Rules, Norms and Basic Knowledge

Brian Weatherson, Rutgers/Arché

May 27, 2009

1 Rules and Basicness

Lewis Carroll's 1895 paper "Achilles and the Tortoise" showed that we need a distinction between rules of inference and premises. We cannot, on pain of regress, treat all rules simply as further premises in an argument. But Carroll's paper doesn't say very much about what rules there must be. Indeed, it is consistent with what Carroll says there to think that the only rule is →-elimination. You might think that modern Bayesians, who seem to think that the only rule of inference they need is conditionalisation, have taken just this lesson from Carroll. But obviously nothing in Carroll's argument rules out there being other rules as well.

One connection between rules and basicness should be obvious. Rules were introduced to be regress stoppers. And basic knowledge is supposed to be a regress stopper as well. So we might suspect that if we found out what Carrollian rules there are, we'd learn a lot about what kind of basic knowledge there is.

Another connection concerns the so-called 'Problem of Easy Knowledge'. Consider the way we might teach an intelligent undergraduate that (1) is a theorem.

$$((p \to q) \wedge p) \to q \tag{1}$$

We might first get the student to assume the antecedent, perhaps reminding them of other uses of assumption in everyday reasoning. Then we'll note that both $p \to q$ and p follow from this assumption. If we're ambitious we'll say that they follow by ∧-elimination, though that might be too ambitious. And then we'll draw their attention to the fact that these two claims together imply q, again noting that this rule is called →-elimination if we're aiming high pedagogically. Finally, we'll note that since assuming the antecedent of (1) has let us derive its consequent, then it seems that if the antecedent of (1) holds, so does the consequent. But that's just what (1) says, so (1) is true. For a final flourish, we might note the generalisation of that reasoning to the →-introduction rule. Or we might call it quits at proving (1); undergraduates can only handle so much logic.

That's a pretty good way, I think, of teaching someone that (1) is true, and indeed why it is true. That is, it is a way the student can learn (1), i.e., come to know (1). But note something special about (1). It's obviously closely related to one of the rules we used to prove it, namely →-elimination. Indeed, when Achilles first tries to do without that rule in Carroll's dialogue, he appeals to an instance of (1). Given how closely related the two are, you might think using →-elimination to come to learn (1) would be viciously circular. But in fact it isn't; the argument we just gave is a perfectly good way of learning (1). I don't know whether we should say it is circular but not viciously so, or simply non-circular, but however we classify the argument it isn't a bad argument. So I conclude that when we are using Carrollian rules, we are immune to worries about easy

knowledge. A Carrollian rule saying *From A infer B* is capable of justifying the conditional $A \rightarrow B$. That is, the rule can be used to come to learn that the rule is truth-preserving.

It might be objected at this point that the undergraduate must have already known (1) if they consent to →-elimination. I think this turns on an unrealistic picture of what undergraduates are capable of. It looks to me that (1) is in fact a very complicated proposition. It has a conditional embedded in the antecedent of a conditional! Such sentences are not familiar in everyday life, and most introductory logic students have problems parsing such sentences, let alone knowing they are true. So I don't think this kind of Socratic objection works.

So rules are regress-stoppers that are immune to certain kinds of easy knowledge objections. That means they have a lot of the characteristics we were looking for in types of basic knowledge. (You might think there's a potential category confusion here, since rules aren't propositions, and basic knowledge is. We'll return to this presently.) This is all well and good, but unless we can identify some rules, we have not made a lot of progress. The aim of this paper is to make a new suggestion for how to do this. Let's consider some rules we might find attractive.

Strong-IBE Given some experiences, infer the truth of the best explanation of those experiences.
Weak-IBE Given some experiences, and knowledge that p is the best explanation of those experiences, infer p.
Strong-Perception Given a perceptual state with content p, infer p.
Weak-Perception Given a perceptual state with content p, and no reason to doubt that that perception is reliable, infer p.

