

Cognitive Mobile Homes

Daniel Greco
Penultimate Draft

Abstract

While recent discussions of contextualism have mostly focused on other issues, some influential early statements of contextualism emphasized the possibility for contextualism to provide an alternative both to coherentism and to traditional versions of foundationalism.¹ In this essay, I will pick up on this strand of contextualist thought and argue that contextualist versions of foundationalism promise to solve some problems that their non-contextualist cousins cannot. In particular, I will argue that adopting contextualist versions of foundationalism can let us reconcile Bayesian accounts of belief updating with a version of the holist claim that all beliefs are defeasible.

Introduction

Epistemic contextualists often motivate their position by arguing that contextualism provides a satisfying resolution of certain skeptical paradoxes.² In this essay, however, I will present one version of a very different strategy for motivating contextualism.³ The

For helpful comments and discussion, thanks to David James Barnett, Jessica Brown, Sinan Dogramaci, Brian Hedden, Sophie Horowitz, John Pittard, Jack Marley-Payne, Ram Neta, Augstín Rayo, Damien Rochford, Miriam Schoenfield, Robert Stalnaker, Jason Stanley, Jonathan Vogel, Roger White, the Corridor reading group, audiences at Edinburgh and Rutgers, and two anonymous referees at *Mind*.

¹See (author?) 1977 and (author?) 3. Arguably, it was also anticipated by (author?) 62. In his recent book-length defense of contextualism, (17) brings up the idea that contextualism has some special relevance to the debate between foundationalism and coherentism, only to set it aside (pp.21-22). I agree with DeRose that there is no entailment from contextualism to either foundationalism or coherentism. But as the rest of this essay should make clear, I do think that contextualism and foundationalism complement each other nicely.

²For paradigm instances of this strategy, see (author?) 52, (author?) 13, (author?) 16, and (author?) 32.

³Actually, as will become clear later, while I will focus on contextualism for ease of exposition, my hope is that the strategy can be employed by non-contextualists of various stripes as well. Epistemic relativists (34; 35) and expressivists (10) have all offered diagnoses of various phenomena that are structurally similar to contextualist diagnoses. My hope is that any view that can mimic familiar contextualist treatments of, e.g., skeptical paradoxes, (e.g., as the relativist can by appealing to assessment-sensitivity in many of the places where the contextualist would appeal to context-sensitivity) can also mimic the contextualist treatment of foundationalism that I will offer here, and that the issues that separate contextualism from a variety of rival views will, for the most part, not be relevant to the present discussion.

strategy I will explore attempts to motivate contextualism by arguing that it provides an appealing way to chart a middle course between coherentism on the one hand, and traditional, non-contextualist versions of foundationalism on the other.

These two strategies are independent of one another—one can hold that contextualism provides a satisfying resolution of skeptical paradoxes without holding that it provides any special advantage to the foundationalist,⁴ just as one can endorse contextualist versions of foundationalism without using them to offer the familiar contextualist treatment of skeptical paradoxes.⁵ I focus on the less common strategy because I believe it points the way to some surprising further potential applications of contextualism that do not suggest themselves quite as readily when contextualism is motivated via skeptical paradoxes.

1 Foundationalism

I will use the term ‘foundationalism’ to refer to a view about the structure of epistemic justification. In particular, I will use it to refer to the view that both of the following claims hold: (1) at least some beliefs are justified even without receiving support from other beliefs—I will call such beliefs, should any exist, ‘foundational’—and (2) all justified beliefs are either foundational, or derive all their support (perhaps indirectly) from justified foundational beliefs.⁶

A natural question for the foundationalist is the following: what criteria must beliefs meet in order to be foundational? Historically influential versions of foundationalism held that beliefs must meet very strict criteria in order to be foundational—perhaps they must be impossible to doubt, or must have contents that are necessarily true whenever believed, or must have some other strong modal properties. It is easy to see why one might find such properties epistemologically interesting. If it were possible to rationally reconstruct our entire body of beliefs from some sparse set of claims that are impossible to doubt, or that must be true (given that we believe them), doing so would be a great intellectual achievement. Criticisms of traditional versions of foundationalism have tended not to target the desirability of completing such a project, but the feasibility. It is not clear whether *any* beliefs have such strong modal properties, and even if some do (e.g., perhaps the conclusion of Descartes’ cogito), they are almost certainly too few to serve as an adequate foundation for the rest of what we normally take ourselves to be justified in believing.

In light of these considerations, contemporary foundationalists tend to defend views on which the requirements for foundational beliefs are more lax, and particular, need

⁴As DeRose does in the passage I refer to in footnote 1.

⁵E.g., while Michael (58) does endorse a sort of contextualist version of foundationalism—he calls it ‘formally foundationalist’—he offers a very different treatment of skepticism from that found in the writings of Stine, Cohen, DeRose, and Lewis. In particular, he does not concede that there is *any* context in which global skeptical challenges are successful.

⁶While ‘foundationalism’ is used in too many different ways to speak of a ‘standard’ use, my formulation is similar to what many writers call ‘minimal’ foundationalism (2).

not be free from the threat of doubt.⁷ In this first section, however, I want to show how one might be motivated to retreat, not to a laxer invariantist form of foundationalism, but instead to a view according to which the criteria beliefs must meet in order to be foundational varies with context.

Suppose we are sympathetic to the Cartesian idea that freedom from doubt really is the *sine qua non* for foundational status. After all, we will not be happy treating the claim that *P* as though it is foundational—e.g., justifying other beliefs by showing how they are supported by the claim that *P*—if we doubt that *P*. If we allow that which beliefs play the foundational role may vary from context to context, however, then we can insist that foundations must be undoubted, without taking the Cartesian position that they must be undoubtable. We might summarize the two ways of respecting the idea that foundations must be free from doubt as follows:

Cartesian: For a belief to be foundational, it must be undoubtable.

Contextualist: For a belief to be foundational (in a context), it must be undoubted (in that context).

Both the Cartesian and the contextualist agree that our justified beliefs rest on a foundation of undoubted premises. The Cartesian thinks there are certain types of beliefs that can never be in doubt, and so can always play a foundational role. The contextualist disagrees—because she thinks all beliefs are dubitable, she does not think that any beliefs can play a foundational role in *every* context.⁸ So while she will agree with the Cartesian in insisting that foundations be undoubted, she only requires *contextually local* freedom from doubt, rather than immunity to doubt across all contexts.

A natural thought for the contextualist foundationalist is that when a belief counts as ‘foundational’ in a context, it is treated by the occupants of that context as (at least provisionally) a default assumption in reasoning and argument. Foundationalism is plausible because we must always treat *some* beliefs as default assumptions in reasoning and argument. But it is also plausible that *no* belief will (or should) be treated as a default assumption for reasoning and argument in *every* context. Rather, any belief that we treat as a default assumption in one context, we can treat as a mere hypothesis in another, to be accepted only if it can be supported on independent grounds—i.e., only if it receives support from what we treat as default assumptions in our *new* context.

While the metaphor of Neurath’s raft is often appealed to by coherentists, I think it is ultimately most congenial to the contextualist foundationalist. The metaphorical way of putting the point of the previous paragraph is as follows: at any given time, we must stand *somewhere* on the raft, and we cannot examine the planks on which we are currently standing (i.e., each context treats *some* beliefs as foundational, so the coherentist is wrong). But still, we can always shift our weight to new planks, so as to

⁷For example, Gilbert (24) defends a view on which *all* beliefs count as foundational. Phenomenal conservatives defend the only slightly less permissive view that all beliefs that *seem* true to their subjects are foundational for those subjects.

⁸Or if she grants that *some* beliefs are indubitable, she thinks they are too few to serve the Cartesian’s purposes.

be able to examine the ones on which we were previously standing (i.e., no belief counts as foundational in every context, so the non-contextualist foundationalist is wrong).

