The Value of Incoherence

Claire Field

I argue that level-incoherence is epistemically valuable in a specific set of epistemic environments: those in which it is easy to acquire justified false beliefs about normative requirements of epistemic rationality. I argue that in these environments level-incoherence is the rationally dominant strategy. Nevertheless, level-incoherent combinations exhibit a distinctive tension, and this tension has been thought by many to indicate that level-incoherence is always irrational. Although this idea has proved resilient, I argue that it is incorrect. I evaluate three candidate explanations for the distinctive tension exhibited by level-incoherent combinations, only one of which is the traditional view (which I call the 'Prohibition View') that epistemic level-incoherence is prohibited by epistemic rationality. I argue instead for the 'Inquiry View', according to which level-incoherence is not rationally criticisable but is a reason to undertake further inquiry.

Keywords: rationality; epistemic akrasia; coherence; epistemic attitudes.

1. Level-Coherence and Epistemic Rationality

Level-coherence is a relation between normative beliefs and first-order attitudes. In general, to be level-coherent is for one's first-order attitudes to be in line with one's normative beliefs. If I believe I ought to go to the gym, but do not intend to do so, I am level-incoherent. If I believe I ought not believe in the Loch Ness monster but find myself nevertheless believing that there is a monster in Loch Ness, I am level-incoherent. Level-incoherence is often thought to be a paradigmatic example of irrationality. This can be expressed using the following wide-scope principle:

Level-Coherence: $O(BO\varphi \rightarrow \varphi)$

Level-Coherence says that you ought to be such that if you believe that you are rationally required to φ, where φ-ing indicates some first-order attitude, you do in fact φ. Here I focus specifically on *epistemic* level-coherence. In the epistemic case, Level-Coherence says that one's normative beliefs about what is epistemically required should be in line with one's first-order doxastic attitudes. Specifically, rationality requires you to be such that if you believe that epistemic rationality requires you to φ, where φ-ing indicates some doxastic activity, you do in fact φ. Since the requirement is wide-scope, it is possible to comply with it in two ways: by refraining from believing that epistemic rationality requires you to φ, or by believing that epistemic rationality requires you to φ and also φ-ing.[i] False normative beliefs can be either general (e.g. "I ought to always believe what reliable testifiers tell me") or specific ("Epistemic rationality requires that on *this* occasion I believe *this* proposition"). Here, I am particularly interested in more general false normative beliefs, because of the wide-ranging consequences they have for our epistemic agency and rationality. However, what I say here applies equally to level-incoherence involving all kinds of normative beliefs.

Level-incoherent combinations of doxastic states are, intuitively, in tension with each other. This tension has often been thought to indicate necessary irrationality. For

example, the apparent tension involved in level-incoherence has prompted many to include a requirement to avoid level-incoherence among a larger package of coherence requirements on epistemic rationality (see Broome (2013), Lee (2022), Lord (2018), Kiesewetter (2017), Reisner (2013), Smithies (2012), Worsnip (2018)). This larger package of coherence requirements traditionally includes requirements to avoid logically inconsistent beliefs, requirements to avoid Moorean absurdity (assertions such as "P, but I do not believe that P") and in general to avoid combinations of attitudes that do not "match" (Reisner (2013)), "don't fit together right" (Worsnip (2021: 3)), or "give rise to a certain normative tension" that is relieved only when one of the attitudes is revised (Kolodny (2007: 229)). It is worth emphasising that one need not think that Level-Coherence is a distinctive structural requirement of rationality to think that rationality always demands compliance with Level-Coherence. Those who emphasise the primacy of substantive rationality (responsiveness to reasons or evidence) have also thought that responding to the demands of substantive rationality, in fact, always means having level-coherent doxastic attitudes (see Kiesewetter (2017) and Lord (2018)).[ii]

I argue here that this is not the only explanation for the apparent tension between level-incoherent attitudes. I suggest instead that there are at least three ways to explain this, only one of which is that level-incoherent attitude combinations are always prohibited by epistemic rationality. I then go on to argue that epistemic level-incoherence has particular value in environments in which it is particularly easy to acquire justified false beliefs about requirements of epistemic rationality and I suggest that it would be preferable if our theories of epistemic rationality could accommodate this. I offer the Inquiry View as an alternative way explaining the tension involved in level-incoherence that allows us to account for the epistemically valuable features of level-incoherence in a specific set of epistemic environments.

Epistemic rationality, as I understand it here, is the positive epistemic status associated with having a rationally intelligible epistemic perspective. Epistemic rationality allows us to make evaluations of an agent's epistemic perspective. Evaluations of epistemic rationality take into account a variety of factors relevant to the agent's epistemic perspective: her doxastic attitudes, her environment, how she responds to evidence, her reasoning patterns, and so on. Epistemic rationality is not interested in evaluating the truth or falsity of the agent's beliefs, or whether or not she has knowledge. However, a plausible theory of epistemic rationality will be such that when all goes well, epistemic rationality helps us achieve epistemic goods such as true belief, justification, and knowledge. These goods are nevertheless not guaranteed. Epistemic rationality is an evaluation that depends on the agent and her epistemic perspective, and the world can always intervene in strange and unpredictable ways and rob even the most epistemically rational agent of knowledge. That said, epistemic rationality should not *reliably* or *predictably* prevent agents from achieving these epistemic goods.

I argue here that epistemic level-incoherence is epistemically valuable in a specific set of epistemic environments: environments in which it is particularly easy to acquire a false justified normative belief about what is rationally required. I call these "normatively misleading epistemic environments". If this is right, then thinking that level-incoherence is always rationally prohibited implies that agents in normatively misleading environments will be reliably prevented from this epistemic value. The epistemology classroom is one example of a normatively misleading environment. In the epistemology classroom, competing views of what epistemic rationality requires are

discussed, debated, and considered by novice philosophers. These novice philosophers have limited access to evidence about what rationality requires - all the evidence they have comes from their teacher, an apparent (but perhaps biased) expert who they have every reason to trust. They also have limited abilities to assess this evidence and determine what it supports. Suppose a teacher who holds a false view about what epistemic rationality requires skews her students' epistemic situations such that the evidence they have access to supports this false view. This epistemic environment is normatively misleading: agents in a situation like this are likely to acquire justified false beliefs about what epistemic rationality requires.[iii]

Online social media (or certain parts of it) is another example of a normatively misleading epistemic environment. The internet affords us access to a huge amount of information, but many agents access that information via platforms structured by algorithms aimed at increasing engagement. This means that the contingencies of what users find engaging determine to a large extent what one sees. Over time, we can expect this to alter what one finds reasonable and convincing. This, we might think, can affect what evidence one has. One example of this is agents becoming unwittingly trapped in “echo chambers” (see Nguyen (2020)). Echo chambers provide a skewed set of evidence that supports not only false beliefs about the world, but also false normative beliefs about how beliefs ought to be managed. They might, for example, provide evidence that discredits information sources outside the echo chamber, emphasise the epistemic superiority of first-hand experience, dismiss the epistemic value of expert testimony (or dispute what counts as expert testimony), or offer apparent defeaters for more mainstream indicators of epistemic adequacy. Inside the echo chamber, one might thus be justified in holding false normative beliefs about what epistemic rationality requires.

In general, normatively misleading environments arise whenever it is possible for the agent's evidence to support something false about what epistemic rationality requires. Note that there are two distinct ways that normatively misleading environments can provide justification for false beliefs about epistemic requirements. First, they can provide evidence for false beliefs about what the true epistemic requirements are. For example, in an epistemology class one might acquire evidence that one ought always to conciliate in response to peer disagreement, or that it is permissible to adopt beliefs for pragmatic reasons, or that one ought to maximise expected epistemic utility. Supposing that these views are false, one could nevertheless acquire evidence for them by considering convincing philosophical arguments in their favour or trusting the testimony of an apparent expert. Likewise, spending too much time engaging with conspiracy theorists’ social media accounts could provide evidence that a theory should only be accepted if it can explain away all coincidences with one coherent story. Second, normatively misleading environments could provide misleading evidence for false normative beliefs about which specific propositions one epistemically ought to believe. For example, in epistemology class one might acquire evidence that one ought never believe that one's lottery ticket will lose (because one can never be completely certain of this). Similarly, spending too much time engaging with conspiracy theorists’ social media accounts can generate misleading evidence for various specific normative beliefs, such as “it is not epistemically rational to trust the mainstream media”, or “I ought to believe that King Charles is a lizard (because this is what the evidence supports)”. In these cases, the agent has true beliefs about what epistemic rationality requires in general (proportioning beliefs to evidence) but has a false view about what her evidence is.