All of these rules except the second would, if they were good rules, generate quite a bit of basic knowledge. That's because the 'inputs' to the rule are not the kinds of thing that require (epistemic) justification. And the rule licences inferring certain propositions from those. The way I'm understanding rules, if the agent makes this inference, and there are no defeaters around (the inferred proposition isn't false, or unsafe, or something similar), the inferrential belief constitutes knowledge. And since it is knowledge that doesn't rest on anything that requires justification, it is genuinely basic knowledge. (The second rule clearly does require something justified as an input, so it will be trickier to use it to get basic knowledge.)

One might object that these inferences rest on something that do require justification, namely the rules. But that would be to miss the point that Carroll was making. Rules like →-elimination are regress stoppers; they are unjustified justifiers. If we had to justify them, we'd be led to an absurd regress. Someone who follows these rules in reasoning is justified even if they can't justify the rule. There are deep questions about what it is to follow a rule, as opposed to merely having one's inferential processes correlate with the rule, but I won't get into those debates here since I'll be looking at agents who are clearly following rules.

Strong-Perception is obviously a fairly strong rule, since it applies even in cases where the agent has reasons to believe her perception is unreliable. **Weak-Perception** is intended to be a translation into Carrollian rule talk of Pryor's dogmatism. It says that we are justified in believing the contents of our experiences unless we have a reason to believe that those experiences are non-veridical. In other words, perceptual experience is a regress-stopper.

The aim of this paper will be to suggest a heuristic for distinguishing real rules from non-rules, and to use it to rule out **Strong-Perception** as a rule, and to suggest that **Strong-IBE** is a rule. To get there, we need to

go via some considerations about cases where moral norms clash.

2 Moral Conflicts

Frequently we have conflicts between various moral rules. For instance, it is a moral rule that it is good to rescue children in imminent danger. It is also a rule that it is bad to injure people. Sometimes these can come into conflict, as in the following case.

> **Child**: Alice sees a child walking into traffic, and sees that he is in imminent danger of being hit by an oncoming truck. She jumps out to rescue him, and safely brings him back to the side of the road. To get to the child she had to push Bob out of the way, and Bob subsequently stumbled and mildly sprained his wrist.

This case is only a conflict in a very weak sense. The reason Alice has to rescue the child justifies pushing Bob out of the way, assuming that in doing so Bob wouldn't be seriously injured. Alice hasn't done anything wrong, and although she's harmed Bob, she hasn't *wronged* him. The following case (which isn't original, though I don't remember where I got it from) is a little different.

> **Hiker**: Colin is hiking through the woods when a very severe blizzard whips up. Colin starts to look for shelter, and finds an empty, but locked, hut. It turns out that the hut belongs to Dana, but she is away right now. Colin's only way of avoiding a likely death in the blizzard is to break into the hut, which he does.

On the one hand, self-preservation is good. Indeed, Colin's reason he has to keep himself alive justifies his actions. But we shouldn't forget that trespassing, and damaging property, is also wrong. In damaging, and trespassing in, Dana's hut, Colin wrongs Dana. That's true, I think, even if all things considered Colin did the right thing. For evidence for this last claim, consider what Colin owes Dana. I think he owes her at least an apology, but also compensation for the damage he did to the hut. This makes the case quite different from **Child**. Alice doesn't owe Bob compensation for his injury, and possibly doesn't even owe him an apology.

The asymmetries between Alice's case and Colin's are striking. What we learn from Alice's case is that the rule *Don't injure people* has to have qualifications; it has to be qualified to exclude cases where the injuries are inevitable in a morally justified rescue. We don't learn that in Colin's case. We have to say that there is a rule that Colin broke. If he didn't break any moral rules, it would be hard to explain why he's under an obligation to Dana to apologise to her and compensate her. His case isn't an *exception* to the *Don't damage other people's property* rule. Rather, it is a case where that rule gets *overridden* by other rules, although the rule still applies.

Colin's case is not a strict moral dilemma, at least as I'd use that term. (Almost everyone has a slightly different use of the term 'moral dilemma' so I have to be a little careful with usage here.) There is a single all-things-considered right thing for Colin to do, namely to break into the hut and shelter from the storm. Doing that is not all-things-considered wrong. But he does break a rule in breaking into the hut. Indeed, he can't help but break a rule, since it seems quite irresponsible to simply freeze to death in the blizzard. Famously, there are cases that draw out even deeper conflicts, such as this famous example from Sartre.

