, which will be my focus in what follows. The problem is whether and how to readjust one's remaining doxastic attitudes once it is settled which possibilities are considered live in a given context. This question has not received much attention in the literature. To my knowledge, there is only one proposal for how thinkers might handle these adjustments, which was made by Harsanyi (1985), and more recently (and independently of Harsanyi) by Clarke (2013).[7] They

---

[6] As an anonymous referee helpfully points out, this problem is closely related to the so-called *frame problem* in AI. A good survey is provided by Shanahan (2016).

[7] I also talk about this problem in my discussion of whether Pettigrew's accuracy-first program is incompatible with a view that takes thinkers to have both credences and outright beliefs. In his reply, Pettigrew considers some alternatives to pseudo-conditionalization, but finds them problematic (Staffel 2017, Pettigrew 2017).

propose that thinkers engage in a strategy I will call 'pseudo-conditionalization' for settling on a particular set of attitudes in a given context. The idea is that thinkers have a global credence function that consists of conditional credences over a set of possibilities. In different contexts, thinkers treat different possibilities as live. By pseudo-conditionalizing on the possibilities she considers live in a given context, it is determined what a thinker's attitudes are in a context. If something is not considered a live possibility in a context, we treat it as if the thinker assigns it credence 0, and its negation is assigned credence 1, and thus believed in that context. The thinker's remaining credences are adjusted as if the thinker were conditionalizing on the live options. In shifting between contexts, the thinker can expand or contract the set of possibilities she rules out, and hence change the set of beliefs and credences she relies on.

When I introduced the tennis match example, I already introduced the pseudo-conditionalization strategy, without calling it that. We assumed that the thinker's full possibility space included three options for where the tennis match might take place, and that she disregarded the LA option because it was so improbable. I modeled this as the thinker's updating her credence in (Boston or NY) from 0.96 to 1, and readjusting her credences accordingly, from 0.48 to 0.5 in each disjunct. Assuming that the thinker's conditional probabilities remain fixed across contexts, it is now easier for her to answer the question of whether it will rain during the match.

It is easy to see why I choose to call this strategy "pseudo-conditionalization." Ordinary conditionalization is the standard Bayesian strategy for rationally updating one's credences on newly learned evidence. The rule says that if you learn that some claim q is true with certainty (i.e. Cr (q) becomes 1), then your new credence in any claim p should be your old credence in p conditional on q, i.e. $Cr_{new}(p) = Cr_{old}(p|q)$. Pseudo-conditionalizing works in the same way, except that the thinker doesn't set her credence in a claim q to 1 because she has learned that it is true, but instead because she is ruling out the possibility of ~q for the purposes of reasoning in that context.[8]

---

[8] The way I have characterized PC might create the impression that proponents of PC are committed to a particular view about how credences and beliefs are stored and/or generated in the mind. The view that seems to naturally accompany PC is one on which some set of background credences are the most fundamental attitudes that are stored by the thinker, on the basis of which outright beliefs are generated when they are needed in a given context. Yet, this view does not seem particularly psychologically realistic. Fortunately, on closer inspection, proponents of PC are not actually committed to it. Proponents of PC are merely committed to the idea that thinkers' doxastic attitudes are *representable* as being generated by PC. This is compatible with different theories of how credences and outright beliefs are stored and generated in our minds. One view that seems relatively uncontroversial in the empirical and philosophical literature is that some of our doxastic attitudes are stored explicitly, whereas others are stored implicitly. Explicitly stored attitudes are immediately and easily available for cognitive processes, whereas implicitly stored attitudes have to be inferred from other attitudes in order to make them available for cognitive processes (see e.g. Harman 1986, Kirsh 2003). Yet, this still leaves open a variety of views about which attitudes are explicitly stored. Do we explicitly store just some of our credences, or some of our beliefs, or some of both? Weisberg (forthcoming) has recently argued for yet another option, namely that we store information, and depending on the situation at hand, this information can be called up as an outright belief or as a credence. Weisberg argues that research on memory beliefs supports this view as the most psychologically realistic

The pseudo-conditionalization strategy seems like an attractive solution to the problem of how thinkers manage their credences and outright beliefs across contexts. It inherits its attractiveness from the appeal of regular conditionalization. Regular conditionalization ensures that a thinker's updated credences are appropriately related to her prior credences, and that the resulting credences are always probabilistic. Both of these points are important in explaining the appeal of pseudo-conditionalization. A thinker who reduces the complexity of a reasoning task by ruling some improbable options out from consideration is presumably still interested in otherwise maintaining the relative proportions of her credences, which reflect both what evidence she has already gathered, and what her take on the impact of this evidence is. Moreover, there are many arguments for the claim that synchronically incoherent attitudes are rationally defective; hence, a simplification strategy that renders the resulting set of attitudes probabilistically coherent is desirable. PC perfectly accomplishes both of these tasks, so it seems like the obviously best strategy for managing beliefs and credences across contexts.

This brings us to the last claim in the puzzle. According to (4), pseudo-conditionalizing is difficult to execute for human reasoners, because it is computationally expensive. The justification for (4) is based on results from research in psychology, computer science, and complexity theory, all of which have investigated how demanding it is to update probability functions by conditionalizing. While psychologists have not studied PC in particular, they have studied whether reasoners' updates on newly learned evidence are representable as updates by conditionalization. Since PC is the same rule as conditionalization, except that the shifts to credence 1 are not evidence-based, psychological research about conditionalization can help us evaluate the feasibility of PC for human thinkers. Relevant psychological evidence includes the following findings: In reasoning tasks that ask subjects to estimate the posterior probability of a hypothesis on the basis of relevant evidence, subjects' estimates tend to be too conservative, i.e. they underestimate how much the evidence impacts the probability of the hypothesis (Edwards 1982, Slovic & Liechtenstein 1971). There have also been studies comparing people's ability to judge the degree to which a piece of evidence confirms a hypothesis and people's ability to judge the posterior probability of a hypothesis given a piece of evidence. (These can come apart: a piece of evidence can have high confirmatory force, yet not lead to a very high posterior probability.) People turn out to be better at judging confirmation than posterior probabilities (e.g. Mastropasqua et al. 2010, Tentori et al. 2016). In investigating people's conditional probability judgments, it has been found that subjects don't strictly conform to the rigidity requirement, which prescribes that

---

theory of how we store and generate outright beliefs and credences. I discuss this view further in section three.

What matters for our purposes is that any of these positions is in principle compatible with the claim that the evolution of a thinker's credences and beliefs across contexts is representable by pseudo-conditionalization. Different views will differ in their explanations of what the underlying cognitive processes are that give rise to the thinker's attitudes being so representable, but differences in these explanations don't affect the argument against PC I make in what follows.

conditional credences should be stable (Zhao & Osherson 2010). Moreover, conditional probability judgments differ (contrary to what would be rational according to Bayesian models) depending on whether subjects learn that some claim p is true, or are asked to suppose that p is true (Zhao et al. 2012). People's conditional probability judgments also don't align perfectly with their unconditional probability judgment as required by the standard ratio formula (Zhao et al. 2009; see also Evans et al. 2015 for further helpful discussion and references). Taken together, the findings in the psychology literature suggest that the way in which people update on newly learned evidence often approximates conditionalization, but mostly doesn't match it perfectly. It is an ongoing research program to identify the reasoning strategies that give rise to the observed updating patterns, but most researchers share the working hypothesis that the underlying reasoning strategies are not equivalent to the conditionalization rule; rather, it seems plausible that reasoners use simplified strategies that approximate conditionalization closely enough.