Some have thought that there can be no normatively misleading environments because the epistemology of rational requirements is such that we *always* have propositional justification for believing the truth about them. On these views, it is not possible to be rationally mistaken about what one is justified in believing. This goes both for mistakes about which specific propositions one is justified in believing (Smithies (2012: 274)), and mistakes about what the requirements of epistemic rationality are (Titelbaum (2015: 253)). If this were true, false beliefs about what epistemic rationality requires could *never* be rational, regardless of the origin of the mistake. However, this view of the epistemology of rational requirements involves some serious and contentious theoretical commitments. It requires us to assume that every agent, regardless of the further details of her epistemic situation, possesses vast resources of propositional justification capable of defeating any possible misleading evidence for false views of what epistemic rationality requires.[iv] Here, I set this view aside, and assume that normatively misleading environments, and rational false beliefs about what rationality requires, are possible.

I am not the first to argue that epistemic level-incoherence is sometimes rational. However, previous arguments for this claim have tended to assume that accepting rational level-incoherence is an epistemic cost, albeit one outweighed by more significant costs. For instance, evidentialists have argued that misleading higher-order evidence sometimes produces situations in which the level-incoherent combination "P" and "P is highly unlikely on the evidence" is supported on the evidence, and so, according to an evidentialist account of rationality, would be rational to believe (see Coates (2012), Lasonen-Aarnio (2014, 2020), Horowitz (2013), Weatherson (2019), Worsnip (2018). However, these evidentialists have usually also thought that there is some other sense in which agents ought not believe these level-incoherent combinations. For Lasonen-Aarnio (2020, 2021), although holding level-incoherent beliefs can sometimes be epistemically required, it manifests a bad epistemic disposition, one that indicates incompetence in dealing with reasons. For Coates (2012), it is not rational to *bet* on level-incoherent beliefs, although it can be rational to hold them. For Horowitz (2013), there is a general requirement of rationality to avoid level-incoherence, but this requirement admits of exceptions.[v] For Worsnip (2018), cases of apparently evidentially supported level-incoherence indicates that the rational demand to believe what the evidence supports represents only one kind of rationality, which must be reconciled with the equally important demands of coherence.[vi] Some have regarded this apparent tension as a theoretical cost that can be avoided by permitting some cases of rational level-incoherence. In previous work, I endorsed accepting rational level-incoherence as a theoretical cost worth accepting as way to resolve apparent tension between the demands of evidence and coherence (see Field (2021)).[vii]

Here, I take a more positive approach to rational epistemic level-incoherence. I argue that in some epistemic environments, agents who are level-incoherent are more rational, more attuned to the truth, and in a position to undertake better inquiries than level-coherent agents. Agents in normatively misleading environments are unlucky: they find themselves in a situation that makes it difficult to believe the truth about normative epistemic requirements, and they are also non-ideal -- ideal agents, perhaps, would be able to immediately construct or intuit a priori arguments for the truth about what is epistemically required. However, they are not rationally criticisable.

This is because in these environments, I argue, level-incoherence is the rationally dominant strategy for non-ideal agents.

2. Tension

Level-incoherent combinations of attitudes exhibit tension: intuitively, they seem difficult for a rational agent to hold at the same time. By tension, I mean a psychological experience associated with things not fitting together. Tension is the *feeling* we have when we try to hold things together that, seemingly, do not fit together. It has often been assumed that this tension indicates an important truth about rationality, namely that rational agents do not hold level-incoherent attitudes (see Broome (2013), Kvanvig (2014), Littlejohn (2018), Titelbaum (2015)). Some have attempted to elucidate it further by noting its similarities to Moorean absurdity (Smithies (2012)), or how it appears to involve self-contradiction (Elga (2007), Littlejohn (2018)), noting that it seems to exert "pressure" on us to revise attitudes (Fogal (2020: 1035), Kolodny (2007)). I do not aim to dispute it this intuition that level-incoherent attitudes are somehow in tension. Instead, I ague here that there are multiple ways to explain it. I consider here three distinct views of how to account for and explain the tension involved in level-incoherence. One of these is the traditional view (the "Prohibition View"), that the tension involved in level-incoherence indicates that level-incoherence is rationally criticisable. I consider two alternative views: the Error View, and the Inquiry View. I ultimately argue that the Inquiry View is the correct way to account for the tension.

First, the traditional Prohibition View.

> **Prohibition View**: Tension indicates that the source of the tension is rationally prohibited.

According to the Prohibition View, the tension involved in level-incoherence indicates that the combination that gives rise to the tension is rationally prohibited. According to the Prohibition View, level-incoherent combinations of attitudes give rise to tension because they are prohibited by epistemic rationality. Agents who have such combinations of attitudes are, in some sense, rationally criticisable. In other words, level-incoherence *seems* irrational because it *is* irrational - for all agents, in all situations.

There are different ways to explain *why* combinations that exhibit this tension are rationally prohibited, all consistent with the Prohibition View. For example, some would take the prohibition of level-incoherence to be explained by reference to a distinctive structural requirement of epistemic rationality (Broome (2013), Lee (2021)). Alternatively, those who deny the existence of distinctive structural requirements of rationality have sought other ways to accommodate the rational criticisability of level-incoherence. A popular way to do this is in terms of what appropriate responsiveness to reasons demands. On this way of spelling out the Prohibition View, agents are always under rational pressure to be level-coherent because, in fact, responding appropriately to one's epistemic reasons always means being level-coherent, and failing to do so is in some sense rationally criticisable (see Kiesewetter (2017), Lord

(2018)). Another way to do this is in terms of competence with reasons (see Lasonen-Aarnio (2020, 2021), Littlejohn (2018)). According to Lasonen-Aarnio, someone who is level-incoherent is criticisable because "they manifest dispositions to fail to be responsive to a special class of conclusive and conspicuous reasons" (2020: 630). These dispositions are "at odds with normative competence" with respect to dealing with one's reasons (2021: 461).

The Prohibition View makes a strong claim. It says that level-incoherence is *always* rationally criticisable: there are no counter-examples. Defending the Prohibition View thus means refuting all possible counter-examples. Not only would this be an onerous task,[viii] but many who defend the idea that level-incoherence is in general rationally criticisable do not actually intend this to be a necessary and universal truth about rationality. For example, in a paper defending the idea that epistemic level-incoherence is in general irrational, Horowitz (2013) explicitly points out that there are exceptions (namely, Williamson (2013)'s unmarked clock cases). Broome, while defending the idea that there are distinctive structural requirements of rationality, including requirements of level-coherence - notes that there are nevertheless exceptions. For example, it is a requirement of rationality to avoid believing contradictions, *unless* one also thinks that some contradictions are “special” - for example, if one is a dialetheist and thinks that instances of the Liar Paradox are true contradictions (Broome (2013: 91)).

Even worse, the Prohibition View appears to generate a paradox of rationality. Suppose that there is at least one true requirement of rationality, R.

R: Doxastic state D is rationally required in situation S.

Suppose also that rationality demands that agents respond appropriately to their epistemic reasons.[ix] Then, suppose that an agent A is in an epistemic situation such that it would be rational for her to believe the opposite of R, R*.

R*: Doxastic state D is rationally forbidden in situation S.

In fact, R* is false. Suppose A is in S, and knows she is in S. What doxastic state is she rationally required to have? According to the true requirements of rationality, she is required to have D, since she is, in fact, in S. However, she is also required to respond appropriately to her epistemic reasons, which in this case support believing that R* is true and R is false. If Level-Coherence is *also* a requirement of epistemic rationality, then epistemic rationality is inconsistent. Recall:

Level-Coherence: $O(BO\varphi \rightarrow \varphi)$

If Level-Coherence is a requirement of rationality, then in this situation rationality requires A to *both* have D and to *not* have D.[x]

The paradox does not arise if we do not insist that conformity to Level-Coherence is a necessary condition of rationality. If it is not, then we are free to think of this as another exception, in which it is rational to have a level-incoherent combination of attitudes: in this case, to believe that D is prohibited while also holding D. Of course, there are other possible responses, but none come without theoretical cost. For example, one could perhaps think that the agent ought not believe R*, but as the situation is set up R* is

supported by the agent's evidence. So, if she does not believe R* she would be failing to believe what her evidence supports. Some have responded by suggesting that perhaps this shows that one's evidence could never really support something like R*, it only *seems* that way (see Littlejohn (2018) and Titelbaum (2015)). Alternatively, one could argue that *if* one's evidence can support something like R* then there is no rational option available (both Littlejohn (2018) and Titelbaum 2015) suggest positions like this. Kvanvig (2014) argues that what rationality requires here is changed by what it is rational for A to believe about rationality. This, as Lasonen-Aarnio (2014) points out, is another way of saying that there are really no traditional universal requirements of rationality at all. Worsnip (2018) argues that this apparent puzzle reveals a conflict between two distinct senses of rationality. Hughes (2022, forthcoming) argues that this reveals a genuine inconsistency within rationality and is an example of an epistemic dilemma. The point I want to note here is that the paradox arises only if we assume that the Prohibition View is true, and we need not think the Prohibition View is true, since there are at least two other ways of explaining the tension that arises from level-incoherence, neither of which generate this paradox.