> **War**: The invading Nazi army has occupied Emil's country and killed his brother. Emil feels drawn to join the resistance, partially out of a sense of patriotic loyalty, partially out of a felt duty

to avenge his brother, and partially out of well-justified anti-Nazi values. On the other hand, Emil's mother "was living alone with him, deeply afflicted by ... the death of her eldest son, and her one consolation was in [him]. But he, at this moment, had the choice between going to England to join the Free French Forces or of staying near his mother and helping her to live. He fully realised that this woman lived only for him and that his disappearance - or perhaps his death - would plunge her into despair."

This case is frequently credited to Sartre as a moral dilemma. That isn't how Sartre thought of it; he thought it was a case that traditional moral theory couldn't answer, and so much the worse for traditional moral theory. I think Sartre is wrong about that. Traditional theories can tell us that either choice Emil makes will be horrible; he will be doing something that he must not do. A little more moderately, whatever he does will be all-things-considered wrong. It is obviously controversial that this is the right way to understand Emil's case, or indeed any case at all, but it seems conceptually coherent to me to describe it this way. (Some might deny this because they think that there are principles of the form *Ought implies can* or *Must implies can*. I doubt any such principle can survive reflection on cases like **War**.

If we accept what I've said so far, there are three kinds of cases where we have an apparent conflict between moral rules.

No Conflict There is a conflict between two *prima facie* duties, but in fact one of the duties is not a genuine duty because of the existence of the other duty.

Weak Conflict There is a conflict between two duties, but there is a clear right thing to do, even though doing it wrongs another person, or violates a duty the agent has.

Strong Conflict There is a conflict between two duties, and whatever one does, it is the wrong thing to do because it violates a duty one must not violate.

We're interested here in epistemology, not ethics, so we'll be interested in the translations of these into epistemological terms. It won't be important to what follows that there are any Strong Conflicts. But it will be important that there are Weak Conflicts, and in fact I believe some of the cases I describe will be epistemological versions of Strong Conflicts.

3 Moral Ignorance

Among the many things we don't know include many truths about ethics. The existence of many intelligent philosophers promoting false theories doesn't help much. (Since ethicists disagree with one another, I'm pretty sure many of them are saying false things!) There have been several recent works on what the *moral* significance of moral uncertainty is. I'm inclined to think that it isn't very great, although moral uncertainty might have some very odd epistemological consequences in cases like the following.

Kantians: Frances believes that lying is morally permissible when the purpose of the lie is to prevent the recipient of the lie performing a seriously immoral act. In fact she's correct; if you know that someone will commit a seriously immoral act unless you lie, then you should lie. Unfortunately, this belief of Frances's is subsequently undermined when she goes to university and takes courses from brilliant Kantian professors. Frances knows that the reasons her professors advance

for the immorality of lying are much stronger than the reasons she can advance for her earlier moral beliefs. After one particularly brilliant lecture, Frances is at home when a man comes to the door with a large axe. He says he is looking for Frances's flatmate, and plans to kill him, and asks Frances where her flatmate is. If Frances says, "He's at the police station across the road", the axeman will head over there, and be arrested. But that would be a lie. Saying anything else, or saying nothing at all, will put her flatmate at great risk, since in fact he's hiding under a desk six feet behind Frances. What should she do?

That's an easy one! The text says that unless someone will commit a seriously immoral act unless you lie, you should lie. So Frances should lie. The trickier question is what she should believe. I think she should believe that she'd be doing the wrong thing if she lies. So she should do something that she believes, and should believe, is wrong. That's OK; by hypothesis her Kantian professors are wrong about what's right and wrong.

So I think the immediate questions about what Frances should do and believe are easy. But the result is that Frances is quite incoherent. For her to be as she should, she must do something she believes is wrong. You might think it is also a rule of rationality/morality that one should not believe that what one is doing is wrong. I guess I think that isn't really a rule. (I think I'm here following Niko Kolodny, but I'm not sure I've understood the scope of his project correctly.) If this version of *Be coherent!* is a rule, then Frances's case would be an instance of Weak Conflict. Whatever she does, she will violate some rule or other. Either she will violate the moral rule that requires lying when necessary to prevent serious immorality, or the epistemic rule that she should believe that theory that is supported by the overwhelming majority of her epistemic reasons, or this coherence rule. It is an instance of Weak Conflict because it is clear how she should resolve it; she should violate the coherence rule.

