

# JUSTIFIED BELIEVING IS TRACKING YOUR EVIDENTIAL COMMITMENTS

Barry LAM

**ABSTRACT:** In this paper, I give an account of the conditions for rationally changing your beliefs that respects three constraints; 1) that rational believing is a matter of respecting your evidence, 2) that evidence seems to have both objective and subjective features, and (3) that our set of beliefs seem to rationally commit us to certain propositions, regardless of the evidential support we have for these propositions. On the view I outline, rationally believing or giving up a belief is a matter of your inferences tracking your rational commitments, and that these rational commitments account for the evidence you must respect. These rational commitments are subjective in that they are relative to the totality of your beliefs, but also objective in the sense that what counts as a commitment is true for everyone everywhere.

**KEYWORDS:** justification, rationality, rational commitments, rational belief-revision

## Introduction

In this paper, I present three problems in the theory of justification: the problem of giving a unified explanation of when we are justified in changing our beliefs, the problem of the objectivity versus subjectivity of evidence, and the problem of giving an account of the normative demands of rational commitments. I outline a unified solution to these three problems according to which justified changes of belief are inferences that are sensitive to all and only pieces of information you are committed to as being evidence for or against *P*. I motivate and state a series of rules that determine when you are committed to a piece of information being evidence for a proposition *P*, and show how the view makes sense of our competing intuitions about the objectivity and subjectivity of evidence.

## 1. Three Questions in the Theory of Justification

In this first section, I present three questions in the theory of justification that I take to be of interest independent of issues concerning knowledge. In the remainder of the paper, I present a theory that attempts to give a unified answer to all three questions.

### 1.1. Justified Belief-Change

Justification was of concern in epistemology because justified belief was long thought to be a component of knowledge. However, beliefs are not the only kinds of things for which we can have epistemic justification. For instance, you can have justification to give up a belief without thereby having justification to believe its negation. In addition to justification to adopt and abandon beliefs, you can have justification to strengthen as well as weaken them. In general, you can have justified *changes* of mind like formation and revision of beliefs in addition to having justified *states* of mind like belief or disbelief. You can even have justification to change the grounds on which you believe that P without having justification to change the belief that P, and without having any change in the justificatory status of your belief. A theory of justification ought to give a unified explanation of the factors that justify changes of mind generally, which includes formation of beliefs. Here we have the first problem of justification at the heart of this paper. Is there a unified explanation of the factors that give us justification to change our minds?

This question is largely independent of questions about knowledge. If I have sufficient justification to give up a belief that P, and I do so, there is no question about whether the doxastic state of mind that results is knowledge, or a component of knowledge. I do not even end up believing that P, so how can I know it? Nor does knowledge enter into the explanation of my justification to give up my belief. It cannot be that failing to know that P is what gives someone sufficient justification to give up a belief that P. If merely lacking knowledge was sufficient to justify the abandonment of belief, then there could be no cases of justified true belief that are not knowledge, for all such cases would be cases where agents ought to give up their beliefs. Knowledge simply is not part of the explanation for the justification of belief abandonment. But yet, identifying the conditions that sufficiently justify me to give up a belief is still of interest even when it has no straightforward link to knowledge. And there does not seem to be anything particularly special about giving up a belief. Forming a belief is just another case of belief-revision. Thus, a unified explanation of justified belief-revision is still of interest even without understanding justification as a key difference-maker to knowledge, or as something that is explained in terms of knowledge. Justification today is as relevant as it was yesterday, even without a straightforward link to knowledge.

## 1.2. Subjective versus Objective Accounts of Evidence

One way to begin understanding justification is in terms of evidence. You have justification to change your belief in a certain way when the evidence you possess points toward or away from the truth of the proposition you believe. But what is it to ‘point’ toward or away from the truth of a proposition you believe? In the theory of probability, evidence for a proposition P is any proposition Q that raises the probability of P. In other words, the conditional probability of P given Q is greater than the unconditional probability of P, provided that P has some initial positive probability.<sup>1, 2</sup> Non-evidential information regarding P is a proposition Q that does not have this feature. This is the general structure of the probabilistic relationships that make one proposition evidence for another. The philosophical substance comes from what we mean by a *probability*. Probabilities might be things out in the world like objective chances, propensities, or dispositions. Probability might instead be subjective degrees of belief, or credences in the mind of a subject. Granting that there are such things as chances, propensities, and credences, exactly which of these things is *the thing* whose structure determines what is evidence and non-evidence for us? The answer tells us whether evidence for a proposition depends on the thoughts of a reasoning subject, or on objective conditional and unconditional features of the world. In other words, to answer the question about probabilities is to give a view about the subjectivity or objectivity of evidence.

The problem is difficult, for we have intuitions that pull us in opposing ways. Consider two cases offered by Alvin Plantinga pointing toward an objective view of evidence.<sup>3</sup>

In case one, person A is captured and altered by Alpha Centauri scientists to have a conditional degree of belief of the Earth being flat, given that there is a

---

<sup>1</sup> In this paper, I am working with this one particular formal characterization of evidence. But many others that are not defined solely in terms of the relationship between conditional and unconditional probabilities are consistent with what I say in this paper. The only important matter for the purposes of this paper is that the structure of certain probabilities on Q and P is what makes Q evidence for P. See Colin Howson and Peter Urbach, *Scientific Reasoning: The Bayesian Approach* (La Salle: Open Court, 1993), chapter 6.