One such alternative is the Error View.

> **Error View**: Tension indicates nothing about what is rationally prohibited or criticisable.

The Error View denies that the tension arising from level-incoherence indicates anything in particular about epistemic rationality. Instead, the right way to explain this tension is by reference to our psychology. Perhaps we feel tension in response to particular combinations because for our ancestors it was evolutionarily advantageous to avoid having those combinations. This does not tell us anything about epistemic rationality, only what was adaptive for our ancestors.

To push this explanation further, we might think that the avoidance of these combinations was not only evolutionarily advantageous, but also often correlated with genuinely epistemically positive traits. If so, this could explain why it has so often been thought that level-coherence is demanded by rationality: it very often correlates with epistemically positive traits. However, correlation does not imply causation, and nor does it imply necessary truths about rationality.

Kolodny (2007) endorses a view close to the Error View. He argues that the “normative tension” we experience in response to incoherence combinations of attitudes is not best explained by appeal to wide scope requirements of formal coherence. Instead, it can be explained entirely by appeal to narrow-scope requirements for attitudes, and facts about the agent’s epistemic situation - her evidence, reasons, and epistemic attitudes. Since our reasons almost always demand that we are coherent, this explains why we might have thought that there were distinctive formal requirements of coherence. Although it is extremely rare, it is nevertheless conceptually possible that sometimes our reasons demand that we adopt incoherent attitudes. That said, Kolodny nevertheless does not think it conceptually possible that our reasons could ever require us to be level-incoherent. Indeed, he endorses a narrow-scope requirement requiring level-coherence between what one believes about what reason requires and one’s first order beliefs (2007: 243).

Finally, there is a third possible explanation for why level-incoherent combinations seem in tension: the Inquiry View. This is the view that I go on to endorse.

> **Inquiry View**: The tension is indication that further inquiry is needed but is not itself rationally criticisable.

According to the Inquiry View the tension involved in level-incoherence does not indicate anything about rational prohibition or criticizability. Instead, it says that attitudinal combinations that are in tension are reasons to inquire further. This is because the tension indicates that there *may* be an epistemic problem with at least one of the attitudes in the combination – they might be irrational, false, or unjustified. However, tension does not *guarantee* this. It is entirely possible that the combination is unproblematic, but only further investigation can determine this. On the Inquiry View, regardless of how things turn out, tension gives us a reason to investigate further.

To use an example, consider warning lights on the dashboard of a car. When these come on, they indicate that there is some problem with the car. Usually, it is not possible to determine exactly what is wrong merely from the fact that the warning light has come on. This is particularly true for older cars, where the warning lights can be very coarse-grained – consulting the manual might reveal that a particular warning light indicates a "problem with the engine". To determine exactly what is the problem with the car, further investigation is needed. It is also possible that there is no problem at all, and the light is malfunctioning. On the Inquiry View, the tension arising from level-incoherence is like the warning light on the car dashboard. Like the warning light, it is a reason to inquire further, but it is not *itself* the problem. The problem is whatever is causing the warning light to come on.

Why does tension give us a reason to inquire further? The tension exhibited by level-incoherent attitude combinations is shared by various other incoherent attitude combinations, for example contradictions and Moorean assertions. This tension gives us a reason to inquire because we know that many of the attitude combinations that exhibit this kind of tension have epistemic problems. For example, contradictions cannot be true, and Moorean assertions cannot be asserted rationally.[xi] Tension very often, but not always, coincides with epistemic problems. This gives us a reason to investigate combinations that are in tension further – it is likely that they are epistemically problematic and thus call for some kind of revision. However, not all instances of tension are epistemically problematic. On at least some accounts of epistemic normativity, level-incoherent combinations of beliefs can sometimes be true, supported by the evidence, and rationally believed. In such cases, further inquiry does not rationalize any revision of attitudes. However, the cases in which level-incoherence are unproblematic are somewhat unusual and rare. Since the unproblematic cases are unusual, tension provides a reason to inquire. It is nevertheless possible that after inquiry, level-incoherence turns out to be entirely appropriate and rational. When this is the case, the Inquiry View permits the agent to retain the level-incoherent combination.

One might wonder how having a reason to inquire further explains the intuitive sense that there is something rationally criticisable about level-incoherence. For instance, suppose you have good evidence for P, and believe P. Someone you trust then says: "Since you believe P, you might want to read today's newspaper". In this case you plausibly have a reason to investigate further, but there is no tension and you do not

seem rationally criticisable. In this case, the observation that you have a reason to inquire does not suggest that you are criticisable, so it is somewhat puzzling how noting that level-incoherent agents have a reason to inquire could explain why level-incoherence has seemed irrational.[xii] This is because reasons to inquire come in different strengths. For weaker reasons to inquire there is no epistemic failing if the agent does not respond to them and inquire. If P is not particularly important, or you have no reason to think that reading the newspaper will radically change your mind, it is plausible that you are not criticisable if you do not inquire. However, this would plausibly change if P was central to your other inquires, or if the person you trust also told you that today's newspaper conflicts with P. In this case the person you trust is giving you evidence that one of your beliefs is false, and this increases the strength of your reason to inquire. The stronger your reason to inquire, the more criticisable you would be for failing to inquire. I do think that some who knows they are level-incoherent and doesn't investigate, is criticisable. However, this is importantly distinct from a position that views level-incoherence as itself criticisble. On the Inquiry View, what is criticisable is the failure to inquire when one has a strong reason to inquire.

The main advantage of the Inquiry View is that it can explain both what appears bad and also what is good about level-incoherence (in particular epistemic environments). The following section explains this distinctive epistemic value of level-incoherence in more detail and argues that the Inquiry View offers the best option for accommodating it.

3. The Epistemic Value of Level-Incoherence

This section outlines how level-incoherence can be epistemically valuable in normatively misleading environments. Normatively misleading environments are epistemic environments in which it is easy to acquire false justified normative beliefs. As I explain further in the following subsection, I am particularly interested here in more general false normative beliefs, because of the wide-ranging consequences they have for our epistemic agency and rationality.

I argue in this section that of the three explanations for tension arising from level-incoherence, only the Inquiry View can accommodate its value. Consider the following case of an agent in a normatively misleading environment:

> **Logic 101.** Suppose that, in fact, rationality requires us to refrain from believing contradictions and there are no exceptions to this. Sam is a typical undergraduate student enrolled in an introductory logic class. Like any rational agent, she has ordinary intuitions about basic truths of logic, However, her professor is a dialetheist who intends to set his students on the right track by exposing them to all the best arguments in favour of dialetheism and all the worst arguments against it. By studying the arguments, Sam comes to believe that sometimes the rationally required doxastic attitude is to believe a contradiction, for example in response to some semantic paradoxes.[xiii] However, when faced with an instance of the kind of contradiction that dialetheism says she is required to believe, she cannot quite bring herself to believe it. She regards this as an epistemic failing.

Sam is level-incoherent. She believes that she ought to believe the contradiction yet does not believe it. Here is another case:

> **Conspiracy Theories.** Through his engagement with social media, Billy has unwittingly fallen into an echo chamber. The information he receives from the algorithmically driven social media platforms he uses consistently appears to convincingly support various false views, both about the world and about how beliefs ought to be managed. For example, the normative belief that rational agents "do their own research", where this means not trusting any testimony they cannot verify for themselves independently. However, every morning Billy is woken up by his radio alarm clock, which plays a true news report from a genuinely reliable news source that Billy is not in a position to verify himself. Billy believes what the news report says, even though this involves trusting an authority rather than "doing his own research".

Billy is level-incoherent. His first-order beliefs do not conform to his normative beliefs about how those beliefs should be managed.