Why think of contextualist foundationalism as a sort of middle course between traditional versions of foundationalism on the one hand, and coherentism on the other? While others have defended this claim at greater length,⁹ the rough idea is that contextualist foundationalism is meant to avoid the most serious pitfalls for the two more traditional views about the structure of justification. With the traditional foundationalist, the contextualist foundationalist can accept the attractive claim that justification by coherence alone is impossible, because it would be circular. But with the coherentist, the contextualist foundationalist can agree that there is not any distinctive class of beliefs well suited to play the regress-stopping role required by foundationalism.¹⁰ Ultimately, however, my aim in this paper will not be to argue for contextualist foundationalism's superiority to coherentism and traditional foundationalism, but to explore an application of it in the context of formal epistemology.

There are many ways of developing the idea of contextualist foundationalism, not all of which would be labeled 'contextualism' in contemporary parlance. The following three views bear a family resemblance to one another, and each, in my opinion, can be motivated by the considerations I will discuss in this essay. While I will generally focus my discussion on what I will call 'orthodox contextualist foundationalism,' this is for ease of exposition; most of what I say should be applicable to each of the views I am about to mention, *mutatis mutandis*.

Orthodox Contextualist Foundationalism:

Sentences of the form '*S*'s belief that *P* is justified' (as well as sentences attributing foundational justification more narrowly) express different propositions when uttered in different conversational contexts. Features of conversational contexts that affect which propositions such sentences express may include the presuppositions made by the participants to a conversation, the purposes of the conversation, and the practical situation faced by the participants in the conversation. In general, a proposition *P*'s being presupposed by the parties in a conversational context *C* tends to make sentences of the form '*S*'s belief that *P* is foundational,' true, when uttered in *C*.¹¹

Relativist Foundationalism:

Sentences of the form '*S*'s belief that *P* is justified' (as well as sentences attributing foundational justification more narrowly) express propositions that

⁹Again, see (author?) 1977 and (author?) 3.

¹⁰The emphasis is on 'distinctive' here. The contextualist will accept that, in any given context, there are regress stoppers. But the regress stoppers in one context may have little in common with the regress stoppers in another. There will not be some distinctive property—e.g., indubitability—shared by all regress stoppers in all contexts.

¹¹See (author?) 52, (author?) 13, (author?) 16, and (author?) 32. For the emphasis on presupposition in particular, see (author?) 4.

are not absolutely true, but only true relative to contexts of assessment. Which contexts of assessment we take on for the purpose of evaluating sentences about fundamentality (and so which such sentences we will be prepared to endorse) may depend on factors such as which presuppositions we and our interlocutors make, our purposes in evaluating the attribution, and our practical situation (along with, perhaps, the practical situation of the subject of the evaluation).¹²

Expressivist Foundationalism:

Sentences of the form ‘*S*’s belief that *P* is justified’ (as well as sentences attributing foundational justification more narrowly) do not express propositions of the usual sort (except perhaps in a deflationist sense), but instead are used to express our acceptance of epistemic norms. As a matter of anthropological fact, which epistemic norms we accept (and so which such sentences we will be prepared to endorse) varies with factors such as which presuppositions we and our interlocutors make, our purposes in evaluating attributions of fundamentality, and our practical situation.^{13,14}

I have characterized the above views as views concerning the truth conditions of sentences of the form ‘*S*’s belief that *P* is justified’ or ‘*S*’s belief that *P* is foundationally justified.’ While this might seem like a simple extension of familiar views about sentences of the form ‘*S* knows that *P*’, in fact things are not so simple; unlike ‘knows’, ‘justified’ and ‘foundational’, at least in the context of epistemology, are arguably technical terms.¹⁵ While it is easy to see how one might appeal to empirical linguistic data to adjudicate debates about the semantics of ‘knows’, it is harder to see how similar debates can even get off the ground concerning terms that do not occur in ordinary usage. My hope is that while I have expressed the views above as views about the truth conditions of a certain sort of sentence, it is not so hard to see how we can understand them as views about what it takes for beliefs to have certain sorts of epistemic statuses—statuses that *could* be attributed by sentences containing words like ‘justified’ and ‘foundational’, but which can play important theoretical roles even if we do not often directly attribute them. For instance, if I take some belief of yours to have one of the relevant statuses (e.g., to be justified), that will likely have implications for how I interact with you—perhaps I will be willing to take the belief for granted in the course of planning our joint endeavors, and will not object when you express the belief in assertion, or assert consequences of the belief—even if I never directly attribute that status. And even if we rarely explicitly attribute fundamentality (or lack thereof) to beliefs, we may manifest judgments about which beliefs are foundational in other ways. For instance, if I assert ‘Bob must be in

¹²E.g., I intend this to be in the spirit of (author?) 35.

¹³E.g., I intend this to be in the spirit of (author?) 10.

¹⁴It may also be the case that epistemic contrastivists such as (44; 45) can take advantage of some of the arguments I will discuss.

¹⁵(15) argues that ‘epistemic justification’ is a technical term, in a ways that contrasts with ‘knowledge’.

his office,' it is plausible that I am not only expressing my belief that Bob is in his office, but also conveying that I do *not* take this belief of mine to be foundationally justified.¹⁶

There are many important differences (and perhaps some not-so-important ones as well) between orthodox contextualism, relativism, and expressivism, and these differences have been much debated outside the context of their relevance to foundationalism.¹⁷ While the debates have usually concerned knowledge attributions rather than attributions of foundationality; the issues are similar. In fact, given certain additional premises, contextualism, relativism, and expressivism about knowledge *just are* versions of contextualism, relativism, and expressivism about foundations. E.g., (60) is naturally interpreted as a foundationalist who holds that a belief that P is foundational for a subject S just in case S knows that P . If we accept this view that knowledge plays this foundational role, then contextualism about knowledge is contextualism about foundations. While Williamson himself has argued against contextualism ((**year?**)), the suggestion that Williamsonian views about the centrality of knowledge in epistemology can be fruitfully combined with contextualism is not unheard of.¹⁸

My strategy in this paper, then, will be to try to remain as neutral as possible on the issues that separate different versions of 'shifty' epistemological views. My hope and expectation is that these differences will not substantially affect how, as a class, these views interact with foundationalism. In particular, I expect that each of these views promises to offer some similar advantages over traditional versions of foundationalism, whatever other drawbacks they may have. For this reason, while I will focus on arguing that orthodox contextualist foundationalism can address certain problems that non-contextualist versions of foundationalism cannot, I expect that similar arguments could be constructed for the other two views listed above.

So far I have briefly motivated the idea that some version of contextualist foundationalism is a natural response to the infeasibility of the Cartesian project. In the next section of this essay I will argue that contextualist foundationalism promises to bear some surprising fruit; in particular, I will argue that it can resolve a persistent puzzle in Bayesian confirmation theory.

Before jumping into the details, however, it may help to start with a bird's eye view of the terrain. There's a certain sort of traditional foundationalism that fits very naturally with Bayesian views about rational belief updating. Very roughly, it goes like this. There is a domain to which we have a sort of privileged epistemological access—at least if we're paying attention, we always have it in our power to know the truth concerning matters in this domain. It constitutes, as (**author?**) (60, p.92) puts it, our

¹⁶I am gesturing towards the view that 'must' in English is an 'evidential', in that (one of) the roles of 'must' in sentences of the form 'it must be that P ' is for the speaker to indicate the nature of her evidence for her belief that P —in particular, to indicate that it is somehow indirect. While the view that 'must' is an evidential is common, it is controversial just what about a speaker's evidence she conveys by using 'must'. See (**author?**) 1 for a general discussion of the category of evidentials.

¹⁷(10) argues that expressivism can reap the benefits of contextualism without some of the main costs. (Though see (41) for a response to this style of argument.) (19) expresses sympathy for both expressivism and relativism in epistemology, and is doubtful that, properly understood, the views ultimately conflict with one another.

¹⁸See (**author?**) 26.