The hypothesis that reasoners use simpler updating strategies than standard conditionalization gains support from results in computer science and complexity theory. If the conditionalization rule were relatively simple to implement from a computational point of view, then it would be puzzling why human reasoners don't conform to it. Yet, it turns out that conditionalization can be very complex to execute, which explains why limited reasoners may rely on simpler rules instead. The most general and abstract result that is relevant for our purposes is that probabilistic inference, including conditionalization, is NP-hard (Cooper 1990). This means that in general, the reasoning task of determining the probability of some claim p can be solved by a non-deterministic Turing machine in polynominal time. In other words, the maximum amount of time it takes a non-deterministic Turing machine to solve the problem is determined by a polynominal function of the size of the input. A thinker's credences (or a Bayesian network, which is Cooper's framework for representing probabilistic information) may contain all the necessary information to determine Cr (p), but many computational steps might be required to actually arrive at the result.[9] This is of course not to say that every single probabilistic reasoning task is this complex. This result about NP-hardness is a result concerning the complexity of probabilistic reasoning (including conditionalization) in general, not a result about particular instances of reasoning. Whether a particular reasoning problem is difficult depends on the information that is explicitly available to the thinker, and the kinds of computations required for arriving at the relevant answer. Consider the tennis example again, in which the thinker wants to answer the question of

---

[9] According to Cooper, the hardest probabilistic inference problems are ones that are representable by multiply connected belief networks that contain many uninstantiated variables. This means that the probability of any given claim in the network can depend on the probabilities of more than one other claim, and that many claims in the network have non-extreme probability values. Probabilistic reasoning task are easier when they are representable by singly connected networks, or when the truth-values of many of the claims in the network are known. The latter finding is in agreement with the thesis that outright beliefs simplify reasoning, since they let the reasoner assume that their truth values are known.

how likely it is to rain during the match. If the thinker already knew, for example, how likely it was *not* to rain during the match, the inference to the probability of rain would be trivial. Similarly for pseudo-conditionalization: If the thinker already had all of the conditional probabilities needed for pseudo-conditionalizing in any context readily available, then pseudo-conditionalizing would not be very difficult, because PC would basically require a trivial one-step inference.[10] But this assumption is not psychologically realistic. Humans have some conditional and unconditional credences readily available to them, whereas others can only be generated by inferential processes (Kirsh 2003). If thinkers had the conditional credences needed for pseudo-conditionalizing readily available to them, then they should also have them readily available for use in regular conditionalization. Yet, as we know from the psychological research I briefly summarized above, human thinkers at best approximate correct Bayesian inferences, including conditionalization in their reasoning. It is a matter of active debate how we can characterize the algorithm human thinkers use when reasoning with and updating their credences, but there seems to be little doubt that it is some kind of heuristic and not a full Bayesian algorithm. Hence, if being a perfect conditionalizer is infeasible for a human thinker, being a perfect pseudo-conditionalizer is, too.

We now have the resources to see why claims (1)-(4) together generate a puzzle. While they are not inconsistent, it is hard to see how they could be true at the same time. Claim (1) asserts that the reason why human thinkers are equipped with outright beliefs in addition to credences is that outright beliefs simplify reasoning tasks, and thereby make them more tractable for limited beings like us. But if human thinkers have to use pseudo-conditionalization for managing their credences and outright beliefs across contexts, then the computational cost of using this strategy is likely to significantly diminish the computational benefits from having outright beliefs in the first place. Hence, it is very implausible that the claims in the puzzle are jointly true. In the next section, I will examine which of the four claims should be rejected.

2. Solving the Puzzle

The first claim in the puzzle states that people have outright beliefs in addition to credences in their inventory of doxastic attitudes, because they help rein in the complexity of our reasoning. There are different strategies for denying this claim. The first option is of course to deny its presupposition, i.e. that people have outright beliefs at all. On such a

---

[10] A similar picture is suggested for consideration by an anonymous referee. If humans had all of their conditional credences explicitly stored, as well as all the claims to which they assign credence 1 or 0, then any unconditional credence could be arrived at via a trivial one-step inference. On this view, computational effort is reduced by drastically increasing the amount of information that needs to be stored by the thinker. The thinker needs to have a conditional credence "ready to go" for every possible evidential situation she might find herself in. Human thinkers don't have this kind of extensive store of conditional credences, so this way of making PC feasible doesn't work for us. This problem was already recognized by Harman, who argues that humans can't update their credences via conditionalization (1986, p. 25). However, Harman mistakenly assumes that this is the only way in which reasoning with degrees of belief could be implemented by humans, and on this basis, rejects the possibility that humans can reason with credences altogether.

view, which has been defended by Jeffrey (1970) and Pettigrew (2016) among others, there is no puzzle to begin with, because the question of what beliefs are for doesn't arise. This solution is of course unattractive for the many defenders of the view that we have outright beliefs. This paper is not the place to settle the debate about whether people really have outright beliefs. I will proceed on the assumption that human thinkers do in fact have outright beliefs that can play a role in their reasoning and decision-making, and note that this view is not universally endorsed. One might also consider challenging the idea that cutting down the space of possibilities simplifies reasoning tasks. Yet, since this result has a solid footing in complexity theory (Cooper 1990), this response is implausible. This leaves us with the option of denying claim (1) by arguing that human thinkers have outright beliefs for a different reason. On such a view, their main function is something other than simplifying reasoning.

Some alternative explanations of the function of outright beliefs have recently been proposed in the literature. One view is that outright beliefs are necessary as a basis for moral judgment. Buchak (2014) argues that merely being very confident that, say, Hans didn't pay for his concert ticket, is not sufficient for judging that Hans did something wrong. This is because high confidence can sometimes be justified based on purely statistical evidence (for example evidence that most concert attendees used forged tickets), but purely statistical evidence is an intuitively (and legally) insufficient basis for a judgment of wrongdoing. Buchak argues that outright beliefs are sensitive to different evidence types, and that one can only rationally form an outright belief in a claim if one possesses non-statistical evidence in its support. Hence, a rational outright belief in a claim p can form the basis of a moral judgment, since it ensures that the thinker possesses non-statistical evidence in support of p. By contrast, a high credence in p may be based on purely statistical evidence, and thus cannot always be an appropriate basis for a moral judgment. While I am sympathetic to Buchak's argument that purely statistical evidence is insufficient as a basis for moral judgment, her position is less convincing as a response to the Bayesian challenge. The worry, in brief, is this: Suppose high confidence based on non-statistical evidence can rationalize outright belief, and outright belief is the appropriate basis of moral judgment. Why, on this view, is outright belief needed as a middleman? Why can't high confidence based on non-statistical evidence directly support moral judgments? This problem with Buchak's view is especially salient when we consider her analogy with legal cases. As Moss (2018) points out, in civil cases, the "preponderance of evidence" standard applies, which means that in order to win a case, the relevant party in the trial must show that the evidence makes it more than 50% likely that they are right. The evidence in question that gives rise to this probability must not be purely statistical. In a case that is decided according to this standard of proof, the resulting probability might not be high enough to warrant belief that the winning party is right. Hence, what matters for whether a legal judgment can be reached in this kind of case is whether the relevant claim can be established to be more than 50% likely based on the right kind of evidence, not whether one can form an outright justified belief in the relevant claim.