Unlike other forms of incoherence such as logical inconsistency, level-incoherent belief combinations can be true, and on at least some accounts, supported by one's evidence (see Christensen (2010), Lasonen-Aarnio (2014, 2020), Weartherson (2019). Even more importantly, level-incoherence can help agents like Sam and Billy undertake better inquiries. Specifically, and as I argue in this section, it can: limit the damaging effects of false normative beliefs (§3.1), enable ordinary deliberation about normative propositions (§3.2) and function as a 'red flag' that can alert agents to possible epistemic problems (§3.3). Level-incoherence is thus highly valuable in leading agents in normatively misleading epistemic environments towards knowledge. Note that the value of level-incoherence is not purely consequentialist. If we take seriously the possibility that we are mistaken about our normative beliefs, then avoiding commitment to level-coherence can be a way of manifesting appropriate concern for and attunement to the truth.

3.1 Damage Control

The first way that level-incoherence can improve inquiry in normatively misleading situations is by limiting epistemic damage done by exposure to misinformation. Agents in normatively misleading environments are often justified in false normative beliefs about how beliefs ought to be managed. These false normative beliefs, if lived up to, have very undesirable epistemic consequences. For example, believing that one ought not trust sources outside one's echo chamber could mean losing out on many true beliefs, acquiring false beliefs (if one's echo chamber contains misleading sources instead), and losing knowledge one has previously acquired from sources outside one's echo chamber (because the false belief that one ought not outside sources would provide an undercutting defeater for knowledge previously obtained). Indeed, others have noticed the potentially damaging effects of false beliefs about what rationality requires. Way (2020) notes that reasoning enkratically from one's normative beliefs to the first order attitudes they recommend does not preserve correctness, and Alexander regards false normative beliefs as an "occupational hazard[s]" (2013: 2-3) of philosophizing about rationality and epistemic justification. A commitment to Level-Coherence means that forming beliefs about what is rationally required, or even

doubting one's current beliefs about these matters, has immediate consequences for the rest of one's beliefs. An agent who has false normative beliefs is significantly worse off. Not only does she have false beliefs, but those false beliefs have further negative epistemic consequences – they contribute to justifying additional false beliefs and could prompt the agent to give up true beliefs and knowledge.

Even worse, this epistemic damage would be *rationally required* when the agent has false normative beliefs about what is rationally required and is in an epistemic situation that strongly supports these false normative beliefs. If the agent's epistemic situation supports a belief to a sufficiently strong degree, and the agent does not have the option of improving her epistemic situation, then it is natural to assume that she does not have the option of rationally giving up the false normative belief. If this is right, then the only way to comply with Level-Coherence would be for the agent to adopt the false beliefs recommended by the false normative belief. Sam and Billy have sufficient available and undefeated evidence for the false normative belief, and they do not have any other available evidence against it, because they lack the relevant capacities to gather evidence that would defeat the belief. Minimally, Sam and Billy are rationally permitted to hold the false normative belief. If we also think that there are positive epistemic duties,[xiv] then they may also be rationally *required* to believe it.

Damage control is particularly important when we are likely to have justified false *normative* beliefs. While any belief can in principle be epistemically damaging in virtue of the further false beliefs it implies, the content of normative beliefs makes them particularly dangerous. Normative beliefs, particularly those that make generalized claims about which kinds of doxastic attitudes are required and permitted, and why, have significant and wide-ranging consequences for an agent's epistemic life. They have the power to guide not only which particular doxastic attitudes an agent has on a particular occasion, but also the more general principles she uses for forming doxastic attitudes on many different topics. This makes it particularly important to minimize potential damage arising from false normative beliefs. For this reason, *level*-incoherence has particular epistemic value for agents in normatively misleading environments.

3.2 Normative Deliberation

False normative beliefs can cause a great deal of epistemic damage. This makes it even more important to deliberate well about them, in the hope that we might end up with *true*, rather than false, normative beliefs. True, rather than false, normative beliefs can help agents in normatively misleading environments overcome those environments.

Rationally replacing false normative beliefs with true ones demands good deliberation about normative matters - it demands that agents consider and respond to evidence about what rationality requires, gather additional evidence, and reason about what it supports. During this process, while the agent is deliberating, it would be inappropriate to attempt to maintain level-coherence. This is because during the deliberation process, one's views are not fixed - or at least, they ought not be. As new evidence about what rationality requires is received, rational agents should update their beliefs. However, if Level-Coherence is a rational requirement that applies to all agents at all times, then agents who are deliberating about what is required would also be rationally

required to update their first-order attitudes as their normative beliefs change. This would be burdensome and hamper the deliberation process.

Note that we cannot avoid this problem even if we thought, as some have, that our epistemic access to the normative truths of epistemology is a priori and universally guaranteed (see Kitcher (2000a, 2000b), Smithies (2012, 2015), Titelbaum (2015)). Even if it is true that we all possess maximally strong propositional justification for the truth about what rationality requires, this would not rule out the possibility of apparent defeating evidence, and this would need to be deliberated about. If Level-Coherence is a rational requirement, then such deliberation would come with burdensome and potentially damaging commitments.

Another way one might attempt to avoid this problem while preserving Level-Coherence is by distinguishing between belief and some other, less epistemically committal state, such as acceptance. Perhaps we need not adjust our *beliefs* while deliberating, only what we accept. And, perhaps level-incoherence between beliefs and acceptances is rationally permissible. Perhaps, though this is somewhat unsatisfying as a solution, because it will still sometimes be desirable to deliberate about attitudes even once they have acquired whichever 'fixed' status we think appropriate for attitudes one has settled on. Once we consider this possibility, the same difficulties return.

I am not the first to suggest that level-coherence requirements of rationality should be relaxed to allow for normative deliberation. Svavarsdottir (1999) argues that level-incoherence can be a rational response to changing one's normative moral beliefs, since when one adopts a new moral view there will predictably be a period of time after one has adopted the belief, but before one has fully regulated one's action in line with it, and this is entirely reasonable and to be expected. For the same reason, Weatherson (2019) restricts his arguments *for* rational level-coherence to cases where the agent has settled views, the idea being that it is too easy to find cases of rational level-incoherence when the agent is mid-deliberation. Similarly Staffel (forthcoming) holds that level-incoherence is permissible during the deliberative process. It is plausible that if we want agents to deliberate successfully about what they ought to believe, as we should, if we accept the possibility of normatively misleading environments and want agents to be able to rationally exit these environments - then we should also avoid thinking that agents are rationally required to maintain level-coherence while deliberating about what they ought to believe.

3.3 Red Flags

Level-incoherent beliefs, exhibit a *prima facie* tension. This tension is epistemically valuable because it can function as a 'red flag': a signal to the agent that further inquiry is needed.

We can characterise this tension by the effects we can expect it to have on a rational agent - it should seem suspicious. For example, in classical logic contradictions work as stopping points in arguments. Since it must be agreed by all who accept classical logic that contradictions can never be true, they are able to perform various useful functions in argumentation. For example, proof by contradiction. If a contradiction is derived from some proposition, P, then it should be clear to all rational agents that P is false. Part of what enables contradictions to have this function is the clear tension

between the two contradictory components. Regardless of what else interlocutors might disagree about, they should be able to agree that a contradiction is a problem.

Level-incoherent combinations might be thought to function in a similar way. For example, Lasonen-Aarnio (2020) holds that it is unreasonable to hold level-incoherent combinations of beliefs. This is not because there is anything necessarily epistemically irrational about level-incoherent beliefs - indeed, Lasonen-Aarnio holds that there are possible cases in which one's total evidence supports level-incoherent beliefs, and in such cases level-incoherent beliefs would be rational. However, she holds that exercising "basic epistemic competence" precludes knowingly having level-incoherent attitude combinations. This is because the normative component in such combinations (e.g. "I ought not to believe P") would, if it were true, be a conspicuous reason to not believe P. Someone who genuinely believes this thus has an apparent, conspicuous reason not to believe P, and would be manifesting an "epistemic vice" by failing to respond to it by revising one of the level-incoherent components (2020: 604).

However, the observation that level-incoherence should be noticed and disapproved of by a rational agent does not necessarily imply that it is a reason to *revise* beliefs. It is not true that an agent with level-incoherent beliefs should always revise her beliefs. Agents who are level-incoherent do not always have a better, level-coherent option rationally available to them. Sometimes, she will not be in a position to do so rationally, even after undertaking a reasonable amount of further inquiry. Instead, I suggest here, level-incoherence provides a reason to *inquire* further, because it indicates that something may have gone wrong.