‘cognitive home,’ where nothing is hidden to view. Because we clearly *can* be mistaken about almost everything, even when doing our best to be careful, this domain will have to concern our own minds. According to this view, while we *can* be mistaken about whether we are injured, or whether we are seeing a red apple, we cannot be mistaken (at least if we’re paying attention) about whether we are in pain, or whether we are having a visual experience as of seeing a red apple. We can think of our belief-forming processes as a journey—we start in our cognitive home (where we can know the truth) and then we venture outside, forming conjectures about further matters based on how things are in our home. E.g., if things are a certain way in our home (we have an experience as of a red apple in front of us), then they are *probably* a certain way outside our home (there probably *is* a red apple).¹⁹

This version of foundationalism view fits neatly with Bayesian accounts of belief updating because together, they suggest a natural answer to the question of how a rational agent should update her beliefs over time; she should strictly conditionalize on propositions concerning matters in her cognitive home. That is, Bayesianism provides a straightforward way of formalizing the idea that hypotheses concerning matters outside our home are credible only to the degree that they are supported by facts about our home. Namely, one’s credence in any proposition should equal the conditional probability of that proposition given all the facts about how things are in one’s home.’²⁰

So the view that we have cognitive homes and the view that strict conditionalization is the appropriate way for a rational agent to update her beliefs together form a neat package. But what if we lack cognitive homes? What if there is no domain in which careful reflection is guaranteed to lead to the truth? In the next section, I will argue that the neat fit between Bayesianism and foundationalism does not require that we have cognitive homes of the traditional sort. There need not be a domain about which we *cannot* be wrong, so much as, in each context, a domain about which it is taken for granted that we’re *not* wrong. We can think of this domain as a *kind* of cognitive home, so long as we recognize the differences between the cognitive homes we’ll be positing, and the more familiar ones. In some contexts, a great deal might count as part of our home (including facts concerning matters beyond our own minds), while in others, even

¹⁹The project of providing an inference to the best explanation style response to skepticism, in which skeptical hypotheses are rebutted by arguing that commonsense hypotheses provide better explanations of some class of mental phenomena—pursued by, among others, (5) and (**author?**) (53)—seems to me easiest to motivate in the context of the above traditional foundationalist picture. The question of whether, or to what degree, commonsense hypotheses provide a better explanation of facts about some class of mental phenomena than skeptical hypotheses is most obviously of epistemological import if we are straightforwardly in a position to know facts about the mental phenomena, and are justified in accepting further hypotheses only to the degree that they are supported by those mental facts (e.g., only to the degree that they can provide good explanations of those facts). Once this assumption is abandoned, it is less clear what the epistemological significance would be of the success or failure of the explanationist project. While (5) does explicitly defend coherentism and argue against foundationalism, it seems to me that the difficulties he encountered in that work in justifying the ‘doxastic presumption’, as well as the fact that he later reversed position and endorsed foundationalism (**(year?)**)—while still defending inference to the best explanation as a response to skepticism—reinforce the idea that explanationist responses to skepticism fit most neatly with foundationalism.

²⁰Thanks to an anonymous referee for help on how to put this point.

facts about whether or not we are in pain might fail to make it in. We have cognitive homes, but because they shift from context to context, we can think of them as mobile homes. These cognitive mobile homes are, I will argue, sufficient to save the package of Bayesianism and foundationalism.

2 Bayesianism and Defeat

A number of writers have argued that there is a tension between Bayesian accounts of belief updating, and the holist epistemological claim that all beliefs are defeasible.²¹ In this section, I will first introduce the topic of defeat, and will try to show why, *prima facie*, Bayesian approaches to belief updating promise to deliver an attractive account of how defeat works. I will then introduce an apparent limitation of such approaches; if we grant the holist claim that all beliefs are defeasible, then there will be cases of defeat that extant Bayesian approaches seem to lack the resources to model. At best, this looks like a gap waiting to be filled with some new extension of the Bayesian framework. At worst, it looks like evidence that the whole approach is misguided. I will argue that we can solve the problem without any novel formal machinery once we view Bayesian theories of belief updating as species of foundationalism, and recognize that this leads to the possibility of a contextualist Bayesian foundationalism. The argument that there *is* a difficult problem for the Bayesian depends, I will argue, on the assumption that the Bayesian is a non-contextualist foundationalist. If the Bayesian is a contextualist foundationalist, however, the problem has a very natural solution.

2.1 Defeat

Suppose you read in a newspaper that the majority of Salemites intend to vote ‘Yes’ on proposition 17. As a result, you come to believe that proposition 17 will pass. Later, however, you read that the ‘No’ movement is about to release an advertising campaign that, based on testing in focus groups, looks to be extremely effective at changing people’s minds.

After you read the first story but before you read the second, you have justification to believe that proposition 17 will pass. But after you read the second story, you no longer have such justification. Epistemologists say that reading the second story ‘defeats’ the support that you had for the claim that proposition 17 will pass.

Can we say anything more precise about defeat? Suppose E is our initial evidence, H is the hypothesis that evidence supports, and E' is our defeater. A natural thought is that in cases of defeat, while E on its own is good evidence for H , $E \& E'$ is not.

Moreover, the probability calculus suggests a natural formalization of these relations. After all, probabilistic confirmation is non-monotonic; E can confirm H , even though $E \& E'$ does *not* confirm H .²² Because of the non-monotonicity of confirmation, we

²¹See (author?) 11, (author?) 55 and (author?) 42.

²²This is not intended to be a *requirement* for defeat—e.g., E' might defeat E ’s support for H without doing so *fully*, so $E \& E'$ might still provide some confirmation for H , even though E' is a defeater for

can treat cases of defeat as ones in which an initial body of evidence supports some hypothesis, but a larger body of evidence—one that includes a defeater in addition to the initial evidence—does not. Applied to the example just presented, the strategy might go as follows.

Intention = The majority of Salemites intend to vote ‘Yes’ on proposition 17.

Pass = Proposition 17 will pass.

AdCampaign = A persuasive ad campaign attacking the proposition is about to be released.

Intention is evidence for *Pass*. We can represent this probabilistically by saying that $P(\textit{Pass} \mid \textit{Intention}) > P(\textit{Pass})$. The second story is a defeater for the claim that the proposition will pass. If we assume that the defeater brings us back to roughly the level of support we had initially for *Pass*, we can represent this as follows: $P(\textit{Pass} \mid \textit{Intention} \ \& \ \textit{AdCampaign}) \approx P(\textit{Pass})$.

While this is a promising beginning of a Bayesian account of defeat, a full story would have to say much more.²³ The challenge to Bayesian accounts of defeat that I will discuss, however, challenges any probabilistic account of defeat that turns on the non-monotonicity of probabilistic confirmation, as the proto-account I have sketched so far does.

2.2 The Ubiquity of Defeasibility

The (putative) difficulty for Bayesian accounts of defeat is the following: just as one can undermine support for a hypothesis by attacking the link between one’s evidence and the hypothesis, one can also undermine support for a hypothesis by attacking the evidence itself. So not all defeat can be understood as involving the accumulation of evidence, where the final, total body of evidence fails to support some hypothesis that was supported by the initial, smaller body of evidence. Arguments of this general form have been given by David (11), Jonathan (55), and James (42). Exactly what this amounts to, and why it is been thought to be a problem for the Bayesian, can be brought out with an elaboration of our earlier example.

As before, you read that the majority of Salemites intend to vote ‘Yes.’ But now, the second story does not concern an ad campaign, but instead alleges that the polls on which the first article was based come from a source with a flawed methodology—their results are not likely to be representative of what the people of Salem, on the whole, intend.

Intuitively, this defeats your support for the hypothesis that proposition 17 will pass. Why does this present a *prima facie* problem for the Bayesian? Let *Intention* and *Pass* have the same meanings as before, and let *FlawedSource* be understood as follows:

E’s support for *H*. Thanks to an anonymous referee for pointing this out.

²³(author?) (Ms.) does an admirably thorough job of displaying the resources the Bayesian has in describing a wide variety of cases of defeat.