Hence, while Buchak argues compellingly that both moral judgments and legal judgments should not be based on purely statistical evidence, it does not follow that the main function of outright beliefs is to ground moral (or legal) judgments.

Another strategy for explaining the function of outright beliefs appeals to the importance of knowledge. Outright belief is usually seen as a necessary condition for knowledge, hence, the importance of having outright beliefs could be justified by appealing to the value of having knowledge. A closer look at how I've defined credences and outright beliefs for the purposes of this paper shows that this response is unpersuasive. I've defined credences as any doxastic attitudes that encode uncertainty, and outright beliefs as doxastic attitudes that don't encode uncertainty. So defined, credences can constitute knowledge on any view of knowledge. It is controversial whether credences, understood as graded attitudes towards simple contents, can constitute knowledge (see Moss 2018, for an argument that they can). Yet, this is not my definition of a credence. On my view, a belief that p is likely falls into the category of credences, and no one disagrees that this type of attitude is a candidate for knowledge (Moss 2018, Hawthorne & Stanley 2008, Hawthorne 2004, Williamson 2000). On this view, the Bayesian challenge can simply be reformulated by asking why a thinker whose credences constitute knowledge needs outright beliefs (i.e. beliefs not encoding uncertainty).

While there may of course be other responses to the Bayesian challenge that I haven't considered here, I will for now conclude that we shouldn't resolve our puzzle by rejecting claim (1). The remainder of the paper thus assumes that the reason why human thinkers have outright beliefs is that they simplify reasoning.

The second claim in the puzzle states that people can reason with mixes of outright beliefs and credences, and that which outright beliefs we rely on can shift between contexts. There are different ways in which we could deny this claim. One way to deny it is to argue that it is fixed or largely fixed across contexts which outright beliefs we reason with. Another way to deny it is to argue that people always reason with either outright beliefs or credences, but never with mixes of them.

The 'no mixing' claim seems extremely implausible from the point of view of a defender of a dual account of the nature of belief. There are plenty of examples of reasoning processes that are most naturally described as involving such mixed attitudes. When we reason with claims we deem uncertain, we usually do so in light of background assumptions we simply treat as true. For example, take the following line of reasoning: "Janet will cook Indian food or Chinese food for dinner. But she probably won't make Indian food today, because we had that last time. Thus, she'll probably make Chinese." This argument is naturally described as having an outright belief in a disjunction as its first premise, and credences as its second premise and conclusion. It is a perfectly ordinary instance of how people reason. Hence, the "no mixing" claim is implausible. Interestingly, there are some formal frameworks that seem to be designed to validate something like 'no mixing.' For example, Leitgeb's stability theory of belief is built in such a way that a coherent thinker who uses standard conditionalization to reason with her

credences, and AGM belief revision theory to reason with her outright beliefs, will always be guaranteed to have rationally permissible combinations of outright beliefs and credences. But the framework is not designed to allow mixed reasoning with beliefs and credences (Leitgeb 2016). Similarly for Lin and Kelly's odds threshold rule for reasoning with beliefs and credences (2013). However, I don't think the proponents of these frameworks are taking themselves to be committed to making empirical claims about whether people can use mixtures of beliefs and credences in their reasoning. In sum, I can't see a good reason why any proponent of the view that people have credences and outright beliefs would want to argue for "no mixing."

A more plausible way of denying the second claim of the puzzle is to argue that while we can technically change what we take to be live possibilities in different contexts, those changes rarely ever happen. And if we need to hardly ever pseudo-conditionalize to switch contexts, having to do so does not make excessive cognitive demands on thinkers like us. We could get the benefit of limiting the space of possibilities we need to engage with within a context, while hardly ever paying the cost of pseudo-conditionalizing. Whether this response succeeds is largely an empirical question. It depends on how often we shift between taking a claim for granted and not taking it for granted. I am not sure how to assess whether this happens infrequently enough to assuage the worry that context-shifting, and thus belief-shifting, via PC is too cognitively demanding. Yet, it is worth noting that endorsing this response is at odds with the motivations that underlie the "belief is credence 1 in context" view in the first place. Moreover, there is a large literature on how to accommodate context shifts, and the proposals range from views that propose shifty standards for which beliefs count as rational (contextualism and subject-sensitive invariantism fall in this category), to views that endorse shifty norms for what attitudes can be relied upon in a given context (see e.g. Levin 2008, Brown 2008). All of these views are supposed to be motivated by the observation that shifts between what we treat as true and what we treat as merely highly likely happen frequently, and not just on very rare occasions. Hence, it would seem odd if proponents of these "shifty" views were attracted to a solution to the puzzle that emphasizes how *rarely* our beliefs undergo contextual shifts. I don't want to rule out that this might be the correct solution to the puzzle, but it doesn't seem to be an attractive route given the motivations that led to the view of belief that gives rise to the puzzle in the first place. Also, the empirical explanation of switching offered by Nagel (2011), which says that people tend to switch to reasoning with credences when they shift from a more automatic to a more deliberate way of thinking, does not suggest that it is particularly rare.[11]

---

[11] An anonymous referee points out that some views of belief, such as the *belief-as-plan view* (see e.g. Dallmann 2017) seem hostile to the idea that we often switch between reasoning with beliefs and credences. The idea is that beliefs are stable in the sense that they are resistant to being constantly reconsidered, which helps us manage our cognitive load. Dallmann shows that it is in fact beneficial for limited thinkers to ignore evidence after a certain point, when it doesn't seem like considering additional evidence will have much impact on one's beliefs (or credences). However, the claim that we should sometimes switch from reasoning with outright beliefs to reasoning with credences is consistent with the idea that we should ignore

This leaves us with the third and fourth claims of the puzzle – that we pseudo-conditionalize on a set of background credences to generate our beliefs and credences in a context, and that doing so is psychologically difficult because it is computationally demanding. The claim that conditionalization is a computationally and cognitively demanding process is well supported by evidence from cognitive science and complexity theory, as I explained in the previous section. This evidence readily applies to pseudo-conditionalization, which is the same formal procedure as conditionalization. Hence, we don't have good grounds to question the fourth claim in the puzzle.[12]

Instead, I propose that we challenge the claim that we determine our outright beliefs and credences in a context by pseudo-conditionalizing. One might think at this point that we already have conclusive grounds to reject (3), since we have good reason to think that PC fails to be an accurate *description* of how thinkers manage their credences and beliefs across contexts. But (3) need not be read as making a descriptive claim about PC being an accurate formal representation of a psychological process. (3) can also be understood as claiming that PC is a normative principle that prescribes how ideally rational thinkers should manage their credences and outright beliefs. If PC is a normative principle, it is one that is too demanding to be complied with by ordinary human

additional evidence in some conditions. While Dallmann doesn't talk about switching in particular, he explicitly mentions that the belief-as plan view is compatible with the view that human thinkers have both credences and beliefs. Hence, I don't take there to be an immediate conflict between views like Dallmann's and the view presented here.