When we move from ignorance about moral norms to ignorance about epistemic norms, the conflicts will get more pronounced.

4 Epistemic Ignorance

In his very interesting paper at the Rutgers Epistemology Conference, Higher-Order Evidence, David Christensen discusses a lot of cases where, in the process of investigating whether p, we learn something about our ability to detect whether p. A lot of the next two sections will be inspired by Christensen's paper, though I will draw quite different conclusions than he draws. Indeed, I'll start with a quite different case to his.

Cautious Drug: Gill is investigating a murder. She gets evidence E, and on the basis of that quite reasonably concludes that it is quite likely the butler did it (her credence in that is $\frac{2}{3}$), a serious possibility that the gardener did it (her credence in that is $\frac{1}{4}$), and very little chance that neither did it (her credence in that is $\frac{1}{12}$). These are all extremely reasonable credences given Gill's evidence. Gill is then told, by a source she knows to be usually reliable, that she has taken some drug that leads to people systematically underestimating how strongly their evidence supports various propositions. So if someone's taken this drug, and believes p to degree $\frac{2}{3}$, then p is usually something that's more or less guaranteed to be true by their evidence. To avoid regresses, people who know they have taken the drug are usually able to correct for its effects by increasing their confidences, but people who do not know this are usually excessively cautious. As it turns out,

> Gill has not taken the drug, and her usually reliable source is mistaken, although she is disposed to believe her source on account of their known record of reliability. What changes should Gill make to her credences about the various propositions about who the murderer is on getting this evidence that she's been drugged?

I think that if there is an answer, it is *nothing*. She can't coherently raise her credences in all of these propositions, after all. And if anything was justified in response to learning about the drug, it would be raising her credence in all the propositions. So no reaction to learning about the drug is justified. But she does have good reason to believe she's taken the drug. So she has good reason to believe that she is being irrationally cautious. That's not true, but she has reason to believe it. So what she should believe is different from what she should believe that she should believe. That's bad, though I think it's a little hard to say why it is bad. We'll return to that it a bit, though first I want to deal with an objection.

It might be argued that Gill is in a position to know she hasn't taken the drug, and so she shouldn't believe that her current beliefs are too cautious. That's because she knows that if she had taken the drug the rational credences in the various salient propositions (*The butler did it*, *The gardener did it*, *Someone else did it*) would sum to more than 1. And she is in a position to know that no rational credences sum to more than 1. But I think we should be a little careful about that last assumption. Assuming Gill does not have access to all philosophical truths, she might not know that credences should not sum to more than 1. Even if she previously believed this, she might well now doubt it, given that her state is evidence against it. If she's rational (which she is, and which she has evidence for) and she's been drugged (which she has reason to believe she has been) then it is rational in her case to have credences that sum to more than 1.

Christensen uses slightly different examples, and draws somewhat different conclusions from them. Here are two of his examples. (I've slightly changed the numbers in the second case.)

> **Sleepy Hospital**: I'm a medical resident who diagnoses patients and prescribes appropriate treatment. After diagnosing a particular patient's condition and prescribing certain medications, I'm informed by a nurse that I've been awake for 36 hours. I reduce my confidence in my diagnosis and prescription, pending a careful recheck of my thinking.

> **Tipping**: My friend and I have been going out to dinner for many years. We always tip 20% and divide the bill equally, and we always do the math in our heads. We're quite accurate, but on those occasions on which we've disagreed in the past, we've been right equally often. This evening seems typical, in that I don't feel unusually tired or alert, and neither my friend nor I have had more wine or coffee than usual. I get $42 in my mental calculation, and become quite confident of this answer. But then my friend says she got $45. I dramatically reduce my confidence that $42 is the right answer, and dramatically increase my confidence that $45 is correct, to the point that I have roughly equal confidence in each of the two answers.