FlawedSource = A second story reported that the polls relied on by the first story come from a source with a flawed methodology.

What we'd *like* to say would be that the following two conditions hold:

1. $P(\textit{Pass} \mid \textit{Intention}) > P(\textit{Pass})$
2. $P(\textit{Pass} \mid \textit{Intention} \ \& \ \textit{FlawedSource}) = P(\textit{Pass})$

This would be to treat *FlawedSource* as playing the same role that *AdCampaign* did in our earlier case. The problem is that the second condition *does not* hold. If the majority of Salemites intend to vote 'Yes', then regardless of how the polls were conducted, the proposition will probably pass. That is, $P(\textit{Pass} \mid \textit{Intention} \ \& \ \textit{FlawedSource}) > P(\textit{Pass})$.

The basic problem is that if we treat the evidence provided by the first story as *Intention*, then that evidence is not defeated by the story about the flawed polling methodology. Rather, if we are to accommodate the possibility of defeat via that story, it looks like we need to treat the evidence provided by the first story as something like *Polls*:

Polls = A poll reported that the majority of Salemites intend to vote 'Yes.'

Only if something like *Polls* was the evidence that the first story provided can we explain why the story about the flawed methodology defeats your support for thinking that the proposition will pass.²⁴

So far this might not seem like a problem at all. What's wrong with treating the evidence provided by the first story as *Polls*? The danger is that just as *Intention*

²⁴One might think that we could stick with treating your evidence as *Intention* rather than as *Polls* so long as we treat you as Jeffrey conditionalizing on the partition [*Intention*, \sim *Intention*], rather than strictly conditionalizing on *Intention*. (11) and (55) convincingly argue that this will not work. While discussing the details is beyond the scope of this paper, the basic problem is as follows. In order to allow that *FlawedSource* undermines the support that *Intention* provides for *Pass*, *FlawedSource* and *Pass* must start out probabilistically independent of one another, but must become probabilistically dependent (in particular, they must become negatively relevant to one another) after Jeffrey conditionalizing on the partition [*Intention*, \sim *Intention*]. But Jeffrey conditionalization cannot induce this sort of probabilistic dependence. More intuitively, conditionalizing on *Intention*—whether strictly, or Jeffrey-style—does not give us a way to represent that the support for *Intention* depends on beliefs about the polls, and so does not give us a way to represent that the support for intention can be defeated by information about polls. To be sure, one *can* deal with the problem if one allows oneself more flexibility in choosing an input partition—see (author?) 54. E.g., we might hold that when we read the newspaper, we do not simply Jeffrey Conditionalize on [*Intention*, \sim *Intention*], but perhaps instead use some more complex partition, including *FlawedSource*, *Pass*, and perhaps many other propositions. Weisberg is unhappy with this sort of response, however: 'Of course we can always set up the input partition to get the results we want, but then our input rule is doing all the work, and the inputs problem becomes the only interesting problem on the table.' (2009, p.805) Ultimately, I think both Wagner and Weisberg are onto something. Wagner is right that, given the right inputs to Jeffrey conditionalization (or, I'd point out, even strict conditionalization), these cases do not pose a problem for the Bayesian. Weisberg is right that, without some constraints on what count as appropriate inputs, this response only saves the Bayesian approach to updating at the cost of robbing it of explanatory power. Ultimately, I think this is where contextualism helps the Bayesian, in ways that should become clear in the latter part of this section.

can be undermined, so too can *Polls*. For instance, suppose the second story is just a correction—it reports that the first story misreported the poll results. If this defeater is to be accommodated, the evidence will need to concern, not what the polls said, but what the newspaper said the polls said.

Perhaps, at some point, we'll get to a way of characterizing the evidence provided by the first story such that it is immune to the sorts of defeat we've been discussing. Maybe the evidence will concern, e.g., the sorts of experiences I have upon reading it. But to many philosophers, especially those sympathetic to broadly holist epistemological positions, this will seem quaintly optimistic. Even claims about our experiences are not immune to reasonable doubt, and can be defeated in the right circumstances.²⁵ To hold out hope that there is some class of beliefs that is immune to rational undermining, and is such that all episodes of learning can be understood as involving the acquisition of some new beliefs in the privileged class (with further changes in our body of beliefs coming via conditionalization) is to hope for something very much like the traditional, 'cognitive home' version of foundationalism discussed at the end of the previous section. To be clear, there are certainly contemporary philosophers who hold such views—perhaps (47; 48) is the most explicit recent defender.²⁶ Moreover, there are projects taken seriously in philosophy that seem to me to presuppose such traditional foundationalist views.²⁷ I have no original arguments to offer in this paper against such views; still, I am interested in exploring the consequences for Bayesian accounts of belief updating once we grant their falsity.

To sum up, the challenge to Bayesian accounts of defeat is as follows. If the Bayesian is too generous about what our evidence is—e.g., if she thinks that our evidence in the cases I have been discussing is *Intention*—then she will not be able to account for certain cases of defeat. But there's no way to avoid being too generous; whatever the Bayesian says our evidence is will turn out to be immune to certain intuitive instances of undermining according to the Bayesian. But nothing is immune to undermining.

Whether this is rightly thought of as a *challenge* at all will depend on what sort of ambitions we have for Bayesian accounts of defeat.²⁸ One response to examples like the ones I have been discussing is simply to concede that they are not cases to which extant

²⁵See, e.g., (author?) 46 for arguments to the effect that we are highly fallible introspectors. See also (author?) 60, especially the discussion of our 'cognitive homelessness' in chapter 4, for more general arguments to the effect that there is no realm of facts which could play the foundational role that we might have hoped facts about experiences would play.

²⁶For Smithies' response to the Schwitzgebelian arguments mentioned in the previous footnote, see (author?) 49.

²⁷In an earlier footnote I mentioned that the project of responding to skepticism by appeal to inference to the best explanation seems to me one such example.

²⁸Lydia (37) argues that the seriousness of the problem I have been discussing, 'depends entirely on what a given philosopher wanted from [Jeffrey Conditionalization] in the first place.' While she grants that certain ambitious hopes are frustrated by the issues we've been discussing, she argues that this need not rob Jeffrey Conditionalization of its interest: 'if, on the other hand, one regarded the [Jeffrey Conditionalization] formula only as a useful epistemic tool for particular propositional updates, then the fact that a great many salient and interesting propositions can be updated using an obvious partition and Jeffrey's rule will be quite sufficient.' (p.581)

Bayesian approaches are meant to apply. While this is fine as far as it goes, it raises the natural question of how we might supplement Bayesian approaches to account for *all* possible cases of defeat, rather than just some.²⁹ To be sure, there are substantial limitations of the Bayesian framework, and if we are to insist on a Bayesian approach to defeat, let alone one that is supposed to handle all possible cases of defeat, it should not be on the basis of the view that *all* phenomena of epistemological interest can be captured in the Bayesian framework.³⁰

Still, there are reasons to be unhappy with the idea that an alternative formalism, or even an extension of the Bayesian approach we've been discussing so far, is what's called for here. First, alternative formalisms available so far run into similar troubles in accounting for undermining defeat.³¹ Second, the Bayesian account of defeat handles some cases—e.g., the first version of the newspaper case—extremely nicely, and these cases do not seem all that different from the cases where the Bayesian account seems to falter. If we rest content with the Bayesian treatment of the original case involving a story about polls and a story about an ad campaign, but look for some alternative approach to handle the variations on the original case, we risk ending up with an account of defeat that looks oddly disjunctive. At least to me, it seems that the sort of defeat provided by the story about the ad campaign is not all that different from the sort of defeat provided by the story about the flawed polling methodology, or the error in reporting the poll results. It would be nice to be able to treat each of these cases as having a similar formal structure, so if we're going to accept a Bayesian account of the first case, (which is hard to resist) it would be nice to have a Bayesian account of the others as well. In the next section, I will argue that such an account can easily be had, so long as our Bayesianism takes a contextualist form.