<sup>2</sup> How we should understand a conditional probability is controversial. Alan Hajek argues in “What conditional probability could not be,” *Synthese* 137 (2003): 273-323, that we should understand it as a basic doxastic state. Bas van Fraassen, in “Belief and the Will,” *Journal of Philosophy* 81, 5 (1984): 235-256, prefers to think of it as the comparative likelihood of A and B to A and not B. I am going to take the concept as basic.

<sup>3</sup> Alvin Plantinga, *Warrant: The Current Debate* (New York: Oxford University Press, 1993).

Barry Lam

picture of a round Earth taken from outer space, to be higher than her unconditional degree of belief that the Earth is flat. According to Plantinga,

it is not the case then that the evidence for the earth's being round really does support 'for [A]' the proposition that it is flat-just as it is not the case that the earth is flat 'for [A]'. If E supports proposition H, then E supports H *simpliciter*, not merely relative to your credence function or mine.<sup>4</sup>

For Plantinga, a picture of a round Earth from outer space is evidence that the Earth is round regardless of the credence function that *anyone* might have.<sup>5</sup>

In case two, Plantinga asks us to consider counter-induction. Person A currently believes that it is as likely as not that John can swim. But A also believes it is very probable that John *can* swim given that he is a member of a group in which 99 out of a hundred members *cannot* swim. A also believes that it is very improbable that John can swim on the assumption that John is a member of a group 99 out of a hundred members of which *can* swim. Clearly A *takes* membership in a group of mostly swimmers to be evidence that a person cannot swim, and A *takes* membership in a group of mostly non-swimmers to be evidence that a person can swim. But, according to Plantinga, regardless of what A 'takes' to be evidence or not, it is clearly a fact that membership in a group of mostly swimmers is evidence *simpliciter* that a person can swim, and membership in a group of mostly non-swimmers is evidence that a person cannot swim. Thus what is evidence for a person is not what she takes to be evidence. What she takes to be evidence is not even 'evidence for her' in any meaningful sense of 'evidence' any more than a proposition is 'true for her' in any meaningful sense of 'true.' At most, these are idioms for A taking something as evidence, and A believing something. But taking something to be evidence, by way of your personal credences, doesn't make it evidence, any more than believing a proposition makes it true, according to Plantinga. Plantinga is pointing here to a very strong intuition that the evidence-for, or support relation is a *subject-independent*, objective one that holds between propositions.<sup>6</sup>

---

<sup>4</sup> Plantinga, *Warrant: The Current Debate*, 131.

<sup>5</sup> Plantinga, *Warrant: The Current Debate*, 131.

<sup>6</sup> Such views are not limited to externalists about justification like Plantinga or Timothy Williamson, *Knowledge and its Limits* (New York: Oxford University Press, 2000). Susan Haack for instance, has a view about evidence like Plantinga's in Susan Haack, *Evidence and Inquiry: Towards Reconstruction in Epistemology* (Oxford: Blackwell, 1993). Also see Michael Titelbaum, "Not Enough There There: Evidence, Reasons, and Language Independence," *Philosophical Perspectives* 24 (2009): 477-528, for an extended discussion relating the issue back to Carnap.

On the other hand, other familiar kinds of cases appear to pull us toward a subject-dependent conception of evidence. Imagine that after long clinical studies of the relationship between cholesterol consumption and heart disease, it is discovered in a lab in Sydney that the development of heart disease has no correlation whatsoever with dietary cholesterol consumption. In viewing evidence-ascriptions akin to truth-ascriptions, the subject-independent view will say that the sentence “cholesterol consumption is not evidence of increased risk in heart disease” is true. Now suppose that you and I are in this Sydney lab, watching a television monitor of the American Dr. Spock asking Freddy about his diet. Upon hearing that Freddy eats nothing but egg-yolks deep fried in clarified butter, Dr. Spock concludes that Freddy is likely to develop heart disease. Dr. Spock does so because his conditional probability that Freddy will get heart disease, conditional on Freddy’s high dietary cholesterol intake, is high. We would conclude that Dr. Spock is engaged in a fully justified piece of reasoning. Assuming that fully justified pieces of reasoning respect evidence, we are led to conclude that cholesterol consumption is evidence (for Dr. Spock) that that Freddy will get heart disease. Thus, evidence is subjective, or subject-dependent. This line of reasoning is difficult to dismiss on the grounds that Dr. Spock is still guilty of being unjustified because there is a sense of ‘unjustified,’ an objective sense, which applies to Dr. Spock. It seems no more unjustified of Dr. Spock to draw his conclusion as it is to say that people are unjustified in forming false beliefs on the grounds that they are false.

Familiar thought-experiments concerning the justification of internal duplicates also pull against subject-independent views.<sup>7</sup> Imagine that you and I are talking about A who inhabits a world in which an evil demon does the following; whenever a certain person is reasoning about a general population and samples that population randomly, the evil demon makes it such that no one but those two randomly sampled members have the feature one is inquiring about. According to the subject-independent view, the sentence spoken in our context, “In that world, random sampling and discovering feature F is not evidence that the general population tends to have feature F” is true. It is also objectively true in that world that random sampling of a population and finding feature F lowers the objective probability that members of that population generally have feature F. So according to the subject-independent view, in that world, “random sampling of a member of a population is not evidence of the features of the general population” is false. Imagine that A, however, is an internal duplicate of B, who in this world truly

---

<sup>7</sup> Earl Connee and Richard Feldman, “Internalism Defended,” in *Epistemology: Internalism and Externalism*, ed. Hilary Kornblith (Oxford: Blackwell, 2001), 231-260.

takes random sampling to be evidence that the general population tends to have feature F. Clearly if B is fully justified in so reasoning, so is A. Indeed, in that world, A is fully justified in reasoning the way she does. And since justified reasoning must be fully evidence-respecting, both A and B seemed to have respected their evidence. Thus, evidence is subject-dependent.