Consider Sam. She is level-incoherent, and in so far as she is aware of this, this should indicate to her that further inquiry into what rationality requires is needed. However, given her epistemic situation, it would take a serious amount of inquiry and intellectual effort for her to arrive at a rational and coherent view. Either she needs to defeat the dialetheist arguments and justify the Law of Non-Contradiction, or she needs to let go of and dismiss the inclination she feels to avoid believing contradictions. As an undergraduate, she is in no position to do the former. If the Law of Non-Contradiction really is a basic logical truth, it is not clear she is in a position to do the latter either.[xv]

Being level-incoherent in a normatively misleading epistemic situation is epistemically valuable because it gives the agent an additional epistemic tool which can, if used correctly, help her escape the normatively misleading epistemic situation. If rationality required us to always resolve level-incoherent tensions immediately, then agents would lack this prompt to undertake further inquiry. Of course, they might through pure luck resolve level-incoherent tensions in the right direction, but there is no guarantee of this. If anything, it seems more likely than not that agents in normatively misleading environments will resolve in the wrong direction. Consider Billy again. Depending on the details of his online activities – that is, how much and what kind of evidence he has – it could be that at the moment he realizes his level-incoherence, the most natural way to resolve the incoherence would be in the direction of the false normative beliefs he has acquired from his online echo chamber, rather than in the direction of the true news report. If Billy does this, he would maintain level-coherence, but end up in a worse epistemic situation: rationally committed to forming further false beliefs and with no reason to reconsider these false beliefs. Having dismissed the news report, all would seem well from his perspective. So, if rationality required level-coherence, Billy would

be under pressure to resolve the conflict, and would be likely to resolve it in the wrong direction.

Alternatively, if Billy believes the news report in addition to the false normative belief, then he is in a better epistemic position. He is aware of a tension that needs resolving, and he has a reason to undertake further inquiry. Of course, there is no guarantee that he will use this tool correctly - he will be able to rationally leave the echo chamber only if he undertakes the right inquiries and acquires further evidence in favour of the true belief. However, there is at least hope.

4. The Inquiry View

I have argued that level-incoherent *belief* has distinctive epistemic value. This epistemic value is something that, ideally, our theories of rationality should be able to acknowledge. They cannot if we endorse either the Prohibition or the Error Views. The Prohibition View excludes the possibility of rational level-incoherence, and the Error View accepts it too easily, and so cannot account for why level-incoherence would function as a red flag indicating the need for further inquiry and re-examination. This leaves the Inquiry View as the best explanation of the intuitive idea that there is something rationally amiss with level-incoherence.

The Inquiry View differs from the Prohibition View in that it does not think that level-coherence is rationally criticisable. Level-incoherent combinations are in tension, and demand investigation, but agents are not rationally criticisable for being level-incoherent. As with the car warning light, the tension arising from level-incoherent combinations is a reason to inquire further, but it is not *itself* the problem. Just as the problem in the car is whatever in the car engine is causing the warning light to come on, the epistemic problem is whatever is causing the level-coherence (for example skewed evidence, misleading higher-order evidence, or imperfect reasoning). And, just like the car warning light, there might turn out to be no problem (it may just be that the light is malfunctioning). It is conceptually possible that after undertaking all the further investigation it is reasonable to expect her to take, the agent still finds that the balance of her epistemic reasons support something level-incoherent.

Indeed, this is likely the situation Sam is in: it's not reasonable to expect her to rationally overturn either the convincing dialetheist argument or her own a priori logical insights. Cases of misleading higher-order evidence may provide additional examples (see Christensen (2010), Lasonen-Aarnio (2010, 2014, 2020), Weatherson (2019)).[xvi] However, as for the car warning light, none of this can be determined in advance of investigation. Whether or not the level-incoherent belief combination is epistemically sound cannot be determined in advance of further inquiry.

The Inquiry View differs from the Error View in that it always views level-incoherence as a reason to inquire. The Error View, meanwhile, merely seeks to explain why we *tend to think* that level-incoherence is irrational. Explanations offered by the Error View can explain our attitude towards level-incoherence, but they do not justify further inquiry, revision, or any other epistemic response. An example of an explanation that the Error View might offer is that due to our evolutionary history we have psychological biases against attitude combinations that are in tension. This accounts for our belief

that there is something wrong with level-incoherence without giving tension any epistemic significance. If we accept the Error View, level-incoherence loses its value as a reason for inquiry and possible tool for epistemic improvement. The Error View can thus not account for the epistemic value of level-incoherence in normatively misleading environments.

The rest of this section anticipates and addresses three possible objections to Inquiry View. Responding to these objections will allow me to clarify further the details of the view.

4.1 Why Believe?

The Inquiry View permits agents to rationally hold combinations of level-incoherent full beliefs. However, one might worry that these full beliefs conflict so much that they risk defeating each other. One might thus reasonably wonder whether the epistemically valuable features described in the previous section could not be just as well using doxastic states other than full belief, such as suspension, partial belief, or credal states. Perhaps, one might think, agents like Sam and Billy ought not *believe* the level-incoherent combinations, but instead take some weaker attitude towards them. In the following paragraphs I explain why I do not think we should take this approach.

First, do the full beliefs involved in level-incoherence combinations always and necessarily defeat each other? In other words, is something like the following principle true?

> **Inter-Level Defeat:** A normative belief that φ-ing is required has defeating force with respect to doxastic attitudes incompatible with φ-ing.

I think Inter-Level Defeat is false. First, note that this is a very strong principle. Where to φ is to adopt some doxastic attitude, it says that having a normative belief can, by itself, rule out the possibility of rationally holding any attitudes incompatible with what that normative belief requires. This rules out the possibility of the agent having independent epistemic reasons to *not*-φ that outweigh the defeater provided by the belief that one is required to φ. For example, it implies that if Sam believes that she ought to believe the contradiction, then she cannot rationally avoid believing the contradiction, even if to do so would be to respond appropriate to genuine truth-tracking a priori insights. Responding to these insights is no longer rational if Inter-Level Defeat is true. One reason to think that Inter-Level Defeat might be true would be a commitment to a perspectivist view of rationality, according to which what one *believes* about what is rationally required completely and conclusively determines what is in fact rationally required (Kvanvig (2014), Whiting (2022)). In other words, merely believing that φ-ing is required makes it the case that φ-ing really is required. This account of rationality incompatible with also thinking that there are robust, mind-independent truths about what rationality requires, and if mistakes are possible about what rationality requires, there must be some fact of the matter about what rationality requires that holds independently of what I believe about it.[xvii] Not only this, but it is very difficult for normative beliefs to act as defeaters for first-order beliefs from a subject's perspective. As Coates (2012: 8) points out, it is hard to see how subjects could regard normative beliefs as defeaters for first-order beliefs. Even when a normative belief such as “I ought not believe P” is *true*, the normative belief itself is not what makes it irrational to hold the first-order belief. Rather, when true, this normative

belief indicates that there is something epistemically amiss with believing P, such that it was never rational to believe P in the first place. For illustration, the reasonable thing for someone who believes P, and then receives good evidence for "I ought not believe P" to think is that she has made some error in her reasoning when coming to believe P. However, it is this error in the reasoning towards P that defeats P, not the normative belief. In other words, the normative belief "I ought not believe P" does not defeat P by itself. It could count against believing P, but only evidence that supports this belief is also evidence that there is some other defeater for P. This is precisely what the Inquiry View predicts.

It is worth noting that Inter-Level Defeat seems even less plausible in other normative domains. For example, consider the moral domain, where level-incoherence has been viewed somewhat positively in cases where the agent has a reasonable false normative belief about morality. Here, some have thought that there can be cases of "inadvertant virtue" that deserve positive evaluation, on the grounds that the agents are responding appropriately to genuine moral reasons, in spite of their false normative beliefs. The classic and widely discussed case is Huck Finn. Huck lives in a slaveowning society and believes that he is morally required to refrain from helping slaves escape - this would amount to stealing property and would be wrong. However, Huck has befriended a slave, Jim, who has asked for his help to escape. Against his "better judgment", Huck feels himself compelled to help Jim escape. Many authors in the moral responsibility literature have classed this case of "inadvertent moral virtue" as deserving of positive moral evaluation. For example, Arpaly (2002) and Arpaly & Schroeder (2013) suggest that the fact that Huck is moved by some deep-seated, intrinsic concern for the right-making features of the situation (Jim's humanity) explains why Huck is praiseworthy, despite his holding false normative beliefs about his action. As Arpaly tells it, although Huck believes he ought to turn Jim the slave in, he finds himself "psychologically unable" to do so. This psychological inability is a response to "data" that Huck has not reflected upon, and perhaps is unable to reflect upon (Arpaly 2002: 77). Since this "data" in fact comprises the reasons that make freeing Jim the right thing to do, and he responds appropriately to them, he is praiseworthy.[xviii] Similarly, some authors have thought that there can be cases of inadvertant *epistemic* virtue (Kearl (2022), Lasonen-Aarnio (forthcoming), Weatherson (2019: 171)). If this is right, then Sam and Billy seem like good candidates - like Huck Finn, both respond to genuine epistemic reasons, despite reasonable false normative beliefs. If the epistemic domain is anything like the moral domain, we should be suspicious of Inter-Level Defeat.