2.3 Bayesianism as (Contextualist) Foundationalism

It is instructive to think of Bayesianism as a species of foundationalism, in which a subject's prior probability function, as well as her beliefs in the evidence propositions she updates on, both have a sort of foundational status. Just as non-Bayesian foundationalists hold that a subject's beliefs are justified just in case they derive support from the subject's foundational beliefs, Bayesians hold that a subject's credences are justified just in case they receive the right sort of support from the foundational elements in the Bayesian epistemological framework—i.e., a subject's credences are justified just in case they result from conditionalizing the subject's prior probability function on the subject's evidence propositions.

²⁹(11) seems hopeful that some extension of Bayesianism will prove adequate. (42) considers some strategies for handling defeat in a version of the Bayesian framework, but also takes seriously the idea that significant departures from that framework might be required to account for defeat.

³⁰To take just one major example, since Bayesianism incorporates an assumption of logical omniscience, it cannot represent logical learning. While there have been attempts to model special cases of logical learning in a Bayesian framework (e.g., (author?) 1988), I think it is fair to say that they are not generally regarded as likely to provide an attractive general treatment of logical learning in a Bayesian framework.

³¹See (author?) 56.

One reason it is helpful to think of Bayesianism as a version of foundationalism is that it makes clear that the Bayesian faces versions of the familiar tricky question for non-contextualist foundationalists. Rather than the generic ‘what criteria must beliefs meet in order to be foundational?’ we can pose two Bayesian-specific versions of this question:

1. What criteria must prior probability functions meet in order to be rationally permissible?
2. What criteria must evidence propositions meet in order to be rationally updated upon?

The first question—referred to in the Bayesian literature as the ‘problem of the priors’—is well known, though it has no generally agreed upon solution. Many historically popular answers to the first question, moreover, can be seen as paralleling historically popular answers to the generic question for foundationalists, and as sharing the difficulties of those answers. For example, **(author?)**’s project of determining a prior probability function on purely logical grounds shares the appeal, and the infeasibility, of the traditional Cartesian foundationalist project.³² And the orthodox subjective Bayesian view that *all* coherent prior probability functions are rationally permissible bears a family resemblance to Harman’s view that all beliefs are foundational, and faces similar objections. From a bird’s eye point of view, both views have the same structure—they say that rationality is a matter of reasoning appropriately *given* one’s doxastic starting points. While they construe the starting points differently—the subjective Bayesian thinks of them as constituted by a subjective probability function, while (23) doubts that we have anything like degrees of belief in the sense countenanced by Bayesians—they both entail that what one ultimately ought to believe strongly depends on what one starts out believing, in ways that can seem implausible.³³

The second question—the question about which propositions subjects can rationally update on—is much less often discussed by Bayesians.³⁴ As the discussion of defeat

³²While there are many reasons to doubt that the Carnapian project can be completed, Nelson Goodman’s **(year?)** objections are perhaps the best known.

³³See **(author?)** 27. While they focus on objecting to the method of reflective equilibrium, they show how the considerations they raise apply to Harman’s general foundationalism. Moreover, the objections they raise can be straightforwardly translated into objections to subjective Bayesianism. E.g., somebody who starts out with a prior probability of 0 that humans will ever walk on the moon, and retains this degree of belief despite having access to the same evidence the rest of us have, seems *irrational*, even if their body of belief as a whole is probabilistically coherent.

³⁴Deborah **(author?)** (36, pp.86-8) identifies the question of when one should accept a piece of evidence—in the Bayesian framework, when one should conditionalize on an evidence claim—as a crucial one for Bayesians, on which they often have little to say.

One notable exception is Timothy (60). Like Mayo, he identifies the question as a key one for Bayesians. Rather than rejecting Bayesianism on the grounds that it cannot handle the question, however, he goes on to defend a version of Bayesianism on which a subject’s evidential probabilities are determined by conditionalizing a prior probability function on each of the propositions the subject knows.

Still, to the extent that this view does *not* give a broadly Bayesian treatment of knowledge, it can be seen as a sort of retreat, at least if one starts from the point of view of an orthodox Bayesian philosopher

above should make clear, however, it is not obviously any less difficult than the first; the arguments discussed above seem to suggest that *any* answer to the second question leads to the implausible conclusion that certain beliefs (namely, whichever ones we are entitled to update on) are immune to certain sorts of undermining.

Both of these questions for the Bayesian, however, can be given contextualist answers.³⁵ And while a contextualist treatment of the first question is a topic for another essay, in what follows I will argue that a contextualist treatment of the second question can let the Bayesian reconcile her treatment of defeat with a version of the ubiquity of defeasibility.³⁶

The basic strategy for saving Bayesian accounts of defeat if we adopt contextualism about what subjects are entitled to conditionalize on, is to argue as follows. First, admit that in a context in which a subject counts as entitled to update on E , E cannot be undermined. So far this looks like biting the bullet and accepting that certain beliefs are immune to undermining—each context will treat certain beliefs as indefeasible. However, the contextualist can account for a version of the ubiquity of defeasibility in the following way: even if we start in a context in which we count as entitled to conditionalize on some evidence E , and in which E thereby counts as immune to undermining, we can always shift to a new context in which we do *not* count as so entitled, and in which our support for E (if any at all) is treated as defeasible, and may be undermined when we learn some new proposition E' .³⁷ Moreover, it may be that we have a systematic tendency

of science. As Mayo notes, ‘the need for a supplementary account of evidence would belie one of the main selling points of the Bayesian approach—that it provides a single, unified account of scientific inference.’ (p. 88) However, if we’re not antecedently committed to Bayesianism answering *all* questions in epistemology and philosophy of science (as I am certainly not) this sort of retreat need not strike us as unwelcome, even if we find the Bayesian framework a generally useful one.

³⁵For some precedent, Jon (author?) (59, p.5) defends a version of Bayesianism on which ‘an agent’s evidence is relative to her current operating context.’ Though ultimately, Williamson’s view has less in common with the sort of contextualism I am discussing than it might seem. In particular, he holds that whatever an agent takes for granted is part of her evidence—there’s no room for us, as attributors of knowledge or justification, to hold that someone is treating something as part of her evidence, when it is in fact not. The sort of contextualism I will be discussing does not have this commitment.

³⁶While I know of no developed contextualist treatment of the problem of the priors, Timothy (60) seems sympathetic to such a position in his treatment of evidential probability:

The discussion will assume an initial probability distribution P . P does not represent actual or hypothetical credences. Rather, P measures something like the intrinsic plausibility of hypotheses prior to investigation; this notion of intrinsic plausibility can vary in extension between contexts. (p.211)

Moreover, if Williamson adopted contextualism about knowledge, then his treatment of evidential probability would amount a (broadly) contextualist answer to the second question, since (as mentioned in the previous footnote) he holds that subjects are rationally entitled to update on (and only on) what they know. Though as previously mentioned, Williamson rejects contextualism about knowledge ((year?)).

³⁷Ram (author?) 38, 39, 40, has defended contextualism about evidence, and the view I express here is heavily indebted to his work. To my knowledge, however, he has never combined his contextualism with Bayesian approaches to belief updating, or used it to address the apparent challenge to such puzzles presented by the ubiquity of defeasibility. The strategy I pursue in this section also bears some similarity to Lewis’ (1996) contextualist strategy for saving closure principles for knowledge from the threat of Kripke’s dogmatism paradox. Like Neta, though Lewis does not apply his discussion of contextualism

to shift contexts in such a way that claims that in an earlier context are treated as indefeasible are, once the question of their defeasibility is raised and taken seriously, treated as defeasible in a later context.³⁸

The strategy that I am suggesting for the Bayesian will need to have two components:

1. An account of how particular contexts do treat certain beliefs as indefeasible, and
2. An explanation of why no particular belief we consider ever seems to be indefeasible (and so, derivatively, an explanation of our temptation to endorse the claim that all beliefs are defeasible)

While developing these components in depth is beyond the scope of this paper, I hope the following will give the reader some idea of how the story might go, as applied to the cases of defeat discussed earlier in this section.