In both of the above cases, the objective matters of fact that raise and lower objective chances are not the things that Dr. Spock and A's reasoning reflects. Rather, they reflect Dr. Spock and A's subjective opinions about what is evidence for what, and yet both are fully justified.

In giving an account of evidence, we are pulled in two competing directions. Plantinga is certainly right in that people can disrespect their evidence even when they don't take certain things to be evidence, due to cognitive malfunction, indoctrination, or otherwise. Yet familiar cases of justified reasoning which involve reasoning contrary to facts about causation, correlation, or objective chances seem to pull toward subject-dependent views. We are now familiar with the distinction between subjective and objective justification.<sup>8</sup> Perhaps there are just two irreducible senses of 'evidence,' one that is subject-independent, and one that is subject-dependent. If at all possible, it would be preferable to have a univocal analysis of evidence. Is there a unified account of the kind of information a reasoner must respect in virtue of which her reasoning is justified? Is this information objective (or subject-independent) or subjective (or dependent on the noetic structure and doxastic attitudes of the reasoner)? This is the second problem of justification.

### 1.3. Rational Commitments

A prominent feature of philosophical argumentation involves claiming that a proponent of a certain view is rationally committed to some other view. This kind of rational commitment is epistemic in nature, and it is also normative. There is some kind of normative demand made of a subject when they are rationally committed to a certain view which they in fact do not hold, or actively reject. Yet, as prominent a role as rational commitments play in philosophy, there have been few examinations of the topic as such. Perhaps there is a presumption that rational commitments can be subsumed under concepts in the neighborhood, like the

---

<sup>8</sup> The distinction is found in, for instance, Jonathan Kvanvig, "Subjective Justification," *Mind* 93, 369 (1984): 71-84, Richard Feldman, "Subjective and Objective Justification in Ethics and Epistemology," *The Monist* 71, 33 (1988): 405-419, and Bruce Russell, "Epistemic and moral duty," in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, ed. Mathias Steup (New York: Oxford University Press, 2001), 34-48.

concept of having justification to believe a proposition. But justification and rational commitment, while related, are not obviously identical. To be rationally committed to a view is not to have sufficient justification to accept it, or to be required to accept it on pain of irrationality. It does not make sense to say that someone who has inconsistent beliefs is justified in believing a contradiction. But saying that a certain set of beliefs rationally commits someone to a contradiction is quite sensible. If you hold a view which entails that there are infinitely many unicorns, I might claim that you should not believe that there are infinitely many unicorns, while at the same time claiming that you are rationally committed to them. To be rationally committed is not the same as being justified in believing, and not the same as being obliged to believe on pain of irrationality.

When we claim that A is committed to P, we are identifying the proposition that P as something to which A bears some kind of epistemic responsibility in her reasoning. But exactly what must A do to carry out such responsibilities? And to which propositions are we rationally committed? We have a series of open questions regarding rational commitments, and the relationship of this concept to epistemic justification. Giving an explanation of the nature and extension of our rational commitments is the third problem in the theory of justification.

## **2. Justified Belief-Revision and Information-Tracking**

Now that we have raised the three questions about justification of interest in this paper, I will now begin building a theory that gives a unified answer to all three, beginning with the problem of justified belief-change.

What separates justified from unjustified belief-revision? Consider the following case, modeled loosely on cases from the empirical study of human reasoning. Two agents, A and B, are presented with testimonial evidence as to whether Jones was drunk on a given night, where one piece of testimony contains information suggesting that she was, and another piece of testimony contains information suggesting that she wasn't. One piece of testimony states that Jones managed to recite Keats and agilely dodge a dart thrown at her on the night in question. The other piece states that she stumbled and knocked a table and its contents over. For A, the vividness and colorfulness of the testimonial details, even when such details are not relevant to the question as to whether Jones was drunk, make a difference to what she concludes about Jones. For instance, when A is given the testimony exactly as I have worded it above, A concludes, given all of the other information, a certain degree of confidence in Jones being drunk. When presented with the same information about Jones stumbling over, but one adds the

detail that Jones stumbled over a *quarter-sawn white oak* table and knocked a *glass bowl of neon green guacamole* onto *curly-shag carpeting*, A concludes to a greater degree of confidence that Jones was drunk. Similarly, if only the other piece of testimony is made vivid and colorful, for instance, that Jones recited a Keats poem *while 14 blondes and three Dobermans* looked on, and dodged a *gold-plated 19th-century African dart* thrown at her, A concludes to a lower degree of confidence that Jones was drunk.

In contrast, B's inferences from the testimony is invariant across the vividness and colorfulness of such details.<sup>9</sup> A is reasoning in an imperfectly rational manner.<sup>10</sup> Comparatively, B's reasoning seems much more rational than A. A's reasoning about Jones is sensitive to irrelevant information in a way that B's reasoning is not. Yet its presence makes a difference to what A concludes about that very matter. Thus, sensitivity of a piece of reasoning to irrelevant information makes that piece of reasoning less than fully rational.