Another way that one might seek to defend Inter-Level Defeat is via appeal to what one's total evidence could support. Worsnip (2018) argues that to violate level-coherence requirements is to "disrespect" one's evidence by failing to take seriously one's own judgments about what the evidence supports. This makes sense only if we assume that evidence can never support level-incoherent combinations. Some have given arguments for this, though they have mostly operated from within an evidentialist framework of rationality, according to which "I ought not to believe P" is functionally equivalent to "my evidence does not support P". For example, Brown (2018) and Skipper (2019) argue that evidence for "my evidence does not support P" defeats any evidence one might have for "P". In similar vein, others have argued that "evidence of evidence is evidence" (Comesana & Tal (2015), Tal & Comesana (2017)), and that normative beliefs about what the evidence supports are "evidence of evidence" that can defeat other evidence once might have for conflicting claims. González de Prado (2019) argues that evidence that one's evidence does not support P can prevent the

agent from *possessing* the evidence for P. If these arguments are successful, then it would be impossible for level-incoherent beliefs to be supported by one's total evidence, and so, at least on a purely evidentialist framework of rationality, these combinations are guaranteed to be irrational.

However, this is too quick. First, the success of these arguments depends on there being a very strong evidential support relation between "P" and "I ought to believe that P". The strongest possible evidential support relation is entailment. If "P" entailed the matching normative proposition "I ought to believe that P". When two propositions entail each other, then evidence for one is also evidence for the other and vice versa. However, if the support relation is anything weaker than entailment, this is not the case: it is possible for one to be true while the other is not (see Roche and Shogenji (2013); Roche (2014)). If "P" and "I ought to believe that P" do not entail each other, then it is possible to generate counterexamples in which the agent has total evidence supporting both "P" and an unmatching higher-order proposition such as "I ought not believe that P". Indeed, no evidential support relations short of entailment will be sufficient to rule out all possible counterexamples in which level-incoherent combinations can both be supported.

Indeed, "I ought not believe that P" and "not-P" do not entail each other. Suppose that last night instead of cycling home as usual and parking your bike in the garage at home, because it was raining you left your bike at your office. You have no reason to suspect that it is anywhere else now. The following proposition is supported by your total evidence: "I ought not believe that my bike is in the garage". You remember leaving it outside your office, your memory is reliable, it is unlikely to have been stolen, and so on. However, none of this entails that your bike is not actually in your garage. Indeed, it is possible, albiet unlikely, that your flatmate borrowed the bike, and took it home, leaving it in the garage. Evidence for "I ought not believe that P" is thus not conclusive entailing evidence for "not-P".

Of course, it is very plausible that evidence for "I ought not believe that P" is *some* evidence that you rationally ought not believe that P, and vice versa. However, there is no guarantee that in every case this will be sufficient to defeat unmatching level-incoherent doxastic attitudes.

Second, these arguments are successful only on a purely evidentialist account of rationality, such that one's total evidence completely determines what it is rational for her to believe. For example, Adler (2002), Owens (2002), Kiesewetter (2017) all define the epistemic forms of level-incoherence they discuss as believing what one does not regard oneself as having sufficient evidence to believe. However, the cases of rational level-incoherence under discussion here are slightly different. They involve agents who adopt doxastic states that go against the recommendations of their normative beliefs about epistemic rationality. Since evidentialism does not exhaust the theories of epistemic rationality that could be rationally believed, these cases are different. For example, we might think that there are epistemic reasons to not believe precisely what our evidence supports. Silva (2018) considers a case in which we have some evidence but know that there is more evidence we could gather. If we think it is possible that this extra, as yet ungathered, evidence could conflict with our current evidence, this could be a reason to avoid believing what our evidence supports.

I do not think we should accept a principle like Inter-Level Defeat. However, even if we did, it is not clear that an alternative doxastic attitude such as suspension could do the

same work as level-incoherent full belief. Note, first, that suspension can also form part of a level-incoherent combination of attitudes. If I believe I ought to believe P, but then suspend on P, this is level-incoherent: I am not adopting the attitude that I believe I am required to adopt. So, for this to work as a genuine alternative to my view, the agent could not merely suspend on the first-order attitude, but must suspend on *both* the normative belief and the first-order belief.

This would mean ignoring one's evidence. As the cases are described, the agents have evidence that supports both the false normative belief, and the conflicting first-order attitude. Suspending or taking some other partial attitude towards these would mean disregarding that evidence. This is not without its own epistemic costs. For example, Simion argues that we are under positive epistemic duties to believe propositions that we have evidence in favour of. Specifically, she argues that "we have an epistemic duty to form a belief that p just in case we have sufficient available and undefeated evidence for p" (2023: 2). Since level-incoherent beliefs do not automatically defeat each other, Sam and Billy are in precisely this position. They have sufficient and undefeated evidence for level-incoherent combinations of beliefs. If there are positive epistemic duties to believe what our evidence supports, then this puts them under an epistemic duty to have level-incoherent beliefs. To not do so would be to fail to meet this epistemic duty, and inappropriately disregard evidence.

Not only this, but it's not clear that suspension can provide as strong a reason to inquire as tension. Perhaps if one is forced, on pain of level-incoherence, to suspend on a question one desires to answer, one has a reason to inquire further. However, the reason to inquire comes from the reasons one has to answer the question, not the suspension itself. Indeed, one may suspend judgment in response to a variety of epistemic situations, but not all of these reasons are situations in which the agent has a reason to inquire further. For example, one might suspend judgment because one's evidence does not settle the matter, and one has no way of gathering more evidence. In this case, one does not have a reason to inquire further – suspension is the result of an inquiry that can be pursued no further.

Nor can partial attitudes or credences provide quite as good a reason to inquire as full beliefs in tension with one another. Level-incoherent beliefs can serve as red flags to prompt inquiry precisely because they are in tension. If the beliefs are only partial, this tension is muted, and so the signal is weakened, and so less useful. Likewise if the attitudes are credal – assuming they are probabilistically coherent, there is no tension (if they are not probabilistically coherent, this just introduces more incoherence). In comparison, if the agent *believes* some of these propositions, and those propositions conflict with other propositions she beliefs, she has immediate reason to inquire and resolve the tension.

The Inquiry View explains all this very well. According to the Inquiry View, level-incoherence is a defeasible indication of epistemic problems. Unlike the Prohibition View, the Inquiry View says that level-incoherence is not rationally prohibited, so level-incoherent agents are not under pressure to revise level-incoherent beliefs. However, nor is it acceptable to simply accept level-incoherence and not undertake any further epistemic activity. Unlike the Error View, the Inquiry View says that level-incoherent agents have a reason to undertake further inquiry.

4.2 Why Level-*Incoherence?*

I have argued that level-incoherence has distinctive epistemic value in normatively misleading environments. However, one might wonder whether my arguments could be applied to incoherence more generally. Are other kinds of incoherent beliefs also useful for agents in unfavourable epistemic environments?

While incoherent beliefs of any kind can serve as a 'red flag' to prompt inquiry, there are reasons to focus on the case of level-incoherence in particular.

First, there are additional reasons to think that rationality requires agents to avoid some kinds of incoherence, that do not transfer over to level-incoherence. For instance, logically inconsistent propositions cannot both be true, and this can be known in advance merely by appreciating the logical form of the proposition. This provides a reason for rational agents to avoid such belief combinations that is not shared by *level*-incoherent combinations. Additionally, logical inconsistencies generate triviality when combined with standard classical disjunction rules - because they logically imply any and all statements. Since it is not rational to believe *everything*, rational agents have a good reason to avoid believing contradictions. For this reason, Lewis (1992) suggests that they must be 'quarantined' if believed, to avoid triviality.[xix]

Second, level-incoherence specifically involves *normative* beliefs, and the environments in which it is useful are those in which the agent's epistemic situation prevents her from forming true normative beliefs. False normative beliefs can be particularly epistemically damaging because they have the potential to influence how the agent forms beliefs in general. This makes avoiding level-coherence particularly important. That said, I do not wish to rule out the possibility that incoherence more generally could also be epistemically useful in some circumstances, this, however, would be a different point requiring a different set of arguments.

4.3 Reason-Responsiveness?

One might worry that the Inquiry View is not significantly different from standard accounts of substantive rationality, according to which to be rational is to respond appropriately to one's reasons. Indeed, the Inquiry View explains level-incoherent tension by appeal to reasons – reasons to inquire further – and denies that there is a rational requirement to be level-coherent. One might think puts my view very close to "reason-responsiveness" or "substantive" views of rationality that deny that there are distinctive requirements of structural rationality, preferring instead to view rationality as primarily a matter of responding correctly to reasons.