As for the first component of our strategy, the short answer is that in any context, certain claims will count as legitimate to update on, and for reasons we've already seen, what you update on will be immune to undermining, in the Bayesian framework. For instance, we can agree in the initial case that you are entitled to conditionalize on *Intention*, and that doing so provides support for *Pass*, but that this support is defeated when you learn *AdCampaign*. The initial, flatfooted Bayesian account of what's going on when you read one story about what the voters intend, and then another story about how those intentions are likely to change, can remain untouched.

What about the next case, where the second story concerns flawed polling methodology? If we started out discussing the earlier case, then bringing up this more complicated one is liable to induce a shift into a new context in which you no longer count as having been entitled to conditionalize on *Intention* in the first place. In this new context—one in which we are taking seriously possibilities in which the polls are unreliable—all you're entitled to conditionalize on is *Polls*—the claim that the polls reported that the majority of Salemites intend to vote 'Yes.' While this claim on its own supports *Pass*, together with *FlawedSource*—which we still take you to be entitled to conditionalize on—it does not.

The next step is similar. When we move to the case in which the second story reports that the first story misreported the poll results, we induce yet *another* context shift. Now that we do not regard you as having unproblematic access to what the polls said, we treat

to Bayesianism.

³⁸This may provide a wedge for various species of anti-foundationalism to sneak back in. Suppose our tendency when *E* is challenged is to retreat to a context in which *E* is not foundational, but *E'* is. And suppose our tendency when *E'* is challenged is to retreat to a context in which *E'* is not foundational, but *E''* is. Suppose that in the series of claims to which we are disposed to retreat, there are no repeats—i.e., no element in the series *E, E', E'', ...* is identical to any other. In that case, there is a sense in which infinitism—of the sort defended by (author?) 28, 29—would have turned out to be right after all, not as an account of the structure of justification *within* a context, but instead as an account of how we're disposed to shift *between* contexts. Alternatively, if there *were* repeats in the series, then a certain sort of coherentism might be vindicated. I will not pursue these questions here. I am more interested in exploring the consequences of the view under discussion than in deciding whether it should be thought of as a version of foundationalism or not.

you as not entitled to conditionalize on claims about what the polls reported. Rather, we'll think you're only entitled to conditionalize on claims about what the newspaper *said* the polls said, and this claim will not support the hypothesis that the proposition will pass when combined with the evidence you get from the second story.

Whenever we regard a claim as having been defeated, we'll give the same formal account; we'll regard somebody as having had an initial body of evidence E that supported some hypothesis H , but as having a final body of evidence E' —a superset of E —that no longer supports E . That is, defeat will always be explained as involving only evidence *accretion*, and never evidence *deletion*. The role that context plays, however, is in determining what gets to count as the initial evidence; rather than giving a context-independent, univocal answer to the question of what a subject's evidence is, we allow that our answer to questions about what a subject's evidence is will change, as our context changes (e.g., we may initially regard you as entitled to update on claims about what the majority of Salemites intend, but we also recognize that if our context shifts, we may no longer regard you as so entitled).

Just what induces these context shifts? This question is closely related to the demand for the second component of our strategy—the demand for an explanation of why no particular belief we consider ever seems to be indefeasible. The answer will depend on the particular form our contextualism takes, and things will of course look different if we instead opt for some version of relativism or expressivism. But certain basic factors may be common to each approach in the family—perhaps subjects will generally not count as entitled to update on some proposition when there are salient possibilities (salient to whom? that may differ depending on the approach) in which they are in error as to the truth of that proposition. Most of the time we're ignoring lots of error possibilities, and we'll treat subjects as entitled to conditionalize on claims about their external environment. But in special cases where certain error possibilities are raised to salience, we'll retreat, and treat subjects as entitled to conditionalize only on other claims—perhaps claims about their experiences, or claims about their inclinations to believe claims about their experiences, or something else. And there is not any principled stopping point to this process of retreat—there's no sort of evidence statement about which no doubts could be raised, though in practice we're usually not inclined to retreat very far.³⁹

We can use these observations to provide the basis for a general story about why no beliefs seem to be indefeasible; raising the question as to whether some belief is defeasible essentially amounts to an invitation to consider possibilities in which the belief might be false, so once the indefeasibility of a belief is in question, the context will typically shift so as to ensure that the belief is defeasible. That is, once we are considering possibilities in which we are in error as to whether E , it will seem irresponsible to simply take E for granted; rather, we'll be inclined to accept E only if it receives adequate support from other claims we *do* take for granted. But this is just what it is for us to treat E as enjoying merely defeasible support. Indefeasibility is 'elusive', in the sense of (32);

³⁹Moreover, the 'retreat' need not always be to weaker and weaker propositions. Rather, as the Neurath's raft metaphor suggests, it might simply be to *different* propositions.

raising the question of whether some claim E enjoys indefeasible support will typically bring us into a new context in which E enjoys merely defeasible support, even if it initially enjoyed indefeasible support in our old context.⁴⁰

The apparent difficulty for Bayesian accounts of defeat required assuming that the Bayesian must give a single, context-independent answer to the question of what you should conditionalize on when you read a particular newspaper story. Whatever answer she gives, as we saw before, there will be cases she cannot handle. But now we see that the Bayesian can reject the demand to tell us, once and for all, what you should conditionalize on in such a case. In most contexts, where we're not worried about certain error possibilities, there's no obstacle to allowing that you're entitled to conditionalize on the content of what the article reports (rather than just, e.g., the claim *that* it reported what it did). But other contexts will demand other treatments, and we need not allow that there's some privileged, stable context that we could never be pushed out of—we need not allow that there's any answer to the question of what you should conditionalize on that we could not come to reject after yet another context shift.⁴¹ The resulting position can be summed up, more or less, by the following remark from Wittgenstein's *On Certainty*: 'the same proposition may get treated at one time as something to test by experience, at another as a rule of testing.' ((year?), 98)

In the absence of more general motivations for contextualist foundationalism, this reply on behalf of the Bayesian might seem ad hoc. But once we admit both that (a) contextualist foundationalism is a viable contender in debates about the structure of epistemic justification, and that (b) Bayesianism is in effect a species of foundationalism in which the beliefs that subjects are entitled to update on play a foundational role, it is hard to avoid the conclusion that (c) we must take seriously contextualist Bayesian foundationalism, according to which the beliefs that subjects are entitled to update on play a foundational role, but which sorts of beliefs get to play that role varies with

⁴⁰I should stress that I don't assume that we *always* retreat once a challenge is raised. Some challenges are not taken seriously; attempts to change the context can be resisted. See, e.g., (4).

⁴¹Earlier in this section I suggested that it would be a defect of a treatment of defeat if it were 'oddly disjunctive.' But one might worry that the sort of account I have suggested on behalf of the Bayesian is just that. We might put the worry as follows: on the story I have been sketching, there are really two quite different types of defeat. The first is the simple sort of defeat that does not involve anything like context shifting. The second is the more complicated sort, which requires a broadly contextualist explanation.

This charge, however, can be resisted. The sense in which the contextualist can deny that she countenances two forms of update (or two associated forms of defeat) parallels the contextualist strategy for affirming 'intellectualism'—the thesis that what one knows depends only on 'truth-conducive factors.' Roughly, while the contextualist will accept different sentences involving epistemological vocabulary (whether those are knowledge attributions, or claims about what subjects are entitled to update on) as factors other than truth-conducive factors change, (perhaps as the practical stakes change, or as she starts considering new error possibilities) in no context will she affirm that some subject would (or would not) know/justifiably believe some proposition if only the non-truth-conducive factors were different (e.g., if only the practical stakes were different, or if only we were considering fewer error possibilities). See (author?) 51, pp. 2-3 and (author?) 17, pp. 24-5, both of whom accept that the sensitive invariantist must deny intellectualism, while the contextualist may accept it. Moreover, I take it that this feature of contextualism is shared by expressivism and relativism, for essentially the same reasons.

context. And once we take that view seriously, we see that it has the considerable virtue of providing the foundations for a general account of defeat.