Reasoning can also be imperfectly rational when it fails to be fully sensitive to relevant information. Suppose A and B are told that Smith and Wesson are two people drawn from the population of Wasilla College students. A and B both observe Smith and Wesson drink heavily and break their beer bottles onto the ground in glee, leaving shards of glass on the sidewalk. Suppose that A and B are asked to give their opinion about the degree to which they believe the population of Wasilla College as a whole generally drink heavily and behave irresponsibly. In Case 1, A and B are told that Smith and Wesson were chosen at random from the population. In Case 2, A and B are told that Smith and Wesson were chosen precisely for their partying ways. A's opinion about the character of the population of Wasilla College students as a whole is invariant between Case 1 and Case 2, whereas B believes the same as in A in Case 1, but to a far less degree that the whole population of Wasilla College students drink heavily and behave irresponsibly in Case 2. A's inference in at least one of these cases seems to be insensitive to relevant information, information about the representativeness of Smith and Wesson of the population they are chosen from. As a result, A's inferences are less reasonable than B's, and thus, imperfectly rational.

---

<sup>9</sup> These examples are exaggerated. For empirical studies suggesting some similar situations in which people appear to place inappropriate or undue weight on vivid and colorful information, see Richard Nisbett and Lee Ross, *Human Inference: Strategies and Shortcomings of Social Judgment* (Englewood Cliffs: Prentice Hall, 1980), chapters 3 and 4.

<sup>10</sup> Throughout the paper, I am treating 'justified' and 'rational' as picking out the same kind of normative epistemic quality. Sometimes, 'rational' as a word seems to work better with qualifiers like 'perfectly' or 'imperfectly,' and in such cases I will use 'rational.' There are other senses of the words in which the two come apart. But I will use them synonymously.



From the cases above, it looks as though sensitivity to all and only relevant information in making up or changing your mind about what is the case is the essential feature that makes for justified belief-change. When you reason in a way that is sensitive to irrelevant information, you are less rational or justified in changing your mind. When you reason in a way that is insensitive to relevant information, you are less rational or justified in changing your mind.

Let us make the account more precise. Let an *inference-type* be a way in which a person reasons, categorized at some appropriate level of generality. In the empirical literature, what goes by the heading 'heuristics and biases' like the patterns I have discussed above are usually a list of inference-types. 'Induction' is a famous type of inference, as are perhaps sub-types of induction like "reasoning from particular facts about people to generalizations about them" or "reasoning from particular facts to generalizations about populations."<sup>11</sup> In the presence of irrelevant information, an inference-type which generates a particular change of belief is sensitive to such information just in case what it generates in the absence of such information, in nearby similar situations, differs from what it generates in actuality. In the presence of non-evidential information, an inference-type is insensitive to that information just in case what it generates, in nearby similar situations, is invariant across the presence and absence of that information. In the presence of certain *relevant* information, an inference-type is insensitive to that information just in case what it generates, in nearby similar situations, is invariant across the presence and absence of that information. In the presence of certain *relevant* information, an inference-type is sensitive to that information just in case what it generates, in nearby similar situations, differs in the absence of that information. This notion of sensitivity is due to Nozick who explains knowledge in terms of a belief being sensitive to truth.<sup>12</sup> I am using it to explain justified belief-changes in terms of relevant information. Like Nozick, I will use 'tracking' as the term for when your inferences are fully sensitive to relevant information.

In the presence of a body of information, a particular inference about whether P will be fully justified only if it is an instance of a type that is fully sensitive to, or fully tracks, all and only the information in that body that is relevant to whether P. This characterization does not identify types of inferences as being justified or unjustified. Our friend A above can employ one and the same inferential-type across two situations, and be perfectly justified in one, but not in

---

<sup>11</sup> Let us set aside for now the famous problem of what counts as an appropriate level of generality. See Earl Connee and Richard Feldman, "The Generality Problem for Reliabilism," *Philosophical Studies* 89, 1 (1998): 1-29.

<sup>12</sup> Robert Nozick, *Philosophical Explanations* (Cambridge: Belknap Press, 1981).

the other, due completely to the fact that irrelevant information is absent in one but present in the other. Similarly, A can employ an inferential-type across two situations, and be unreasonable in one but fully reasonable in the other, due completely to the fact that certain evidential information is present in one but absent in another. It is incorrect to say that 'induction' or 'the availability heuristic' is an unjustified form of reasoning. Rather, certain instances of them can be unjustified, while others can be perfectly justified.

### **3. Commitments as Justifiers**

Justified belief-revision must track all and only relevant information. There is a reading of this view that is simply trivial. What is 'relevance' except 'justification-giving'? This makes justified belief-revision a trivial matter of belief-revision that is sensitive to justification-giving factors! The move to 'evidence' talk does not help. Almost by definition, information that is relevant to whether P is evidence for or against P, where this can simply be defined in terms of conditional probabilities as in section 1.2. But we still need a solution to our second problem so as not to collapse into triviality. Is the proper characterization of probability subjective or objective? In other words, does our reasoning need to respect what our subjective probabilities say is relevant, or what is in fact relevant independent of our subjective probabilities? If the answer is our subjective probabilities, then we have something non-trivial; what is relevant is whatever it is a subject takes to be relevant. However, recall that in response to intuitive cases, our intuitions pull two different ways. How can the evidence that rationality demands we respect in our reasoning be, at the same time, independent of our personal conditional probabilities, or what we 'think' is evidence for what, and also dependent on our personal probabilities?