[12] According to recent psychological research, human thinkers can engage in different types of reasoning, sometimes called "System 1 reasoning" and "System 2 reasoning." While the exact details of these views are controversial, the idea is that System 1 reasoning is fast, automatic, and can handle large quantities of information. System 2 reasoning, by contrast, is under the thinker's deliberative control, takes up working memory, and can handle only small quantities of information (for some helpful recent discussion of the distinction, see Kahnemann 2011, Evans & Stanovich 2013, Mugg 2016). It has been suggested to me that perhaps the puzzle can be dissolved as follows: Since System 2 reasoning is constrained by the limitations of working memory, we need outright beliefs in order to simplify this type of reasoning. However, the processes that determine the credences and outright beliefs that a given act of System 2 reasoning employs are executed by System 1, which can handle more complex processing tasks. Hence, it is not a problem if the strategies for managing our outright beliefs and credences are complicated, because they are handled by System 1 processes. If this view were correct, it would ease the tension between the claim that outright beliefs have the function of simplifying reasoning, and the claim that PC is a very complicated strategy for managing outright beliefs and credences.

Unfortunately, this response to the puzzle is not empirically plausible. While it is plausible that the processes by which outright beliefs and credences are selected for use in a given reasoning task are automatic and don't use working memory, it is not plausible to assume that our System 1 computational resources can handle PC. The psychological evidence we have about whether human reasoners are conditionalizers allows for people to rely on System 1 reasoning in solving the relevant inference tasks. As explained earlier, the research shows that people at best approximate Bayesian reasoning in a lot of cases, but they are not perfect conditionalizers. If our automatic reasoning capacities can at best approximate conditionalization, they can at best approximate pseudo-conditionalization, as those are formally equivalent rules of inference. Moreover, it is also not empirically plausible to assume that outright beliefs only play a role in simplifying reasoning tasks that are executed deliberately and employ working memory. Most of our everyday reasoning is executed by automatic processes, and these processes need to treat many highly likely claims as true in order to make the necessary computations feasible.

thinkers, as we just saw. But perhaps this is unproblematic: other Bayesian principles of rationality, such as the requirement to have probabilistically coherent credences, are also unattainable by ordinary thinkers, but this does not disqualify them. They are usually characterized as rational ideals, which can function as regulatory aims for non-ideal thinkers.

Unfortunately, the move toward interpreting PC normatively is unsuccessful, as I will now argue. There is a crucial difference between requirements such as coherence on the one hand, and PC, viewed as a normative requirement, on the other hand. On the view expressed by claim (1) in the puzzle, the ability to reason with outright beliefs is a kind of heuristic that limited thinkers are equipped with to better cope with their limitations. Credences, by contrast, are not viewed as having a heuristic function in the same sense. Their role is independent of the cognitive limitations of the thinker who has them. But if outright beliefs, unlike credences, are essentially a heuristic tool, then whatever norms apply to them can't plausibly undermine their heuristic function. In other words: If there is a norm that is supposed to apply to some state or activity, the norm can't possibly be a true norm if it prevents the state or activity from serving its purpose or functioning correctly. Too see this point more clearly, we can consider some analogies to other devices or activities that are intended to function as heuristics.

First, imagine that Maximus University has recently made bikes available in various places, which are supposed to help students travel more quickly across the large campus. Now, suppose further that there is a rule for riding bikes on this campus, which states that to prevent accidents, riding faster than normal walking speed is not permitted. It's easy to see that if this were the rule for how to properly ride bikes on campus, there would be no point in having the bikes in the first place, because if people followed the rule, their travel times across campus would be just the same as if they walked. Hence, the bikes can't serve their purpose of making it faster to get across campus if bike riding is governed by this norm.

Another good example is a recent advice video that was posted on the Food Network's Instagram page.[13] The aim of the video was to give the audience a "life hack" for making sandwiches in a more practical way. Here's the "improved" method of making peanut butter sandwiches the video recommends: Take a piece of parchment paper and add a big spoonful of peanut butter. Fold the paper over, and roll over it with a rolling pin to flatten the peanut butter between the parchment. Freeze overnight, then cut into squares with kitchen shears. Then peel the paper off the peanut butter squares and put them on top of your bread.

The "life hack" proposed in the video quickly became the subject of ridicule online. People pointed out that instead of making it easier to prepare a peanut butter sandwich, which is generally the point of a "life hack," this method actually made it more

---

[13] https://www.youtube.com/watch?v=T7DsHMga92g
Readers who are looking for additional illustrations of this point may try googling "pointless life hacks."
The "Quick Tips" section in the magazine *Cook's Illustrated* is also a rich source of examples.

difficult and time consuming. Hence, the frozen parchment method can't possibly be an improvement over the standard method for making peanut butter sandwiches, since it does not simplify the task.

The analogy to managing our beliefs via PC is easy to see: the point of having beliefs is supposed to be that they help simplify reasoning. Any plausible norm that governs managing our beliefs and credences across contexts must be compatible with beliefs serving their intended function. But since PC is too complicated for human reasoners to use as a managing strategy for their beliefs and credences, it is not a plausible norm governing this activity, since it hinders outright beliefs' function of simplifying reasoning. Just like the speed limit prevents bikes from serving the purpose of decreasing travel times across campus, and the parchment paper method prevents people from efficiently making peanut butter sandwiches, requiring thinkers to manage their credences and beliefs via PC prevents outright beliefs from effectively simplifying reasoning processes. Hence, interpreting claim (3) in the puzzle normatively instead of descriptively does not enhance its plausibility.

In sum, I've argued that we should resolve the puzzle by rejecting claim (3). It is implausible regardless of whether we interpret is as a descriptive claim or as a normative claim. But this means that the defender of the remaining claims in the puzzle, and in particular claim (1), lacks an explanation of how our credences and outright beliefs are determined in a given context of reasoning. Of course, it is ultimately an empirical question of how human thinkers manage their beliefs and credences across contexts. My aim is not to resolve this empirical question here. Instead, I will consider some possible strategies human might use to manage their beliefs and credences, which suggest that simpler reasoning strategies often carry the cost of leading to synchronically and diachronically incoherent attitudes. I will moreover argue for the normative claim that my solution to the puzzle has implications for the debate about normative constraints on belief – endorsing the heuristic view of outright belief requires that putative norms on belief must be vetted for their complexity and feasibility for human reasoners.