In the terms I've been using so far, Christensen thinks that these are cases of Weak Conflict. He thinks the things the narrator does in each case are, all things considered, the right thing to do. But he also thinks they are imperfect solutions. This is easiest to see in **Tipping**. Let's say the bill was $70. Then the narrator's share was $70 ÷ 2, i.e. $35 plus 20% for the tip, so plus $7, so it is $42. The narrator is competent with simple arithmetic,

so he has access to all of this evidence, and not believing something that is clearly entailed by one's evidence is bad, especially when you are trying to figure out whether it is true and have worked through the computation. But Christensen thinks it is even worse to, immodestly, take oneself to be the one who is correct in a dispute like this. His primary motivation for this, I think, comes from the following variant on Sleepy Hospital.

> Or consider the variant where my conclusion concerns the drug dosage for a critical patient, and ask yourself if it would be morally acceptable for me to write the prescription without getting someone else to corroborate my judgment. Insofar as I'm morally obliged to corroborate, it's because the information about my being drugged should lower my confidence in my conclusion. (Christensen, pg 11)

I think the last line here isn't correct. Let's get clearer about the case. The medical evidence suggests that the prescription should be, let's say, 100μg, and that's what the narrator at first intends to prescribe. But the narrator also has evidence that he's unreliable at the moment, since he's been awake so long. Christensen thinks that this evidence is evidence *agains* the claim that the prescription should be 100μg, and the narrator should not believe that the prescription should be 100μg. The alternative view, my view as it happens, is that the narrator should believe that the prescription should be 100μg, although he shouldn't, perhaps, believe that he should believe that. The second conjunct is because he has good reason to think that his actual judgment is clouded, because he has been awake so long. So both views agree this is a case of Weak Conflict, but they differ on whether the appropriate all-things-considered judgment is primarily sensitive to the evidence about reliability, as Christensen thinks, or the ordinary medical evidence, as I think.

Christensen's argument against my position, as I understand it, is that if that were so, then the narrator should make the 100μg prescription. But I think that relies on too simplistic an understanding of the relation between norms of belief and norms of action. After all, my view implies that in believing the prescription should be 100μg, the narrator violates a norm. His beliefs are incoherent in a way, since he believes something he has reason to think he shouldn't believe. And it might be that whenever we have a choice between something that's intuitively an act and something that is intuitively an omission, we should only perform the act if the belief that motivates the act is not in violation of *any* norm. Put another way, when there is a chance to double-check before acting, the doctor should act only if both the medical evidence, and the evidence about the doctor's own reliability, point in the direction of acting. Here's a case that supports that explanation of the case.

> **Cautious Hospital**: A doctor has been on duty for 12 hours. In the world of the story, at that stage in a shift, doctors are typically excessively cautious about their diagnosis. The initial feelings of drowsiness cause them to second-guess themselves even though they are capable of making reliable confident judgments. Helen, a doctor, knows these facts, and has been on duty for 12 hours. Helen is in fact immune to this general tendency of over-caution, though she does not have any prior reason to believe this. She looks at the symptoms of a patient who is in some discomfort, and concludes that probably he should be given 100μg of drug X, although more tests would confirm whether this is really the best action. That's the right reaction to the medical evidence; there are realistic explanations of the symptoms according to which 100μg of X would be harmful, and the tests Helen considers would rule out these explanations. Had she only just come on duty, she

would order the tests, because the risk of harming the patient if the probably correct diagnosis is wrong is too great. But Helen now has reason to worry that if she does this, she is being excessively cautious, and is making the patient suffer unnecessarily. What should Helen do?

If Christensen is right that agents should act on 'higher-order' beliefs about what they should believe, then Helen should prescribe 100μg of X. She should think, "I think probably 100μg of X is the right treatment here, and people in my position are generally a little cautious, so there's excellent reason to hold 100μg is the right treatment. So I'll do that." But that's a horrible breach of good medical practice. She has great evidence that X might be harmful, and general consideration about the credence forming practices of slightly drowsy doctors isn't the kind of thing that could defeat that evidence.