Rational commitments, I submit, are things which have all the right features for being those things which characterize evidence. A rational commitment depends on what you believe. Two people who believe different things may be committed to different things. Commitments are therefore subject-dependent. But a commitment need not itself be something that you believe. You can be committed to Q being evidence for P without thereby also believing that Q is evidence for P. In this way, what you are committed to is not necessarily something that you yourself may believe or have any opinions about. And if your rational commitments make up the extension of the probabilities that determine what information is evidentially relevant to P for you, then it is possible that Q is evidence for P for you even if you do not 'think' it so. At the same time, what is evidence for you still depends on the totality of your beliefs and subjective

probabilities. Rational commitments have subjective and objective features that make them prime candidates for determining what is evidence for you. And while the notion of ‘commitment’ is still a normative one, it is not the same as the notion of justification. Therefore, saying that what counts as relevant and irrelevant information is a matter of your rational commitments is nontrivial. This is the theory I will articulate in this section. In the next section, I will show how it helps to settle the intuitions concerning objective and subjective evidence.

According to what I will call the Commitment-Tracking Theory, given the set of actual conditional and unconditional credences held by a person, let’s call it  $pr_{@}^{(*)}$ , there is a set of possible degrees of belief for a person, call it  $pr_{\odot}^{(*)}$ , whose structure determines what is evidentially relevant for her. That is, whatever A’s personal credences in P ( $pr_{@}(P)$ ) or Q ( $pr_{@}(Q)$ ), P is evidence for Q for A just in case  $pr_{\odot}(P|Q) > pr_{\odot}(P)$  provided  $pr_{\odot}(P) > 0$ . In other words, the structure of the probabilities on P and Q in the set of possible degrees of belief to which an agent is *committed* is what determines the evidential relevance of Q to P. Yet, what is in the set of commitments will depend on the subject’s actual credences together with whatever set of rules generate the set of commitments of A. On this view, two people with difference subjective credences can differ in terms of what is evidence for what for them. However, a subject’s own actual subjective credences, conditional or unconditional, do not exhaust what is evidence for what for her.

### 3.1. Rules of Commitment

What we require is a motivated set of rules that determine a subject’s commitments. A logical consequence of a theory or set of beliefs is considered a commitment of that theory or subject of those beliefs.<sup>13</sup> This logical consequence conception of commitment extends to the probabilistic case as well. A certain weather forecaster may believe that rain is twice as likely as no precipitation, and that given any further drop in barometric pressure before midnight, it would be thrice as likely as no precipitation. Imagine that the following three other degrees of belief are not ones he currently possesses: (1) that a drop in barometric pressure makes rain half as likely as no drop, (2) that it is twice as likely to snow than rain,

---

<sup>13</sup> This is even accepted by the most skeptical of philosophers about any relationship between logic and justified reasoning like Gilbert Harman, “Internal Critique: A Logic is not a Theory of Reasoning and a Theory of Reasoning is not a Logic,” in *Handbook of the Logic of Argument and Inference: The Turn Towards the Practical, Volume 1 in Studies in Logic and Practical Reasoning*, eds. D.M. Gabbay, R.H. Johnson, H.J. Ohlbach, and J. Woods (Amsterdam: Elsevier Science B.V., 2002), 171-186.

(3) given no precipitation tomorrow, it is likely that there will be no drop in barometric pressure after midnight. Some of these possible degrees of belief, (1), are inconsistent with an agent's actual degrees of belief. Some of them, (2), are consistent with them, and some of them, (3), are ones to which an agent may not possess, but is nonetheless *committed*.

These notions of commitment are closely related to the notion of logical and probabilistic coherence.<sup>14</sup> But our commitments extend beyond formal coherence. Rational commitments appear to be those propositions, probabilities, and conditional judgments that make maximal sense of all your beliefs and judgments. Consider the following non-epistemic case. It is consistent for the park to prohibit skateboarding on the grounds that it causes costly damage to the concrete. Now imagine that stunt bicycling causes similar costly damage to concrete and is just as prevalent. The park is committed to prohibiting stunt bicycling also. Now it is perfectly consistent to prohibit one but not the other. However, when you do things on certain grounds (and only on those grounds), you are committed to coherently applying those grounds to new cases unless there are justifiable exceptions. The same is true in the epistemic case. If you believe that P on certain testimonial grounds, and only on those testimonial grounds, then you are committed to any other proposition Q that rests on those and only those exact same grounds. I call this principle the Rule of Parity.

### Rule of Parity

The Rule of Parity is not limited to commitments to believe certain claims. You are also committed to things being evidence of other things. If Repub thinks that John Kerry's past 'flip-flopping' is evidence of Kerry's untrustworthiness, where what this means is that Repub has the right structure of subjective conditional and unconditional probabilities on these propositions, then he is committed to Mitt Romney's 'flip-flopping' to be evidence of Romney's untrustworthiness. This is a commitment that is overridden only if such an agent has significant enough opinions about how Romney is an exception. Conversely, if for Demo, Kerry's past 'flip-flopping' is irrelevant to Kerry's trustworthiness, then Demo is committed to

---

<sup>14</sup> These are prominently featured in Laurence Bonjour, *The Structure of Empirical Knowledge* (Cambridge: Harvard University Press, 1985), chapters 5 and 6. But the notion is not identical to coherence. As we have already seen, it seems to be quite a robust judgment that people with inconsistent beliefs are committed to contradictions, even though a contradiction in no way makes maximal sense of the things you believe. I will discuss the issue of inconsistency more in section 5.

Romney's 'flip-flopping' as being irrelevant to Romney's trustworthiness, unless such an agent has significant enough opinions about how Kerry is an exception.

Let us apply the Rule of Parity to a case we've already seen. An agent A knows that an urn contains some colored balls, and that she must decide the percentage of balls that are red. She knows that Jones will draw two balls at random without replacement. Jones does so, and both balls are red. A believes that the likelihood that all of the balls are red, given that Jones draws two consecutive red balls at random, is much higher than the likelihood that all of the balls are red, given that Jones draws two consecutive red balls specifically because they are red. A's conditional probability reveal that she takes random sampling to be stronger evidence of the properties of a general population than targeted sampling.