3. The Descriptive and Normative Consequences of Rejecting PC

Using PC is infeasible for limited human reasoners because it requires that we update the remaining credences whenever the number of possibilities that are considered live is adjusted. This is an excessively taxing task for thinkers like us, who don't store all of our beliefs and credences explicitly. Hence, human reasoners must be relying on a simpler procedure for managing their beliefs and credences. We know from the study of approximative Bayesian reasoning that any rule that is simpler than PC will fall short of producing the same outcome as PC in at least some contexts, even if it is some kind of approximating algorithm (Kwisthout et al. 2011, see also Predd et al. 2008).

It of course ultimately an empirical question how people in fact manage their beliefs and credences across contexts. Which procedure works well for human thinkers depends, among other things, on facts about their cognitive architecture into which we

currently have very little insight. Still, thinking about different possible strategies for managing beliefs and credences across contexts can give us some clues about their potential costs and benefits for human thinkers.

We observed earlier that PC has some desirable features. Thinkers who are representable as pseudo-conditionalizers have stable conditional credences, which means that their assessment of the impact of the evidence they may receive doesn't change. Moreover, employing PC guarantees that a thinker has coherent credences within each context of reasoning. Still, using PC inevitably introduces some diachronic incoherence, in the sense that the thinker changes her credence without receiving new evidence. This of course happens by design, since the whole point of PC is to treat options as ruled out when the evidence deems them to be merely unlikely. Hence, no alternative procedure would avoid introducing this kind of diachronic incoherence. By contrast, alternative strategies for managing beliefs and credences might introduce instability (and thus incoherence) in the thinker's conditional credences across contexts, as well as incoherence within a context of reasoning.

To see more concretely how this might play out, we will consider two slightly different proposals that have recently been made in the literature about how credences and beliefs are stored by human thinkers, and we'll see what kinds of procedures for managing beliefs and credences across contexts might naturally pair with them. One such proposal has been made by Norby (2015), who argues for the following view: Thinkers represent a space of possibilities, of which a subset is selected by some automatic filtering process to be considered in a specific context of reasoning. Once such a subset is selected, credences are then assigned to the possibility space under consideration. For example, if the possibility space that is called up includes the possibility that I have milk in my fridge, but not the possibility that I don't, then it is treated as true (and thus assigned credences 1 in context) that I have milk in my fridge. It is a bit less clear how Norby thinks intermediate credences are determined in a given context. He suggests that the underlying full possibility space does not include an assignment of credences. Rather, different possibilities in the space have different probabilities of being selected for consideration in a given context. However, this is of course not the same thing as the thinker's credence in a possibility: The fact that a possibility is likely to be included for consideration in a variety of contexts is compatible with it being given low credence in those contexts. Norby emphasizes that the same possibility can receive different credence assignments in different contexts, but it remains a bit unclear what determines a possibility's non-extreme credence assignment in a given context.

Weisberg (forthcoming) offers a more concrete answer to this question. He shares the view that thinkers don't explicitly encode beliefs or credences, but instead represent possibilities or pieces of information at a fundamental level. Citing empirical evidence about the recall of memories, he suggests the view that outright beliefs and credences are generated partly on the basis of features of the recall process itself. He explains:

Confidence in memory-based beliefs appears to be constructed at the time of recall, rather than stored. If you're asked what the capital of Iceland is, the more easily the answer (Reykjavik) comes to mind, and the more related information comes to mind, the more certain you will be that your answer is correct. So your confidence that Reykjavik is the capital of Iceland doesn't appear to be stored in memory, at least not directly. […] If nothing about Iceland is stored in memory, nothing will come to mind at the time of recall, and you will have virtually no confidence that the capital of Iceland is a place called 'Reykjavik'. (p.26)

Of course, Weisberg's view refers specifically to memory-based beliefs, so we need to be careful not to overgeneralize it. Combining Norby's and Weisberg's views, we arrive at the following rough and ready picture of our mental architecture: We represent a space of possibilities, and in a given context of reasoning, a subset of this space is selected for consideration. Some possibilities are recalled without their complements – they constitute what is believed or disbelieved in that context. Factors like salience and easy recall presumably influence which possibilities are awarded this status. Other possibilities are given less than full confidence, and how much confidence they get is at least sometimes and at least partly determined by how smoothly we can call up these possibilities and what evidence is recalled in favor and against the relevant possibility.

Accepting this picture as our working hypothesis, we can see several ways in which deviations from PC might crop up. First, it is not clear how conditional credences enter the picture. If they are not stored or generated in a stable manner, but also made up "on the fly", then we might encounter diachronic incoherence in the thinker's conditional credences, i.e. they might change between contexts. This would constitute a deviation from PC. Another type of diachronic incoherence has been observed to occur when possibilities are represented with a different fineness of grain in different contexts. Norby illustrates this with data on the well-known unpacking effect: Suppose I can represent an event either as the coarse partition $\{E, \sim E\}$, or in a finer-grained manner as $\{E_1,…., E_n, \sim E\}$. For example, this could be {I buy a car, I don't buy a car} vs. {I buy a VW, I buy a Porsche, …, I don't buy a car}. Thinkers tend to assign E a lower probability when it is presented as one option, compared to when it is unpacked, i.e. the sum of the credences assigned to E_1 through E_n is larger than the credence assigned to E. Norby takes this data to provide evidence for his view of how credences are represented, according to which thinkers lack stable credences across contexts. We can thus observe that thinkers who lack stable background credences are vulnerable to a kind of diachronic incoherence which does not arise for adherents of PC.

We may also consider whether the Norby/Weisberg picture should lead us to expect incoherence in a thinker's credences within a given context of reasoning. Again, my remarks here will be somewhat speculative, since we don't have detailed knowledge of how the process of assigning credences in a given context of reasoning works. Suppose, as Weisberg claims, that each possibility is assigned a credence based on features of the

thinker's memory and the recall episode. If this happens for each possibility separately, we could end up with slightly incoherent credences even within a context of reasoning. Here's how this might happen. Suppose a thinker's full possibility space regarding some issue E (say, where the tennis match will take place) contains five possibilities, $E_1 - E_5$. If all possibilities were considered, E_1 and E_2 would be considered very unlikely, say 5% likely each. E_3 would be considered the most likely, say 50%, and the remaining two possibilities are each considered 20% likely. Now, in some contexts, the thinker might not even consider E_1 and E_2 at all. But if the mechanism for assigning credences to the remaining three options remains the same as in a context in which all five possibilities are considered, then the thinker might end up assigning credences to E_3, E_4 and E_5 that sum to slightly less than 100%. Unless some procedure is in place that ensures that the credences in the options under consideration are normalized, so that they sum to 100%, the "recall and assign each credence on the fly" procedure does not guarantee coherence within a context. Such a normalization procedure requires computational effort, so it is a place in which our mind might cut corners. Of course, it is important to emphasize again that whether we employ such a normalization procedure or not is ultimately an empirical question. But as described, the Norby/Weisberg picture does not rule out the possibility of incoherence within a context. If there is no normalization procedure that coherentizes the credences in a given context, we again arrive at deviations from PC, this time within contexts of reasoning.