So when the regular medical evidence, and the evidence about our cognitive capacities, point towards different actions, it isn't that we should always do the action suggested by the evidence about our cognitive capacities. Rather, we should do the cautious action, especially if we have a duty of care to the people who will be harmed by the action should it all go awry. So Christensen doesn't have an argument here for taking this 'second-order' evidence to be ruling when it comes to what we should, all-things-considered, do or believe.

5 Explaining Norm Conflicts

I've been stressing so far my disagreements with Christensen, but actually we agree on a lot around here. In particular, we agree that both **Sleepy Hospital** and **Cautious Hospital** are cases of Weak Conflict. And I think we have the same explanation for each case. In both cases the 'first-order' evidence supports a conclusion p. But the agent has reason to believe that the evidence does not support that conclusion, but rather an alternative conclusion q. Given that, if the agent is sufficiently self-aware, she is guaranteed to break one of the following rules.

Evidence You should believe what the evidence supports.
Immodesty You should be able to believe that you're responding to the evidence appropriately.

If the agent has evidence E, which supports believing p to degree x, and she has strong reason to believe that E supports believing p to degree $y >> x$, and she knows more or less what her degree of belief in p is, then she'll have to violate either **Evidence** or **Immodesty**. But both of those are rational norms. And they aren't norms that have ceteris paribus clauses; they are simply norms of belief. So whatever our agent does, she will violate an epistemic rule. Christensen and I both think that there is an all-things-considered right answer to the question of which rule we should violate, though we disagree about what that is.

Christensen suggests, and I agree, that there are also cases of Strong Conflict, though they are rarer than cases of Weak Conflict. The following may be an example.

> **Counter-Indicator**: Ian has very good reason to believe the following proposition: *p is true iff Ian doesn't believe it*. Ian is also very interested in the truth of p. And Ian knows whether or not he believes p. What should Ian believe?

Whatever Ian does, it is seriously irrational. Either he believes p, then has excellent reason to believe that a particular belief of his, namely p, is false. Or he fails to believe p, and hence fails to believe something that

he has excellent reason to believe is true, and which he knows is entailed by things he has excellent reason to believe. Either way, he violates what look like norms of reasoning one must not violate. So, like Emil, he is in a Strong Conflict.

If there are Weak Conflicts, and even Strong Conflicts, we might wonder what explains the existence of these conflicts. In general, the fact that we have some evidence pointing one way, and more evidence pointing another, does not produce a conflict, as in this case.

> **Investigation**: Jack, a detective, is investigating a murder. The signs at the crime scene all point to Kate. She has a motive to want the victim dead, the crime fits her prior methods, and her blood and hair were found at the scene. But Kate has a rather good alibi. At the time known to be at the time of the murder, she was being interviewed by Jack over a previous murder.

Jack has a lot of evidence that Kate was the murderer, but he should not believe that she is the murderer. Unlike in the hospital cases, it isn't as if there is some reason for him to believe Kate is the murderer, and some to believe that she isn't. Rather, Jack's evidence gives him a compelling reason to believe Kate is not the murderer, even though a lot of evidence suggests that she is the murderer.

If epistemic Conflicts exist, and equivocal evidence cannot explain the existence of epistemic Conflicts, then a natural suggestion is that epistemic *rules* explain the existence of Conflicts. Indeed, on a picture of epistemology where all we have are rules and evidence, if evidence can't explain something, rules must.

The hypothesis that Conflict arises because of a clash in rules makes some predictions. In particular, it predicts that we'll see Conflicts in cases where two rules collide. Now it's hard to test this prediction without knowing just what the rules are, but I think the following two principles are either rules or at the very least consequences of rules.

→-elimination If you are investigating whether q, and you know p, and you know $p \to q$, then infer q.

Use Trustworthy Rules Don't use a rule of inference that you have all-things-considered reason to believe is a bad rule.

These two rules clash in cases like the following.

> **Leaping Larry**: Larry is trying to tell whether q is true. He knows, and knows that he knows p. And he knows, and knows that he knows $p \to q$. And he knows that that's all he knows relevant to q. A few weeks ago, he would simply have inferred q. But his logic teacher, McGee, has been presenting him with a lot of reasons that →-elimination is a bad rule. And McGee's arguments that it is bad are much stronger than anything Larry has been able to come up with in reply. Indeed, given the reasons available to him, he has an all-things-considered reason to reject →-elimination. What should Larry believe?