The Rule of Parity states that if randomness or non-randomness counts as evidence for agent A in the case of urns with red balls, then agent A is committed to it being evidence for general propositions about all things and their features, unless A has sufficient opinions about how a certain general proposition must be an exception. Specifically, since taking a sampling of Wasilla College students to indicate some features of the general population is no different from taking a sampling of balls to indicate the features of a general population of balls, the features of the sampling method count as evidence in both cases.

### Rule of Integrity

Suppose a subject has conditional probabilities such that Q is not taken to be evidence for P. At the same time however, such a subject evaluates or is disposed to evaluate anyone positively or negatively according to a norm in which Q is evidence for P. Such evaluations commit the agent to Q being evidence for P. For instance, if you do not have an opinion that skin color is evidence of criminal dispositions, but you evaluate certain people negatively for not thinking some person is likelier to be a criminal on the basis of their skin color, then you are committed to skin color being evidence of criminal disposition. I call this the *Rule of Integrity*, which states that if you evaluate yourself and others according to a set of epistemic standards, then you are committed to the doxastic states that make sense of those standards. Any conditional probability whose presupposition explains an epistemic evaluation an agent makes, or would make given the agent's epistemic standards, is one to which she is committed.

### Rule of Belief

The *Rule of Belief* states that one is committed to any doxastic state that one possesses. This captures the idea that if one already takes something to be evidence

Barry Lam

and something as non-evidence, it counts as something that the subject's inferences ought to respect or ignore. For instance, imagine that A believes that rain is likelier given a drop in barometric pressure than with no drop. Yet upon viewing nothing but a drop in barometric pressure, A actually lowers his opinion about the likelihood of rain. His inference here is imperfectly rational, since it fails to respect what is evidence by his own lights, in a very straightforward way. Notice that the same is true even if A mistakenly believes, because of his incompetent meteorology professor, that rain is likelier given a *rise* in barometric pressure. Seeing nothing but a rise in barometric pressure, A ends up decreasing his confidence in rain. A's line of reasoning would still be imperfectly rational in the same way.

### Rule of (Probabilistic) Coherence

Let me end with a formal characterization of the first rule I discussed in this section, the *Rule of Coherence*. According to this rule, any degree of belief, conditional or otherwise, which is coherence-contributing to an agent's existing degrees of belief is one to which she is committed. A certain possible degree of belief may be one whose adoption would increase the overall coherence of an agent's system of beliefs. In that case, this is a possible degree of belief to which the agent is committed.

There are various technical ways of capturing this idea of coherence. Suppose a person has a certain finite set of credences and opinions. Some of these opinion can be merely comparative and qualitative, as in "I am more confident in P than in Q," "P and Q is much likelier than P and not Q, which is itself likelier than Q," or "P given Q is a lot likelier than P given not Q." If such an agent's credences are consistent, there will be a set of functions satisfying the axioms of probability theory such that each function satisfies every opinion of the agent. Call this set of functions the agent's *Representor*.<sup>15</sup> Here is one analysis of 'coherence-contributing.' A certain probability in a proposition or pair of propositions is coherence-contributing if every function in the Representor contains it. That is, if every probability function consistent with your current degrees of belief has the probability for P being  $n$ , then you are committed to the probability of P being  $n$ . This analysis can be generalized. If a certain structure of probabilities is present in every probability function consistent with your current degrees of belief, then you are committed to having opinions with that structure. For instance, if every probability function consistent with your degrees of belief

---

<sup>15</sup> See Bas van Fraassen, "Fine-grained Opinion, Probability, and the Logic of Belief," *Journal of Philosophical Logic* 24 (1995): 349-377, for formal details about Representors.

make the probability of P given Q higher than the probability of P, then you are committed to having the conditional probability of P given Q higher than the probability of P. In other words, you are committed to Q being evidence for P.

Though this formal analysis is on the right track, it is incomplete. For one, the above account only takes into consideration the coherence-commitments of agents who begin with consistent opinions. This is because a function in your Representor must satisfy the Kolmogorov axioms, which simply *are* the consistency constraints on degrees of belief. Nonetheless, people can have inconsistent degrees of belief and yet still be committed to possible probabilities. For instance, a person might be committed to possible probabilities in virtue of those possible probabilities being coherent with a large enough consistent subclass of her degrees of belief. Why should one inconsistent pair of opinions spoil the party?

Another reason such an account is incomplete is that the other rules already dictate a set of possible degrees of belief to which a person is committed. Your commitments should include those opinions which are coherence-contributing to what you are already committed to. For instance, suppose that A is committed, by the Rule of Parity say, to the opinion that rain is likelier tomorrow than snow, where this opinion is not something A possesses. The Representor for A's actual degrees of belief will differ from the Representor for the set of opinions to which A is committed in virtue of the the other rules of commitments. For example, suppose because of A's other opinions, adding the opinion that rain is likelier tomorrow than snow probabilistically commits A (by coherence) to the opinion that the walk to work tomorrow is likelier to be longer than the walk to work today. The Representor of A's actual degrees of belief will not commit A to the opinion that the walk to work tomorrow is likelier to be longer than the walk to work today, but the Representor of the set of opinions to which A is committed, in virtue of the other rules, will. And the latter seems correct to me as capturing what A is in fact committed to.