In the literature on the relationship between credences and beliefs, some philosophers have proposed threshold rules that determine what is believed in a context. For example, Foley (2009) proposes a descriptive version of the Lockean thesis, which says that each context is associated with a credence threshold, such that "one believes that P just in case one is sufficiently confident of the truth of P." Call this rule *General Rounding (GR)*. Alternatively, there could be a rule according to which a thinker treats only some claims as true (false) that are assigned a credence above (below) some threshold. Call this rule *Selective Rounding (SR)*.[14] Questions concerning the complexity of the implementation of such rules, and how this affects the coherence of a thinker's attitude, don't usually get explicitly discussed. But it is worth noting that if those rules were implemented without employing an additional renormalization procedure, they would also lead to incoherence within contexts. We can see the effects that failures to renormalize have on the coherence of a thinker's credences by considering an example. Here are some credences that constitute a subset of some thinker's full credence function. Within this subset, no possibilities have currently been ruled out for the purpose of simplifying reasoning:

$Cr(A\&B) = 0.86$
$Cr(A\&\sim B) = 0.1$
$Cr(\sim A\&B) = 0.01$

---

[14] I am grateful to Sylvia Wenmackers for helping me see that GR and SR should be distinguished.

Cr(~A&~B) = 0.03

Cr(A) = 0.96
Cr(~A) = 0.04
Cr(B) = 0.87
Cr(~B) = 0.13

Suppose that the thinker eliminates some of the possibilities by using either *GR* or *SR*. If the remaining, unrounded credences were determined in the same way, regardless of whether some possibilities are eliminated from consideration, the thinker is left with an incoherent credence function within a context. Although this kind of procedure introduces slight incoherence, it eliminates the computational cost of renormalizing the remaining credences depending on which possibilities are considered live, and is thus less difficult to execute than PC.

Below, the table shows how the different rules would determine the thinker's credences in a given context. The columns PC(A) and PC(B) show the thinker's attitudes if she pseudo-conditionalized on A and B, respectively. The columns GR(>0.85) and GR(>0.9) show how the thinker's attitudes would turn out as a result of applying the General Rounding rule with different thresholds. The columns SR(A) and SR(B) show the thinker's attitudes resulting from selective rounding, treating either only A or only B as true.

Credences only	PC(A)	PC(B)	GR (>0.85)	GR (>0.9)	SR(A)	SR(B)
Cr(A&B) = 0.86	0.895	0.989	1	0.86	0.86	0.86
Cr(A&~B) = 0.1	0.105	0	0	0.1	0.1	0
Cr(~A&B) = 0.01	0	0.011	0	0	0	0.001
Cr(~A&~B) = 0.03	0	0	0	0	0	0
Cr(A) = 0.96	1	0.989	1	1	1	0.96
Cr(~A) = 0.04	0	0.011	0	0	0	0.04
Cr(B) = 0.87	0.895	1	1	0.87	0.87	1
Cr(~B) = 0.13	0.105	0	0	0.13	0.13	0

Looking at these results, we can easily see that since they don't renormalize the remaining non-extreme credences, they tend to generate some incoherence in the thinker's credence function within a context of reasoning.

There is much more to be learned about how humans store beliefs and credences, and how they reason with them. Examining some sample theories has been instructive in shedding light on open questions in this area. Studying the Norby/Weisberg proposal is instructive, because we saw how their sketch of the mental architecture of human belief

suggests processes for managing beliefs and credences across contexts that can give rise to violations of PC. First, incoherence between contexts can take hold if there is no stable basis of conditional credences that remains fixed over time. The human mind might lack such a stable basis in order to avoid excessive storage demands. Second, the view on which each credence is generated in a somewhat individualized manner based on what is stored in memory and features of the recall episode suggests that incoherence might also arise within a context of reasoning. If the assignment of individual credences is somewhat insensitive to what other possibilities have been ruled out, this might generate failures to renormalize credences within a context, which in turn generates incoherence. We also considered threshold rules for determining what is believed within a context, and noted that if those rules are not implemented alongside a renormalization procedure, they, too, generate incoherence within a context.

The main point to take away from this discussion is that there are a variety of different ways in which human reasoners might manage their beliefs and credences across contexts, and depending on exactly how their mental architecture is set up, there will be different heuristics that make this task manageable. Different heuristics will lead to different types of deviations from PC, and thus different forms of incoherence within a context or between contexts of reasoning. (Remember that any strategy that avoids incoherence is equivalent to PC, and thus just as difficult to execute.)

Using a heuristic for managing beliefs and credences across contexts that introduces some synchronic or diachronic incoherence might seem initially like a bad deal – after all, there are numerous arguments in the literature showing that incoherent credences are accuracy-dominated (Joyce 2009, Pettigrew 2016), and vulnerable to Dutch books (see e.g. Hájek 2008). However, incoherence comes in degrees, and the bad consequences of being incoherent equally come in degrees (Schervish et al. 2000, 2002, 2003, De Bona & Finger 2015, De Bona & Staffel 2017, 2018). We can intuitively think of the degree of incoherence of a credence function as the distance between it and some closest coherent credence function. Small divergences from coherence lead to small deviations from optimal accuracy, and to relatively little Dutch book vulnerability.[15] If thinkers generally only treat claims as true that they take to be highly probable, the incoherence generated by not renormalizing the remaining credences within a context of reasoning is generally not going to be very dramatic. This is because, if the thinker's credence in some claim shifts by only a little, say by 0.05, then the adjustments to the thinker's remaining credences this would normally require according to PC are also quite minimal, and so if those adjustments aren't made, the divergence from coherence created this way is not very large. Similar claims hold for slight incoherence between different contexts of reasoning. As a result, the loss in accuracy, and the increased Dutch book

---

[15] This holds provided that we measure Dutch book vulnerability in a normalized manner, i.e. we control for the size of the bets involved. For details regarding incoherence measures, and how degrees of incoherence related to degrees of inaccuracy and Dutch book vulnerability, the reader may consult the cited sources.

vulnerability to which this amount of incoherence gives rise is also not very dramatic. These considerations suggest that trading simplicity for coherence can actually be beneficial for thinkers, provided these tradeoffs don't generate incoherence that is too significant. Whatever method thinkers actually use for managing their credences and outright beliefs, we know, based on empirical research of the kind I cited in section two, that (with some exceptions such as base-rate cases) human reasoners tend to be decently good at *approximating* coherent Bayesian reasoning, even if they are usually not perfectly coherent.

One might still worry, however, that in some cases, ignoring a small probability of error or being slightly incoherent could have catastrophic consequences. For example, the probability that I might have a grave accident while on vacation might be very small, but if it happened, the cost of treatment could bankrupt me. Hence, when I decide whether to buy travel insurance, I should not simply treat it as true that I won't have an accident, even if my credence in this claim is extremely high. Even if ignoring small error probabilities and being slightly incoherent are unproblematic in many contexts, in high stakes situations doing so is extremely inadvisable. Fortunately, human thinkers seem to be sensitive to this issue. It has been extensively documented in the literature on pragmatic encroachment that reasoners tend to switch to relying on credences in relevant propositions in high stakes contexts (Roeber 2016). The famous bank case illustrates this (De Rose 1992): When nothing much rides on whether the bank is open on Saturday, people seem to be inclined to treat it as true for the purposes of reasoning if they are fairly confident in it. But when it's extremely important that they get to the bank, people start paying attention to the small chance that they might be mistaken. Hence, in cases where small error probabilities matter, people often seem to attend to these small error risks more carefully. But of course, more empirical research is needed to pin down the reasoning strategies that give rise to this data. We learn from these observations that for human thinkers, using a heuristic for managing beliefs and credences that differs from PC need not be disadvantageous. Given a smart approximation strategy, the cost of the resulting incoherence need not be significant.