Indeed, there are possible accounts of substantive rationality that can agree with my claim that sometimes one's situation makes it rational to have level-incoherent beliefs, though this depends on the details of how those views think one's reasons could be. In fact, many proponents of reasons-responsiveness views of rationality deny that one's reasons could support believing something level-incoherent. For example, Lord (2018) argues that, in fact, our appropriate responsiveness to our reasons always demands that we are coherent (where this includes level-coherence). Many proponents of reasons-responsiveness views regard this as a theoretical advantage, since it can accommodate the intuitive sense that rationality requires level-coherence without appealing to rational requirements (see, for example, Kiesewetter (2017)).

However, this point is typically made in the context of a wider dispute about what the fundamental components of rationality are: reasons or requirements. I am not taking a stand on this issue. Instead, I am interested in what best explains the widespread intuition that level-coherence is something that rationality demands (while remaining neutral on whether it is reasons or requirements that do this demanding). So, in so far as reasons-responsiveness views seek to vindicate the idea that rationality demands level-coherence, I cannot agree with those views. Nevertheless, it remains theoretically possible to defend a reasons-responsiveness view that does not think that one's reasons necessarily demand level-coherence.[xx] The Inquiry View is compatible with such a view, but not identical with it - its focus is different. First, the Inquiry View is a account of the tension arising from level-incoherence, not a view about what the fundamental components of rationality are. Second, the Inquiry View seeks to explain the specific epistemic value of level-incoherence in normatively misleading environments. Third, the Inquiry View is focused on level-incoherence specifically, for reasons outlined in the previous subsection. While there may be scope to combine the Inquiry View with reasons-responsiveness views of rationality, the views are not identical.

5. Summary

In summary, the Inquiry View comes out as the best account of the tension arising from level-incoherence. This is because of the three views, only the Inquiry View is able to capture both what is epistemically valuable about level-incoherence and what has made it appear rationally criticisable.

According to the Inquiry View, level-incoherence provides a reason to undertake further inquiry, but is not rationally criticisable. The tension arising from level-incoherence signals *potential* problems, but does not guarantee them. This permits some cases of rational level-incoherence, which I have argued is particularly valuable for agents in normatively misleading situations. The Prohibition View cannot accommodate this epistemic value, since it is committed to level-incoherence *always* being rationally criticisable. Nor can the Error View accommodate this, since tension tells us nothing about rational criticizability; on this view agents have no reason to investigate. The Inquiry View explains why so many have found the idea that level-incoherence is rationally criticisable intuitive while acknowledging its epistemic value for agents in normatively misleading environments.

References

Adler, J. (2002). Akratic Believing? *Philosophical Studies* 110, 1 – 27. Doi: 10.1023/a:1019823330245

Alexander, D. (2013). The Problem of Respecting Higher-Order Doubt. *Philosophers' Imprint* 13, 1-12. Doi: http://hdl.handle.net/2027/spo.3521354.0013.018

Arpaly, N. (2000). On Acting Rationally Against One's Best Judgment. *Ethics* 110, 488-513. Doi: 10.1086/233321

------- (2003). *Unprincipled Virtue: an inquiry into moral agency.* Oxford University Press.

Arpaly, N. & Schroeder, T. (2013) *In Praise of Desire.* Oxford University Press.

Bennett, J. (1974). The Conscience of Huckleberry Finn. *Philosophy* 49,123 –134. Doi: 10.1017/s0031819100048014

Broome, J. (2013). *Rationality Through Reasoning.* Wiley-Blackwell.

Brown, J. (2018). *Fallibilism: Evidence and Knowledge.* Oxford University Press.

Coates, A. (2012). Rational Epistemic Akrasia. *American Philosophical Quarterly* 49, 113-24.

Comesana, J. & Tal, E. (2015). Evidence of Evidence is Evidence. *Analysis* 75, 557–559. Doi: 10.1093/analys/anv072.

Elga, A. (2010). How to Disagree About How to Disagree. In R. Feldman & T. Warfield (Eds.) *Disagreement (*pp. 175-186). Oxford University Press.

Feldman, R. (2005). Respecting the Evidence. *Philosophical Perspectives* 19, 95–119. Doi: 10.1111/j.1520-8583.2005.00055.x

Field, C. (2019). It's OK to Make Mistakes: Against the Fixed Point Thesis. *Episteme* 16, 175-185. Doi: 10.1017/epi.2017.33

------ (2020). Anti-Exceptionalism About Requirements of Rationality. *Acta Analytica* 36, 423-441. Doi: 10.1007/s12136-020-00450-0

------ (2021). Giving Up the Enkratic Principle. Logos and Episteme 12, 7-28. Doi: 10.5840/logos-episteme20211211

González de Prado, J. (2020), Dispossessing Defeat. *Philosophy and Phenomenological Research* 101: 323-340. Doi: 10.1111/phpr.12593

Harman, E. (2011). Does Moral Ignorance Exculpate? *Ratio* 24, 443–468. Doi: 10.1111/j.1467-9329.2011.00511.x.

— (2015). The Irrelevance of Moral Uncertainty. *Oxford Studies in Metaethics* 10. Doi: 10.1093/acprof:oso/9780198738695.003.0003.

Hawthorne, J.; Isaacs, Y.; & Lasonen-Aarnio, M. (2021). The rationality of epistemic akrasia. *Philosophical Perspectives* 35, 206-228. Doi: 10.1111/phpe.12144

Horowitz, S. (2013). Epistemic Akrasia. *Noûs* 48, 718-744. Doi: 10.1111/nous.12026

Hughes, N. (2022). Who's Afraid of Epistemic Dilemmas? In S. Stapleford, M. Steup, & K: McCain (Eds.) *Epistemic Dilemmas: New Arguments, New Angles.* Routledge.

-------- (forthcoming). Epistemic Dilemmas Defended. In N. Hughes (ed.) *Epistemic Dilemmas*. Oxford University Press.

Johnson-King, Z. (2019). Praiseworthy Motivations. *Noûs* 54, 408–430. Doi: 10.1111/nous.12276

Kiesewetter, B. (2017). *The Normativity of Rationality.* Oxford University Press.

Kitcher, P. (2000a). A Priori Knowledge. In S. Bernecker & F. Dretske (Eds.) *Knowledge: Readings in Contemporary Epistemology.* Oxford University Press.

------ (2000b). A Priori Knowledge Revisited. In P. Boghossian & C. Peacocke (Eds.) *New Essays on the A Priori*. Oxford University Press.
Kolodny, N. (2005). Why Be Rational? *Mind* 114, 509-563. Doi: 10.1093/mind/fzi509
---------- (2007). How Does Coherence Matter? *Proceedings of the Aristotelian Society* 107, 229-263. Doi: 10.1111/j.1467-9264.2007.00220.x
Kvanvig, J. (2014). *Rationality and Reflection: How to Think About What to Think.* Oxford University Press.
Lasonen-Aarnio, M. (2010). Unreasonable Knowledge. *Philosophical Perspectives* 24, pp. 1–21. Doi: 10.1111/j.1520-8583.2010.00183.x
------- (2014). Higher-Order Evidence and the Limits of Defeat. *Philosophy and Phenomenological Research* 88, 314–345. Doi: 10.1111/phpr.12090
------- (2020). Enkrasia or Evidentialism? Learning to Love Mismatch. *Philosophical Studies* 177, 597–632. Doi: 10.1007/s11098- 018-1196-2.
------ (2021). Coherence as Competence. *Episteme* 18 (3):353-376. Doi: 10.1017/epi.2021.33.
------ (forthcoming). *The Good, the Bad, and the Feasible: Knowledge and Reasonable Belief.* Oxford University Press.
Lee, W. (2022). The Real Myth of Coherence". *Erkenntnis* 87, 1211–1230. Doi: 10.1007/s10670-020-00239-y
Lewis, D. (1982). Logic for Equivocators. *Noûs* 16, 431-441. Doi: 10.2307/2216219
Littlejohn, C. (2018). Stop Making Sense? On a Puzzle about Rationality. *Philosophy and Phenomenological Research* 96, 257–272. Doi: 10.1111/phpr.12271
Lord, E. (2018). *The Importance of Being Rational*. Oxford University Press.
Markovits, J. (2014). *Moral Reason*. Oxford University Press.
Nguyen, T. (2020). "Echo Chambers and Epistemic Bubbles". *Episteme* 17, 141–161. Doi: 10.1017/epi.2018.32
Owens, D. (2002). Epistemic Akrasia. *The Monist* 85, 381–397. Doi: 10.5840/monist200285316
Priest, G. (1979). Logic of Paradox. *Journal of Philosophical Logic* 8, 219–241. Doi: 10.1007/bf00258428
-------- (1986). Contradiction, Belief, and Rationality. *Proceedings of the Aristotelian Society* 86, 99-116. Doi: 10.1093/aristotelian/86.1.99
Raz, J. (2005). The Myth of Instrumental Rationality. *Journal of Ethics and Social Philosophy* 1. Doi: 10.26556/jesp.v1i1.1
Reisner, A. (2013). Is the Enkratic Principle a Requirement of Rationality? *Organon F* 20, 436–462.
Roche, W. (2014). Evidence of evidence is evidence under screening-off. Episteme 11, 119-124. Doi: 10.1017/epi.2013.40
Roche, W. & Shogenji, T. (2013). Confirmation, transitivity, and Moore: the Screening-Off Approach. *Philosophical Studies* 3, 1-21. Doi: 10.1007/s11098-013-0161-3
Silva, P. (2018). Explaining Enkratic Asymmetries: Knowledge-First Style. *Philosophical Studies* 175, 2907–2930. Doi: 10.1007/s11098-017-0987-1.
Simion, M. (2023). Resistance to Evidence and the Duty to Believe. *Philosophy and Phenomenological Research*. Doi: 10.1111/phpr.12964
Skipper, M. (2019). Reconciling Enkrasia and Higher-Order Defeat. *Erkenntnis* 84, 1369–1386. Doi: 10.1007/s10670-018-0012-x
Sliwa, P. (2015). Moral Worth and Moral Knowledge. *Philosophy and Phenomenological Research* 93, 393–418. Doi: 10.1111/phpr.12195