To capture these considerations, let me give a revised version of the Rule of Coherence: an agent A is committed to all possible probabilities (and structures of probabilities) which are coherence-contributing for her. Being coherence-contributing will be a matter of being a value of every function in the Representor of A's actual degrees of belief, or a large enough consistent subclass of those degrees of belief, and a value of every function in the Representor of A's commitments, or a large enough consistent subclass of those commitments.<sup>16</sup>

---

<sup>16</sup> What counts as 'large enough' I will leave vague and undefined here, with the hope that either something more precise is possible, or that the vagaries do not undermine the view.

Barry Lam

I do not pretend that these are exhaustive rules of rational epistemic commitments, but they are a start. From the agent's own opinions  $pr_{@}^{(*)}$ , the rules of commitment generate a class of probabilities,  $pr_{\odot}^{(*)}$ . The first thesis of the Commitment-Tracking theory is that Q is evidentially relevant for P for a subject just in case  $pr_{\odot}^{(*)}$  makes Q evidentially relevant for P. The rules of commitment show how Q gets to be evidence for P in virtue of a subject's opinions, without getting to be evidence P solely because a subject believes it to be so. The Commitment-Tracking Theory then states that a subject in the process of reasoning about whether P must track all and only relevant information about P, where what is relevant is given by the structure of the probabilities to which she is committed.

#### 4. The Objective Subjectivity of Justified Belief-Change

Now that I have fully given the Commitment-Tracking Theory of justified belief changes, we are in a position to see the way in which evidence can be subject-dependent. Differences in  $pr_{@}^{(*)}$  will oftentimes generate differences in  $pr_{\odot}^{(*)}$ . This is what happens in the Dr. Spock and Evil-Demon cases above. Dr. Spock's ignorance of the relevant causal facts means that he is only committed to dietary cholesterol being evidence of heart disease. On the other hand, knowledge of such causal facts renders such facts evidence for the scientists in Sydney. The same holds for subjects in Evil-Demon worlds.

We also see the sense in which what is evidence can be objective, or subject-independent. What is evidence for what for an agent *can* come apart from what a subject merely takes to be evidence. In this sense, evidence can be 'independent' of a subject's credences or noetic structure. Your evidence can differ from what you merely believe to be evidence, but this does not threaten the claim that Q is evidence for P *in virtue* of your beliefs and conditional probabilities. As a result, the Commitment-tracking theory has the tools to make sense of Plantinga's cases. The mere fact that a person takes something to be evidence or non-evidence, and reasons accordingly, does not ipso facto make such reasoning fully justified. There are many pieces of evidence a person's reasoning must respect outside of what that person believes to be evidence. But it does not follow from this fact that the kind of evidence a person's reasoning must respect in order to be justified is completely subject-independent, as Plantinga states.

Let us look back at Plantinga's cases of the flat-earthier and counter-inductivist. We seem to have a judgment that there is a kind of unreasonableness



exhibited by people whose noetic structure makes it out so that a picture of a round Earth is evidence for the Earth being flat, and who reason accordingly. However, Plantinga's cases are very misleading in that all he states of the cases is that the flat-earther and counter-inductivist have subjective credences of the right probabilistic structure. One way of filling in the rest of the details of the case saves the intuition that they are unreasonable in the way they reason, but not because evidence is completely subject-independent.

When we imagine agents in thought-experiments, we are not imagining them in all detail, so we make certain (justified) assumptions about what they are like. Typically people are generally committed to taking a picture of an *o* that is F as evidence that an *o* is F. They are committed to this, by the Rule of Parity, as long as they take *one* instance of a picture of an *o* being F to be evidence that an *o* is F. Or they are committed to this, by the Rule of Integrity, so long as they in any way deem as unreasonable people who do not so reason. By making an assumption that Plantinga's agent is typical, we take them to be committed, according to the various rules, to taking the picture of the Earth being round to be evidence that the Earth is round. When he concludes that the Earth is flat, he is thereby *insensitive* to evidence. Such an agent of course, by the Rule of Belief, is also committed to the picture being evidence that the Earth is flat. But that doesn't defeat our judgment that he is also insensitive to evidence. He simply has inconsistent commitments. Inconsistent commitments about a matter render it impossible for an agent to be fully evidence-respecting in her reasoning about that matter. I will have more on this below in section 5.

On the other hand, we could fill in details for Plantinga's agent to make them committed and only committed to a picture of a round earth being evidence that the Earth is flat, or membership in a set of mostly swimmers being evidence that a person does not swim. But as soon as we fill out all of the details, such a person looks fully justified in changing their minds according to their commitments. For instance, we can build into the case that the flat-earther has sufficient beliefs about why pictures of *the Earth being round* is relevantly different in evidential value to pictures of anything else having any other properties. Suppose he thinks that the Round-Earth Society will fabricate millions of pictures of a round Earth and distribute them in his neighborhood only if the society learns that the Earth is flat. He is wrong, as his source for believing this, unbeknownst to him, is completely unreliable. Now he sees a picture of a round earth. Isn't he now fully reasonable in becoming more confident that the Earth is flat? And isn't it because his reasoning is fully respecting the evidence he has?

Counter-induction is no different. On the assumption that A is a lot like us and other people, we evaluate as rational ordinary people who take membership in a group who are mostly F to be evidence that o is F. We are thus committed by the Rule of Integrity that such membership is evidence of an o's being F. By Parity we are committed to membership in a group of mostly swimmers being evidence that John is a swimmer. Absent any further details about A, our normal assumptions about people very reasonably make us conclude that A is not respecting evidence that John is a swimmer in her reasoning about John.