Still, it is worth acknowledging that the sort of contextualism we must appeal to in order to employ the strategy I have been discussing is, in important ways, a departure from some familiar forms of contextualism. (author?) (50, pp.102-105) distinguishes between *superficial* and *deep* contextualism. Roughly, the sort of context-dependence posited by the superficial contextualist is eliminable, while the sort of context-dependence posited by the deep contextualist is not. The deep contextualist is contextualist ‘all the way down,’ while the superficial contextualist thinks that, once you start digging, you eventually hit some context-independent epistemological bedrock. Without worrying too much about exactly how far ‘all the way down’ is, we can compare versions of contextualism for ‘depth’ depending on just how many of our epistemological notions they take to be context sensitive. For concreteness, it will help to consider some examples.

Consider the following simple form of contextualism about justified belief. A subject’s belief that P is properly called ‘justified’ in a given context only if the epistemic probability of P for the subject is above a certain threshold. Which threshold? That depends on the context. The epistemic probability of P , however, does *not* depend on context—context only comes into the picture in determining how probable P must be in order for a subject’s belief that P to merit the label ‘justified.’ The view defended by (author?) 14 is not *too* far from this.⁴² This is a relatively superficial version of contextualism in that it takes there to be a central epistemological notion—epistemic probability—for which contextualism does not hold. Moreover, this sort of superficial contextualism clearly cannot avail itself of the strategy I have been suggesting, since the strategy I have been suggesting requires that epistemic probability *itself* is context-sensitive; in contexts where E counts as evidence, we may strictly conditionalize on it, so it will have maximal probability. In other contexts, it will not. So the sort of contextualism Bayesianism requires to deal with the ubiquity of defeasibility goes deeper than at least one influential and familiar species of epistemological contextualism.

David (32) defends another version of contextualism, which also turns out to be insufficiently deep for the Bayesian’s needs. According to Lewis, knowing that P requires eliminating all contextually relevant counterpossibilities to P . His view is contextualist because the set of relevant counterpossibilities varies with context. But what does the eliminating—what one’s evidence is—does not.⁴³ On Lewis’ view, one’s evidence is always given by the content of one’s experience and one’s memory ((year?), p. 425). If we put things in Bayesian terms, this amounts to the claim that one is always entitled to update on the content of one’s experience and memory, no more and no less—and Lewis is no contextualist about the contents of experience and memory. For my purposes, this

⁴²Cohen advances a contextualist version of the relevant alternatives theory of knowledge; on his view epistemic probability is taken as fixed independently of context, and how probable an alternative has to be before it counts as relevant is determined by context. Also, while they’re sensitive invariantists rather than contextualists, (18) employ a similar strategy on which epistemic probability is stakes-invariant, but the cutoffs for knowledge and justification are not.

⁴³(author?) (50, pp.102-5) gives further reasons why Lewis’ view is best thought of a superficial form of contextualism.

is a crucial juncture at which Lewis' contextualism takes a superficial rather than a deep form, crucial because the Lewisian contextualist cannot appeal to my strategy for reconciling Bayesian accounts of updating with the apparent ubiquity of defeasibility; Lewis' view straightforwardly implies that our knowledge of the contents of our experience and memory is indefeasible in every context, so if we want to reject that assumption we'll need to look elsewhere.

Most writers who've defended versions of contextualism have only officially committed themselves to relatively superficial forms contextualism, since they've only defended contextualism about one or another epistemological notion. Moreover, insofar as they invoke some context *insensitive* epistemological machinery in explaining their position—e.g., a context insensitive notion of epistemic probability—they're committed to rejecting deep forms of contextualism.⁴⁴ Still, while deeper forms of contextualism are less familiar in epistemology than their more superficial cousins, I think they deserve to be taken seriously.⁴⁵ Where relatively superficial forms of contextualism are unable to reconcile Bayesian accounts of defeat with the ubiquity of defeasibility, deeper forms can.

3 Conclusion

The idea that contextualism can provide a sort of middle way in the debate between coherentism on the one hand, and traditional foundationalism on the other, is not new. But with the exception of some discussion of traditional skeptical problems, there has been relatively little work done exploring how *other* epistemological debates look once we take this middle way.⁴⁶

While I have focused on questions about defeat and belief updating, I hope to also have given the flavor of a general strategy we might use to explore further potential applications of contextualist foundationalism. First, take some view that might seem inextricably linked to traditional versions of foundationalism—versions that require a fixed cognitive home. Then, explore whether the idea can be defended while appealing

⁴⁴As an aside, I am inclined to think that many versions of superficial contextualism are not particularly interesting qua epistemological theses (though they may be interesting qua semantic theses); to the extent that the contextualist allows a realm of 'pure,' context-insensitive epistemology, epistemologists can largely ignore context and focus their investigations on topics to which it is not relevant. Moreover, the existence of a set of context-insensitive epistemological questions limits the scope of the epistemological work that can be done by contextualist theses. For instance, if we try to solve skeptical problems by appeal to some version of superficial contextualism, such problems will likely reappear concerning epistemological notions where we concede that contextualism will not help. If we are contextualists about knowledge but not epistemic probability, then even if our contextualism helps us explain why utterances of 'I know that I have hands' might be true in ordinary contexts, it will not help us explain why the epistemic probability that I have hands is any higher than the epistemic probability that I am a handless brain in a vat.

⁴⁵See (author?) 50 for further motivation.

⁴⁶Much of Michael Williams' work involves using a sort of contextualist foundationalism to argue that global skeptical challenges are somehow illegitimate. He has not, to my knowledge, discussed the relationship between contextualist foundationalism and Bayesianism. Ram Neta's work ((year?); (year?); (year?)) also involves something like contextualism about which beliefs are foundational, but to my knowledge he has not applied his version of contextualism to the topic of defeat.

only to mobile homes. The idea that rational updating proceeds by strict conditionalization is one example where I have argued this strategy is fruitful. But I suspect there are others. For instance, there has been much recent work on the relationship between binary epistemological notions (belief, knowledge) on the one hand, and graded ones (credence, epistemic probability) on the other.⁴⁷ Many of the writers working on this topic presuppose that a certain flatfooted account of the relationship—on which belief requires credence 1, and knowledge requires epistemic probability 1—cannot be right, since they assumed that few if any propositions are believed to degree 1, or have maximal epistemic probability. They take the flatfooted view to be incompatible with the post-Cartesian realization that *all* claims are, at least to some degree, doubtful. But the perspective of the present paper suggests a different way of accommodating that idea. While it may be that no propositions have maximal probability in *every* context, it can still be that many propositions have maximal probability in *some* contexts. Once we're happy thinking of graded notions like credence and probability as themselves context-sensitive, the obstacle to the flatfooted approach to the relationship between binary and graded epistemological notions may very well dissolve.⁴⁸

G.K. (author?) (1910) famously claimed that Christianity had not been tried and found wanting, but instead had never been truly tried. Similar claims have been made in defense of communism, the war in Afghanistan, the Reagan revolution, and various other ideas that have, in fact, been tried. Nevertheless, this is the plea I offer on behalf of contextualist foundationalism. Because it promises to illuminate debates about Bayesianism, defeat, and perhaps many other topics as well, before we conclude that it is an idea whose time has come and gone we should give it a real try.

Bibliography

- [1] Aikhenvald, Alexandra, 2004. *Evidentiality*. Oxford University Press.
- [2] Alston, William P. 1976. "Has Foundationalism Been Refuted?" *Philosophical Studies* 29:295.
- [3] Annis, David B. 1978. "A Contextualist Theory of Epistemic Justification." *American Philosophical Quarterly* 15:213–219.
- [4] Blome-Tillmann, Michael. 2009. "Knowledge and Presuppositions." *Mind* 118:241–294.
- [5] Bonjour, Laurence, 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.