Sliwa, P. & Horowitz, S. (2015). Respecting All the Evidence. *Philosophical Studies* 172, 2835–2858. Doi: 10.1007/s11098-015-0446-9
Smithies, D. (2012). Moore's Paradox and the Accessibility of Justification. *Philosophy and Phenomenological Research* 85, 273–300. Doi: 10.1111/j.1933-1592.2011.00506.x
----- (2015). Ideal Rationality and Logical Omniscience. *Synthese* 192, 2769–2793. Doi: 10.1007/s11229-015-0735-z
Staffel, J. (forthcoming). *Unfinished Business*. *Examining our Thoughts in Progress*. Oxford University Press.
Svavarsdottir, S. (1999). Moral Cognitivism and Motivation. *Philosophical Review* 108, 161–219. Doi: 10.2307/2998300.
Tal, E. & Comesana, J. (2017). Is Evidence of Evidence Evidence? *Noûs* 51, 95 – 112. Doi: 10.1111/nous.12101.
Weatherson, B. (2019). *Normative Externalism.* Oxford University Press.
Williamson, T. (2013). Gettier Cases in Epistemic Logic. *Inquiry : An Interdisciplinary Journal of Philosophy* 56. 1–14. Doi: 10.1080/0020174x.2013.775016
-------- (2017). Ambiguous Rationality. *Episteme* 14, 263–274. Doi: 10.1017/epi.2017.24.
Way, J. (2020). A Puzzle About Enkratic Reasoning. *Philosophical Studies* 178 (10):3177-3196. Doi: 10.1007/s11098-020-01575-z
Whiting, D. (2022). *The Range of Reasons.* Oxford University Press.
Worsnip, A. (2018). The Conflict of Evidence and Coherence. *Philosophy and Phenomenological Research* 96, 3–44. Doi: 10.1111/phpr.12246
------- (2021). *Fitting Things Together: Coherence and the Demands of Structural Rationality.* Oxford University Press.
Zimmerman, M. (1997). Moral Responsibility and Ignorance. *Ethics* 107, 410–426. Doi: 10.1086/233742.
— (2008). *Living with Uncertainty: The Moral Significance of Ignorance.* Cambridge University Press

[i] Note that believing that epistemic rationality requires one to φ is not necessarily the same thing as believing that φ-ing is supported by one's evidence, unless one is assuming that evidentialism is the complete theory of epistemic rationality. Put another way, I am concerned with what Silva (2018) calls 'normative' enkratic principles, rather than the 'evidential' enkratic principles that have been frequently discussed in the literature on epistemic level-incoherence (see Horowitz (2013), Lasonen-Aarnio (2014, 2020), Sliwa & Horowitz (2015), Weatherson (2019)).

[ii] I take no stand here on whether rationality is structural or substantive. The difference between the two views lies in what underpins the normative pressure to be level-coherent: structural requirements prohibiting level-incoherence, or something more substantive.

[iii] Note that normatively misleading environments are *always* normatively misleading, even if sometimes the agent somehow manages see through their misleading aspects and believe the truth.

[iv] In previous work, I raised suspicions about the plausibility of this (see Field (2019, 2020)).

[v] Specifically, on her view, Williamson's unmarked clock cases (Williamson 2013).

[vi] Others who defend the idea that cases of apparent rational level-incoherence represent a deep and unresolvable tension in our notions of rationality include Hughes (2022, forthcoming), Williamson (2017).

[vii] Feldman (2005) and Littlejohn (2018) note this as an option but do not endorse it. See also arguments for reasons to be incoherent in the practical case (Broome (2013), Kolodny (2005),

Raz (2005). Kolodny (2007) also argues that there are reasons to be incoherent in the epistemic case, though he departs from most other views in *not* regarding this incoherence as a cost and instead arguing for an error theory of coherence requirements (though, interestingly, not *level*-incoherence requirements (2007: 243)).

[viii] Generating potential examples of rational level-incoherence is a growing cottage industry. See Arpaly (2000), Christensen (2010), Coates (2012), Elga (2007), Feldman (2005), Field (2021), Hawthorne, Isaacs, and Lasonen-Aarnio (2021), Lasonen-Aarnio (2014, 2020), Weatherson (2019), Williamson (2013).

[ix] I take these to be two fairly innocuous assumptions about rationality, though I defend them in more detail in previous work (Field (2019, 2021).

[x] Similar arguments have been made by Lasonen-Aarnio (2020) and Worsnip (2018).

[xi] That is, according to a standard view, though Whiting (2022: 129, 194) argues otherwise.

[xii] I am grateful to Maria Lasonen-Aarnio for pushing me to consider this objection.

[xiii] For further explication of this view, see Priest (1979, 1986).

[xiv] As, for example, Simion (2023) argues.

[xv] There is much more to say here about the normativity of logic and the significance of a priori reasons in favour of basic logical truths. Whether Sam is in a position to reject her first order belief depends on this. However, note that if she *is* in a position to amend her first-order doxastic attitudes in line with her normative beliefs, her problems do not end here. If we think she is rationally required to be level-coherent, and the only available option she has for becoming level-coherent is to give up her belief in a logical truth, then this is some serious epistemic damage. As per §3.1, this is damage that we should want to avoid, and for that reason we should avoid the Prohibition View.

[xvi] I say 'may' because I am not attempting to wade into debates about what, exactly, misleading higher-order evidence supports. I am, however, leaving it as an open possibility that level-incoherent combinations are the correct and rational response to misleading higher-order evidence.

[xvii] Of course, it could be that there *are* mind-independent truths but they change based on what I believe. However, is the requirements of rationality are like this, they would be very different to the kinds of putative requirements of rationality epistemologists are accustomed to thinking about – they would be highly specific, complex, and agent-relative. Indeed, as Lasonen-Aarnio (2014: 332) points out, endorsing an account of epistemic rationality that is like this would mean committing oneself to an error theory about more traditional requirements of epistemic rationality, such as avoiding contradiction, believing what the evidence supports, or even level-incoherence requirements themselves.

[xviii] For agreement on the point that Huck is praiseworthy because he does the right thing for the right reasons, see Harman (2011, 2015), Johnson-King (2019), Markovits (2014), Sliwa (2015), Weatherson (2019). It is important to acknowledge various nuances within this agreement. For example, not all have thought that Huck is fully praiseworthy (Sliwa (2015) argues that he would be even more so had he known what was right), and some have thought there is something morally objectionable to doing what one believes is wrong (see Bennett (1974)). Though, only Zimmerman (1997, 2008) has argued explicitly that Huck would deserve blame.

[xix] Of course, this does not apply to other more general coherence requirements that do not involve logical inconsistency, such as requirements to avoid Moorean absurdity. I remain neutral about whether these are rationally prohibited or could also be thought of as exhibiting useful epistemic tension.

[xx] Indeed, this possibility was suggested to me by Benjamin Kiesewetter (pc).