Yet, as soon as we fill in enough details to make A unlike normal people, and therefore not committed to the claims of ordinary people, our negative judgment of A's reasoning disappears. Consider the (very reasonable) conditional probability that a certain day, the hundredth day in a series of days, is dry, on the condition that 99 of the previous days have rained, is higher than one's conditional probability that the hundredth day is dry. The background beliefs behind such a conditional probability is the idea that given enough consecutive days of rain, the chances of a dry day increases, not decreases. Here is one very simple example that makes perfectly reasonable high conditional probabilities on a certain member of a class failing to have a property P, given that the class is such that 99 out of 100 members have P. For one to get such a reasonable conditional probability, we need a certain story, a story about the kind of class it is (an ordered class of days), a story about the sampled day (the last member in that order), a story behind the ordering (consecutive days) and the kind of property it is (raining). This kind of background ends up making perfectly plausible that being a certain member in a class consisting of mostly F members lowers the likelihood that a certain member is F.

If we fill out these kinds of details in the case of A's reasoning about John by giving A similar kinds of background beliefs about John, swimming, and the class of swimmers we are interested in, then even if these background beliefs are all false, it seems that A's inferences about John are fully tracking relevant information, and therefore fully justified. On the other hand, we can construct a case in which we take someone to be changing her beliefs unjustifiably, yet she reasons in a way that appears to be consistent with her conditional probabilities. The Commitment-tracking view will imply that there will always be some hidden presuppositions about an agent's commitments that make sense of why she is unjustified, and such commitments trace back to her other beliefs and normative judgments. The Commitment-Tracking view states that it is possible for Q to be evidence of P for a subject even if the subject does not believe or judge it so, but in

all such cases, Q will be evidence of P for her *in virtue* of her beliefs and judgments.

## 5. Inconsistency

What if an agent begins with inconsistent opinions? The inconsistency manifest in such a way that according to the rules of commitment as I have stated them, the agent becomes committed to a set of degrees of belief that somehow make P both evidence and non-evidence for Q for him. How is he to rationally reason then? Let's imagine an agent, A, who is both hypocritical and ad hoc; A believes one thing is evidence for another, evaluates people in a way inconsistent with such a belief, takes random sampling as evidence of typicality among inanimate objects and their features but not at all as evidence in the case of people and their features. Would such an agent now be committed on my account to something both being and not being evidence? If that is the case, is such a person doomed to imperfect reasoning no matter what he does, since he would necessarily be sensitive to non-evidence for him while being sensitive to evidence for him?

The answer to this question is, yes, you bet! The very reason why you should not have all matters of hypocritical epistemic standards and ad hoc beliefs about evidence is precisely because the facts about normative commitment doom you to reasoning in an imperfectly rational way. This, in my opinion, is the central normative importance of consistency.<sup>17</sup> When you have inconsistent commitments with respect to what is evidence for Q, you cannot reason in a fully evidence-tracking way with respect to that proposition. Any way you change your mind about that proposition will track evidence while also tracking non-evidence, or will be insensitive to non-evidence while failing to be sensitive to evidence. Inconsistent commitments doom you to imperfectly unjustified reasoning.

Almost, but not quite. Aren't we all inconsistent in some way or another? Does that mean no one ever reasons in a fully justified way? Of course not. We might have fully consistent subsets of commitments which make evidence and non-evidence for us in a certain domain perfectly consistent. It would then be possible for us to reason in a perfectly evidence-tracking way in that domain. Secondly, there is clearly a sense in which some people with inconsistent commitments can be more reasonable in their reasoning than others who have inconsistent commitments. A person can be *more* committed to P's being evidence

---

<sup>17</sup> In my dissertation, *The Dynamic Foundations of Epistemic Rationality* (PhD. Diss., Princeton University, 2007), I argue against the idea that the normative importance of inconsistency is that it is a necessary condition of static rationality, or rational belief-states. Also see David Christensen, *Putting Logic in its Place* (Oxford: Oxford University Press, 2005).

Barry Lam

for Q than not, even though she is committed to both. In that case her reasoning, when it is sensitive to P, would be more reasonable than it would be if it were insensitive to P. So a person can be more or less reasonable in their reasoning even if they are imperfectly so, depending on how much more she is committed to one side of an inconsistency than the other.<sup>18</sup>

Suffice it to say, the fact that people can have inconsistent commitments is, I think, a virtue and not a problem for my view. People with all manners of ad hoc and hypocritical epistemic standards and beliefs about evidence cannot but fail to change their minds justifiably. The fact that people who have inconsistent commitments can still reason justifiably in certain domains does not threaten the view, and the fact that people with inconsistent commitments can reason more or less justifiably also does not threaten the view.

## 6. Conclusion

I began by raising three questions about justification that are of interest independent of issues concerning knowledge. I then offered a theory that justified belief-change is a matter of tracking relevant information and failing to track irrelevant information in the way you reason about what is the case. What makes some piece of information evidentially relevant or irrelevant for a subject is the structure of the class of opinions to which she is rationally committed. A subject's rational commitments depend on her knowledge, beliefs, opinions, and epistemic standards, but they are not identical to them. There are rules of commitment that tell us how you can be committed to things you do not believe. This Commitment-Tracking theory of justified belief-change makes sense of our intuitions that we can fault subjects for unjustified reasoning even when such reasoning succeeds according to their own lights. But what is evidence for what still depends on subjective factors.

---

<sup>18</sup> Obviously more needs to be said here about this topic, and how to make precise sense of this view.