A further lesson from rejecting (3) and endorsing a heuristic view of outright belief concerns constraints on what consistency or coherence norms on belief can plausibly look like. Standardly, justifications of synchronic and diachronic norms of belief appeal to purely epistemic grounds. Such grounds include, for example, whether the norms promote particular aims such as believing the truth and not believing falsehoods, having properly based or well-justified beliefs, and so on. These defenses thus tend to consider the epistemic credentials of these norms, but they usually don't attend to how easy or difficult it would be for human thinkers to abide by them. But, as I have argued above, if the main function of outright beliefs is to be a heuristic tool that simplifies the reasoning processes of limited thinkers, then, if N is a norm that governs beliefs, N must not prevent outright beliefs from serving this function. This means that, if one endorses the heuristic view of the function of outright belief, one is thereby committed to thinking that suitable

norms of belief must meet both epistemic constraints *and* feasibility constraints. It is insufficient to argue for norms of belief based on epistemic constraints alone. An example of a view on which my argument puts pressure is Leitgeb's stability theory of belief (Leitgeb 2016). In the first chapter of his book, he endorses the view that outright beliefs simplify reasoning (although it is not entirely clear that he is committed to the view that this is their primary function). He then goes on to endorse norms on belief, as well as norms on how credences and beliefs may be connected in rational thinkers, that share many features with standard ideal consistency and coherence norms familiar from Bayesian epistemology. Leitgeb readily admits that these ideal norms can't realistically be obeyed by human thinkers, yet does not recognize that endorsing such norms is in tension with his claim that the function of beliefs is to simplify reasoning. I will leave it open here how Leitgeb and philosophers who hold similar views should resolve this tension. It is an important and unexplored question for future research which norms on belief are compatible with their simplifying role.

Conclusion

In this paper, I presented a new puzzle about outright belief. The puzzle arises from combining a popular answer to the Bayesian challenge – that outright beliefs have the function of simplifying reasoning – with some additional plausible claims about outright belief that have been recently defended in the literature. Those claims are that we can mix outright beliefs and credences in reasoning, that we can switch between relying on a high credence and relying on an outright belief in different contexts, and that thinkers manage their credences and outright beliefs across contexts via the pseudo-conditionalization strategy. The tension between these claims arises from the fact that PC, while having some features that appear desirable from a normative point of view, is computationally demanding. PC can be interpreted descriptively or normatively, but it is implausible on either interpretation. Its descriptive version is incompatible with known results about human reasoning, specifically the result that they aren't perfect conditionalizers. The normative version of PC is implausible, because it would require outright beliefs to be governed by a norm that impedes their function as a simplifying heuristic.

I then considered which lessons we can draw from this solution to the puzzle. A natural question to ask is what alternative strategies are available to human thinkers for managing their beliefs and credences, given that PC is an inadequate description of their thought processes. While it is ultimately an empirical question which strategies thinkers use, I argued that we should expect tradeoffs between simplicity and coherence. PC already induces some diachronic incoherence in the thinker's attitudes, and any strategy that is not equivalent to PC will incur additional incoherence either within or between contexts of reasoning. Incoherence is generally taken to be problematic, since it leads suboptimal accuracy and Dutch book vulnerability. Yet, since the kinds of feasible strategies thinkers might use to manage their outright beliefs and credences are likely to

introduce only a fairly minimal amount of incoherence, the tradeoff can be beneficial for the thinker. I further argued that we can draw an important normative lesson from this solution to the puzzle. Endorsing the heuristic view of outright belief captured in claim (1) of the puzzle entails that putative norms on belief can't be defended on purely epistemic grounds. They must also be vetted for their complexity and feasibility for human reasoners. If outright beliefs' primary function is to simplify reasoning, they can't possibly be governed by norms that prevent them from accomplishing this.

References

Brown, Jessica (2008): "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning," *Noûs* 42 (2), 167–189.

Buchak, Lara (2014): "Belief, Credence, and Norms," *Philosophical Studies* 169 (2), 285–311.

Carr, Jennifer (forthcoming): "Subjective Probability and the Content/Attitude Distinction," *Oxford Studies in Epistemology*.

Clarke, Roger (2013): "Belief Is Credence One (In Context)", *Philosophers' Imprint* 13, 1–18.

Cooper, Gregory F. (1990): "The Computational Complexity of Probabilistic Inference Using Bayesian Belief Networks", *Artificial Intelligence* 42 (2-3), 393-405.

Dallmann, Justin (2017): "When Obstinacy is a Better (Cognitive) Policy", *Philosophers' Imprint* 17 (24), 1–17.

De Bona, Glauber & Finger, Marcelo (2015) "Measuring Inconsistency in Probabilistic Logic: Rationality Postulates and Dutch Book Interpretation," *Artificial Intelligence* 227, 140–64.

De Bona, Glauber & Staffel, Julia (2017): "Graded Incoherence for Accuracy-Firsters," *Philosophy of Science.* (online first)

De Bona, Glauber & Staffel, Julia (2018): "Why be (approximately) coherent?" *Analysis.* (online first)

DeRose, Keith (1992): "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52 (4), 913–929.

Easwaran, Kenny & Fitelson, Branden (2015): "Accuracy, Coherence, and Evidence," *Oxford Studies in Epistemology*, Vol. 5, 61–96.

Edwards, Ward (1982): "Conservatism in human information processing," in: Daniel Kahnemann, Paul Slovic & Amos Tversky: *Judgment under Uncertainty: Heuristics and Biases*, Cambridge: CUP, 359–369.

Evans, Jonathan St. B. T. & Stanovich, Keith E. (2013): "Dual-Process Theories of Higher Cognition. Advancing the Debate," *Perspectives on Psychological Science* 8 (3), 223–241.

Evans, Jonathan St. B. T., Thompson, Valerie A. & Over, David E. (2015): "Uncertain deduction and conditional reasoning," *Frontiers in Psychology* 6, 398.

Foley, Richard (2009): "Beliefs, Degrees of Belief, and the Lockean Thesis," in: F. Huber & C. Schmidt–Petri (eds), *Degrees of Belief*, Synthese Library 342, Springer, 2009, 37–47.

Greco, Daniel (2017): "Cognitive Mobile Homes," *Mind* 126 (501), 93-121.

Greco, Daniel (2015): "How I Learned to Stop Worrying and Love Probability 1," *Philosophical Perspectives* 29, 179–201.

Harman, Gilbert (1986): *Change in View*, Cambridge: MIT Press.

Harsanyi, John C. (1985): "Acceptance of Empirical Statements: A Bayesian Theory Without Cognitive Utilities," *Theory and Decision* 18 (1), 1–30.