I think Larry faces a Weak Conflict here. He should conclude q. His evidence directly entails q, and refusing to go on and accept it would violate a central rule of inference. But using that rule would violate **Use Trustworthy Rules**. I say that's too bad, the all-things-considered thing to do is to infer q. Others might think that **Use Trustworthy Rules** trumps, and the all-things-considered thing to do is to decline to infer q, even though it is immediately entailed by things Larry knows. That actually doesn't matter a lot here; what does matter is that there is a Conflict.

6 Discovering Rules

If I'm right that rules clashes are the best explanation of Conflicts, then we can use the existence, or non-existence, of Conflicts, to diagnose whether there are rules of a certain kind. For instance, the following vignette suggests that **Strong-IBE** is a rule.

> **After the Flood**: The basement in Nora's house frequently floods after heavy rains, but fortunately she has a pump to pump out the water. Unfortunately, the pump often fails. It rained heavily yesterday, and when Nora is leaving for class she notices that the outlet pipe attached to her pump is dry. The pipe is usually damp if the pump has been extracting water, but it is dry if the pump is broken. The best explanation of her observations is that the pump has failed again. Nora has been taking some philosophy classes, and they have been discussing Hume and Popper's arguments against induction. The sceptical arguments presented in class seem very strong to her; in fact she doesn't have any idea what could count as a good response to these arguments. What should Nora believe?

I think Nora faces a Conflict here. On the one hand, the best explanation of her observations, and her background knowledge of the facts about her pump, the rain, etc, is that the pump has failed. On the other hand, given her philosophy classes, using IBE would violate **Use Trustworthy Rules**. If she concludes the pump has failed, as she should given her evidence, then reflected on her belief-forming practices, she would conclude that she'd done something irrational. That is epistemically regrettable, even if it isn't all-things-considered wrong. So I think Nora faces an epistemic Conflict here, and the best explanation of that Conflict is that there is a clash between two rules, one of them being **Strong-IBE**.

On the other hand, we don't get an argument for **Strong-Perception** being a rule. Indeed, we don't see Conflicts there where we'd expect to see them if **Strong-Perception** were a rule, as in cases like this one.

> **Poor Lighting**: Oscar is looking at a bright red wall, and having a clear and distinct impression of redness. He is told by Pete that the reason the wall looks red is that it has a very bright red light shining on it. This is false, the wall is in fact red. But Pete has been a reliable testifier for a long time, and the wall does, due to the freshness of the paint, have a kind of shininess that Oscar usually associates with walls being illuminated by red light. What should Oscar believe?

That's easy; he should withhold judgment until he he can investigate whether Pete is right that there's a red light. Even though he does directly perceive the redness of the wall, he has good reason to believe his perception is mistaken in this case. And in such cases we should suspect that we are being deceived. This case feels less like cases of genuine Conflict, like **Leaping Larry**, and more like cases where we have clashes between different evidential sources, like **Investigation**. But if **Strong-Perception** were a rule, then it would be a case of rule clashes. So **Strong-Perception** is not a rule.

It would be nice if we could use this method to judge whether **Weak-Perception** is a rule. But the method I've been using here won't extend to those cases. That's because the method relies on case where the agent does have reason to believe that perception is unreliable. The closest I can get to a case where **Weak-Perception** generates a Conflict is the following.

> **Demon**: Quinn is in the stands at Liverpool, and sees Steven Gerrard score a goal. He immediately forms the belief that Gerrard just scored a goal. But Quinn has been taking a lot of epistemology classes, and given the philosophy he knows, he has good reason to doubt the accuracy of visual perception. He realises that Descartes' arguments for the reliability of perception are unsound, but doesn't have any idea about what could replace them. On reflection about his current experiences, he isn't sure what to believe. What should he believe?

If (a) this is not a case where the qualification on **Weak-Perception** applies, and (b) not a case of Conflict, then we'd have an argument against **Weak-Perception**. But I'm not sure whether either of these are true. So the question of whether **Weak-Perception** is a genuine rule of inference is a question for another day.