⁴⁷See, e.g., (author?) 7, (author?) 43, (author?) 31.

⁴⁸See (author?) 12 and (author?) Forthcoming for arguments along these lines. Though it is not quite clear to me whether Clarke's view should be thought of as a species of contextualism, rather than a species of interest-relative or sensitive invariantism, in the sense defended by (author?) 25, (author?) 51, and (author?) 18.

- [6] ———. 1999. “Foundationalism and the External World.” *Philosophical Perspectives* 13:229–249.
- [7] Buchak, Lara. 2013. “Belief, Credence, and Norms.” *Philosophical Studies* pages 1–27.
- [8] Carnap, Rudolf, 1967. *The Logical Structure of the World [and] Pseudoproblems in Philosophy*. London, Routledge K. Paul.
- [9] Chesteron, G.K., 2007(1910). *What’s Wrong With the World*. Cossimo Classics.
- [10] Chrisman, Matthew. 2007. “From Epistemic Contextualism to Epistemic Expressivism.” *Philosophical Studies* 135:225–254.
- [11] Christensen, David. 1992. “Confirmational Holism and Bayesian Epistemology.” *Philosophy of Science* 59:540–557.
- [12] Clarke, Roger. 2013. “Belief Is Credence One (In Context).” *Philosophers’ Imprint* 13:1–18.
- [13] Cohen, Stewart. 1987. “Knowledge, Context, and Social Standards.” *Synthese* 73:3–26.
- [14] ———. 1988. “How to be a Fallibilist.” *Philosophical Perspectives* 2:91–123.
- [15] ———. 1995. “Is There an Issue About Justified Belief?” *Philosophical Topics* 23:113–127.
- [16] DeRose, Keith. 1995. “Solving the Skeptical Problem.” *The Philosophical Review* 104:1–52.
- [17] ———, 2009. *The Case for Contextualism*. Oxford University Press.
- [18] Fantl, Jeremy, and Matthew McGrath, 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- [19] Field, Hartry. 2009. “Epistemology Without Metaphysics.” *Philosophical Studies* 143.
- [20] Garber, Daniel, 1982. “Old Evidence and Logical Omniscience in Bayesian Confirmation Theory.” In Earman, John, editor, “Testing Scientific Theories,” University of Minnesota Press.
- [21] Goodman, Nelson, 1983. *Fact, Fiction, and Forecast*. Harvard University Press.
- [22] Greco, Daniel. Forthcoming. “How I Learned to Stop Worrying and Love Probability 1.” *Philosophical Perspectives* .
- [23] Harman, Gilbert, 1986. *Change in View: Principles of Reasoning*. Cambridge, MA: The MIT Press.

- [24] ———. 2001. “General Foundations Versus Rational Insight.” *Philosophy and Phenomenological Research* 63:657–663.
- [25] Hawthorne, John, 2004. *Knowledge and Lotteries*. Oxford University Press.
- [26] Ichikawa, Jonathan Jenkins. 2013. “Basic Knowledge and Contextualist “E = K”.” *Thought: A Journal of Philosophy* 2:282–292.
- [27] Kelly, Thomas, and Sarah McGrath. 2010. “Is Reflective Equilibrium Enough?” *Philosophical Perspectives* 24:325–359.
- [28] Klein, Peter. 1998. “Foundationalism and the Infinite Regress of Reasons.” *Philosophy and Phenomenological Research* 58(4):919–925.
- [29] Klein, Peter D. 1999. “Human Knowledge and the Infinite Regress of Reasons.” *Nous* 33:297–325.
- [30] Kotzen, Matthew. Unpublished Manuscript. “A Formal Account of Evidential Defeat.” .
- [31] Leitgeb, Hannes. 2014. “The Stability Theory of Belief.” *Philosophical Review* 123:131–171.
- [32] Lewis, David. 1996. “Elusive Knowledge.” *Australasian Journal of Philosophy* 74:549–67. Reprinted in (author?) 33.
- [33] ———, 1999. *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
- [34] MacFarlane, John, 2011. “Epistemic Modals Are Assessment-Sensitive.” In Egan, Andy, and Brian Weatherson, editors, “Epistemic Modality,” Oxford University Press.
- [35] ———, 2011. “Relativism and Knowledge Attributions.” In “Routledge Companion to Epistemology,” Routledge.
- [36] Mayo, Deborah, 1996. *Error and the Growth of Experimental Knowledge*. University of Chicago Press.
- [37] McGrew, Lydia. 2013. “Jeffrey conditioning, rigidity, and the defeasible red jelly bean.” *Philosophical Studies* .
- [38] Neta, Ram. 2003. “Contextualism and the Problem of the External World.” *Philosophy and Phenomenological Research* 66:1–8211.
- [39] ———. 2004. “Perceptual Evidence and the New Dogmatism.” *Philosophical Studies* 119:199–214.
- [40] ———. 2005. “A Contextualist Solution to the Problem of Easy Knowledge.” *Grazer Philosophische Studien* 69:183–206.

- [41] Plunkett, David, and Tim Sundell. 2013. “Disagreement and the Semantics of Normative and Evaluative Terms.” *Philosophers’ Imprint* 13.
- [42] Pryor, James, 2013. “Problems for Credulism.” In “Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism,” Oxford University Press.
- [43] Ross, Jacob, and Mark Schroeder. 2014. “Belief, Credence, and Pragmatic Encroachment1.” *Philosophy and Phenomenological Research* 88:259–288.
- [44] Schaffer, Jonathan. 2004. “From Contextualism to Contrastivism.” *Philosophical Studies* 119:73–104.
- [45] ———. 2007. “Knowing the Answer.” *Philosophy and Phenomenological Research* 75:383–8211.
- [46] Schwitzgebel, Eric, 2011. *Perplexities of Consciousness*. MIT Press.
- [47] Smithies, Declan. 2012. “Moore’s Paradox and the Accessibility of Justification.” *Philosophy and Phenomenological Research* 85:273–300.
- [48] ———, 2012. “A Simple Theory of Introspection.” In Smithies, Declan, and Daniel Stoljar, editors, “Introspection and Consciousness,” Oxford University Press.
- [49] ———. 2013. “On the Unreliability of Introspection.” *Philosophical Studies* 165:1177–1186.
- [50] Stalnaker, Robert, 2008. *Our Knowledge of the Internal World*. Oxford: Oxford University Press.
- [51] Stanley, Jason, 2005. *Knowledge and Practical Interests*. Oxford University Press.
- [52] Stine, Gail. 1976. “Skepticism, Relevant Alternatives, and Deductive Closure.” *Philosophical Studies* .
- [53] Vogel, Jonathan. 1990. “Cartesian Skepticism and Inference to the Best Explanation.” *Journal of Philosophy* 87:658–666.
- [54] Wagner, Carl. 2013. “Is Conditioning Really Incompatible With Holism.” *Journal of Philosophical Logic* .
- [55] Weisberg, Jonathan. 2009. “Commutativity or Holism? A Dilemma for Conditionalizers.” *British Journal for the Philosophy of Science* 60:793–812.
- [56] ———. 2015. “Updating, Undermining, and Independence.” *British Journal for the Philosophy of Science* 66.
- [57] Williams, Michael, 1999—First edition published in 1977. *Groundless Belief*. Princeton University Press.

- [58] ———. 2013. “Skepticism, Evidence and Entitlement.” *Philosophy and Phenomenological Research* 87.
- [59] Williamson, Jon, 2010. *In Defense of Objective Bayesianism*. Oxford University Press.
- [60] Williamson, Timothy, 2000. *Knowledge and its Limits*. Oxford University Press.
- [61] ———. 2005. “Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge.” *Philosophical Quarterly* 55:213–235.
- [62] Wittgenstein, Ludwig, 1969. *On Certainty*. Basil Blackwell, Oxford.