Hájek, A. (2008): "Dutch book arguments," in P. Anand, P. Pattanaik, & C. Puppe (Eds.), *The Oxford Handbook of Rational and Social Choice*. Oxford: Oxford University Press, 173–195.

Hawthorne, John (2004): *Knowledge and Lotteries*, Oxford: Oxford University Press.

Hawthorne, John & Stanley, Jason (2008): "Knowledge and Action," *Journal of Philosophy* 105 (10), 571–590.

Jeffrey, Richard (1970): "Dracula meets Wolfman: Acceptance vs. Partial Belief," in: M. Swain (ed.), *Induction, Acceptance, and Rational Belief*, Dordrecht: D. Reidel, 1970, 157–185.

Joyce, J. M. (2009): "Accuracy and coherence: Prospects for an alethic epistemology of partial belief," in F. Huber and C. Schmidt–Petri (Eds.), *Degrees of belief*, Vol. 342 of *Synthese Library*, Berlin: Springer, 263–297.

Kahnemann, Daniel (2011): *Thinking Fast and Slow*, New York: Farrar, Straus and Giroux.

Kirsh, David (2003): "Implicit and Explicit Representation," in: Lynn Nadel et al. (ed.), *Encyclopedia of Cognitive Science.* Wiley, 478–481.

Kwisthout, Johan, Wareham, Todd & van Rooij, Iris (2011): "Bayesian Intractability is not an Ailment that Approximation Can Cure." *Cognitive Science* 35(5), 779–784.

Lance, Mark Norris (1995): "Subjective Probability and Acceptance," *Philosophical Studies*, Vol. 77, 147–179.

Leitgeb, Hannes (2016): *The Stability of Belief. How Rational Belief Coheres with Probability*, Oxford: Oxford University Press.

Levi, Isaac (1964): "Belief and Action," *The Monist* 48 (2), 306–316.

Levin, Janet (2008): "Assertion, Practical Reason, and Pragmatic Theories of Knowledge," *Philosophy and Phenomenological Research* 76 (2), 359–384.

Lin, Hanti (2013): "Foundations of Everyday Practical Reasoning," *Journal of Philosophical Logic* 42 (6), 831–862.

Lin, Hanti (2014): "On the Regress Problem of Deciding How to Decide," *Synthese* 191, 661–670.

Lin, Hanti & Kelly, Kevin (2012): "Propositional Reasoning that Tracks Probabilistic Reasoning," *Journal of Philosophical Logic* 41 (6), 957–981.

Mastropasqua, Tommaso, Crupi, Vincenzo & Tentori, Katya (2010): "Broadening the study of inductive reasoning: Confirmation judgments with uncertain evidence," *Memory & Cognition* 38 (7), 941–950.

Moss, Sarah (2018): *Probabilistic Knowledge*, Oxford: Oxford University Press.

Mugg, Joshua (2016): "The dual-process turn: How recent defenses of dual-process theories of reasoning fail," *Philosophical Psychology* 29 (2), 300–309.

Nagel, Jennifer (2011): "The Psychological Basis of the Harman-Vogel Paradox," *Philosophers' Imprint* 11 (5), 1–28.

Norby, Aaron (2015): "Uncertainty Without All the Doubt," *Mind and Language* 30 (1), 70–94.

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

Pettigrew, Richard (2017): "Précis and Replies to Contributors for Book Symposium on *Accuracy and the Laws of Credence*," *Episteme* 14 (1), 1–30.

Predd, J.B., Osherson, D., Kulkarni, S. & Poor, H. V. (2008): "Aggregating Probabilistic Forecasts from Incoherent and Abstaining Experts," *Decision Analysis* 5 (4), 177–189.

Roeber, Blake (2016): "The Pragmatic Encroachment Debate," *Noûs*. (online first)

Ross, Jacob & Schroeder, Mark (2014): "Belief, Credence, and Pragmatic Encroachment," *Philosophy and Phenomenological Research*, Vol. 88, No. 2, 259–288.

Schervish, Mark, Seidenfeld, Teddy & Kadane, Joseph (2000): "How Sets of Coherent Probabilities May Serve as Models for Degrees of Incoherence," *International Journal of Uncertainty, Fuzziness and Knowlegde-based Systems* 8, 347–355.

Schervish, Mark, Seidenfeld, Teddy & Kadane, Joseph (2002): "Measuring Incoherence." *Sankhya: The Indian Journal of Statistics* 64, 561–587.

Schervish, Mark, Seidenfeld, Teddy & Kadane, Joseph (2003): "Measures of Incoherence: How not to Gamble if You Must," in: *Bayesian Statistics 7*, ed. José Bernardo, et. al., 385–401. Oxford: Oxford University Press.
Shanahan, Murray (2016): "The Frame Problem," *The Stanford Encyclopedia of Philosophy*, Spring 2016 Edition, ed. Edward N. Zalta.

URL = <https://plato.stanford.edu/archives/spr2016/entries/frame-problem/>

Slovic, Paul & Liechtenstein, Sarah (1971): "Comparison of Bayesian and Regression Approaches to the Study of Information Processing in Judgment," *Organizational Behavior and Human Performance* 6, 649–744.

Staffel, Julia (2017): "Accuracy for Believers," *Episteme* 14 (1), 39–48.

Staffel, Julia (2016): "Beliefs, buses and lotteries: Why rational belief can't be stably high credence," *Philosophical Studies* 173 (7), 1721–1734.

Sturgeon, Scott (2015): "The Tale of Bella and Creda," *Philosophers' Imprint* 15 (31), 1-9.

Tang, Weng Hong (2015): "Belief and Cognitive Limitations," *Philosophical Studies* 172 (1), 249–260.

Tentori, Katya, Chater, Nick & Crupi, Vincenzo (2016): "Judging the Probability of Hypotheses Versus the Impact of Evidence: Which Form of Inductive Inference is More Accurate and Time-Consistent?" *Cognitive Science* 40, 758–778.

Weatherson, Brian (2016): "Games, Beliefs and Credences," *Philosophy and Phenomenological Research* 92 (2), 209–236.

Wedgwood, Ralph (2012): "Outright Belief," *Dialectica* 66 (3), 309–329.

Weisberg, Jonathan (forthcoming): "Belief in Psyontology", *Philosophers' Imprint*.

Weisberg, Jonathan (2013): "Knowledge in Action," *Philosophers' Imprint* 13 (22), 1-23.

Williamson, Timothy (2000): *Knowledge and Its Limits*. Oxford: Oxford University Press.

Zhao, Jiaying, Crupi, Vincenzo, Tentori, Katya, Fitelson, Branden & Osherson, Daniel (2012): "Updating: Learning versus supposing," *Cognition* 124, 373–378.

Zhao, Jiaying & Osherson, Daniel (2010): "Updating beliefs in light of uncertain evidence: Descriptive assessment of Jeffrey's rule," *Thinking & Reasoning* 16 (4), 288–307.

Zhao, Jiaying, Shah, Anuj & Osherson, Daniel (2009): "On the provenance of judgments of conditional probability," *Cognition* 113, 26–36.