Coherence, Truthfulness, and Efficiency in Communication[1]

Sherrilyn Roush

Abstract

Why should we make our beliefs consistent or, more generally, probabilistically coherent? That it will prevent sure losses in betting and that it will maximize one's chances of having accurate beliefs are popular answers. However, these justifications are self-centered, focused on the consequences of our coherence for ourselves. I argue that incoherence has consequences for others because it is liable to mislead others, to false beliefs about one's beliefs and false expectations about one's behavior. I argue that the moral obligation of truthfulness thus constrains us to either conform to the logic our audience assumes we use, educate them in a new logic, or give notice that one will do neither. This does not show that probabilistic coherence is uniquely suited to making truthful communication possible, but I argue that classical probabilistic coherence is superior to other logics for maximizing efficiency in communication.

KEYWORDS: consistency, probabilistic coherence, truthfulness norm, communication, efficiency

Introduction

Why should I be consistent? Why should I not believe both p and not-p? More generally, why should I be probabilistically coherent, with my degrees of belief conforming to the axioms of probability? Dutch Book Theorems tell us that if my belief states violate the probability axioms then I will be susceptible to bets in which I will lose no matter the outcome, and if I adhere to the probabilistic conditions then I will avoid this fate. (Hàjek 2008) Critics wonder why having beliefs should force us to bet at all, and whether there are so many bookies skulking around corners. (Glymour 1980, 63-93) What is supposed to conform to probability constraints are belief states, and one might think that the aim of belief is truth, so some are dissatisfied with the pragmatic quality of the betting

[1] Thanks to Sandy Goldberg, Kim Johnston, Seana Shiffrin, Michael Rescorla, Clark Glymour, Walter Carnielli, Steve John, and audiences at Oxford University, Shandong University, and the 13th meeting on Logic and the Foundations of Game and Decision Theory (LOFT 2018) for helpful discussions of this material.

argument. They want a reason to think adherence to the probability axioms will make us better off epistemically. Recent work has given an attractive way of meeting this demand, for if I am probabilistically coherent then no one with my evidence whose beliefs are not probabilistically coherent will be more accurate than I am. (Joyce 2009)

A striking feature of both of these popular arguments is that they are self-centered, basing the demand for coherence of my beliefs on self-interest, the consequences of coherence and incoherence for myself. I will explore a complementary type of justification of an obligation to be coherent that is based on the consequences for others. The Dutch Book and Accuracy arguments could themselves be re-cast from this perspective. Buying and selling bets involves another party, and what I do often has consequences for her as well as myself. If I am a bettor with probabilistically coherent beliefs, then for every proposition p there is a price at which I would be willing to buy or sell a bet indifferently, that is, at which I would be willing to change places with the other party to the transaction. By contrast for a person with a sub-additive belief function – in which for some p his confidences in p and not-p sum to less than 100% – there are propositions on which the highest price at which he is willing to buy the bet is not one at which he would be willing to sell it. (Kerkvliet and Meester 2017) If this style of believing were universalized it could be bad for business: if lottery tickets were being sold on p, and everyone were sub-additive on p, and everyone had the same buy price and sell price as each other, then no tickets would be sold. The accuracy advantage that probabilistically coherent beliefs bring will also have consequences for others: such a subject – provided she is also sincere, with her assertions mirroring her belief states – will be the best informer she can be to those who trust her testimony.

It could be fruitful to think about many kinds of consequences for others of a person's logic of belief. The consequences that I will focus on here concern an audience's ability to identify what a speaker believes. If I'm inconsistent or violate the probability axioms in some other way, such as failing to believe what obviously follows logically from what I believe, then, I will argue, others are liable to be misled about my beliefs, or else we are both liable to waste resources avoiding this miscommunication.

Here I will argue first that an audience's default assumption that a speaker follows some logic has an important role in efficient communication. Second, not managing our belief states in accord with the logic an audience expects us to follow is liable to make our assertions misleading about our belief states, thereby violating the moral norm of truthfulness. Then I will discuss the scope of the obligation that truthfulness imposes on the management of our beliefs. Since truthfulness will require at most adhering to the logic one's audience happens to expect, it does not single out probabilistic coherence as

uniquely suited to communication. However, I will argue that if we compare the alternatives for how efficient they can make our communication, probabilistic coherence is superior. Thus, in addition to the familiar self-interested reasons we have to be probabilistically coherent we have reasons based on our obligations to and interactions with other people.

1. Logic is marvelous.

One of the useful roles logic plays in investigation of the world is allowing us to extend our knowledge without investigating every proposition individually. (Roush 2017) If I know p, and know that p logically implies q, then I may know q on that basis alone.[2] Logic has a similar role in enabling an audience to fill out a picture of a speaker's beliefs from the meagre evidence of his assertions. My outright assertions are a small subset of the beliefs of mine that those cooperating with me need to know about even in a single context. They will be able to identify many of my beliefs from generalizations about all human beings – I usually believe the obvious, and I don't believe that cats grow on trees – and my behavior is some guide to my beliefs if they already know my values. But some of my beliefs are individual to me. Had we but world enough and time, they could ask and I could make assertions about all relevant propositions as to whether I believe them or not, but we typically don't have time for all that before they need to decide whether or how to act on the basis of my assertions. They may not even know in advance which propositions will become relevant, and so not know which propositions they should ask after my beliefs about.

Things become less difficult if they can make assumptions about my logic, because then they can infer from beliefs I have asserted to relevant beliefs, or lack of belief, on matters that I have not made assertions about. Logic is a way for us to extend our knowledge of the world efficiently, acquiring new beliefs by thought rather than toil. Assumptions about a speaker's logic are a way for an audience to efficiently extend their knowledge of the speaker's beliefs.

Numerous philosophers have claimed that one of the default assumptions we make as interpreters is that the speaker is coherent, by which some explicitly mean probabilistically coherent (Davidson 1973) and all definitely include classical deductive consistency (Quine 1960, ch. 2; Lewis 1974; Dennett 1998, 13-42; Mildenberger 2018).

[2] I say "may", because many, including myself, deny that knowledge is closed under known implication. However, that only implies that there are exceptions to the transmission of knowledge over deductive inference, not that they are typical. (See Roush 2017 and Mayo-Wilson 2018.)

The Principle of Charity of which, for some thinkers, this assumption is a part has even been taken by some to be a necessary condition for interpreting the meanings of a speaker's utterances, and indeed for interpreting a speaker as having beliefs at all (Davidson 1973, 137; Lewis 1974, 112), the idea being that in order to understand the speaker of an unfamiliar language we must pare down the possible interpretations of the language, which we can only do by assuming he is rational, and coherence is required for rationality. Some (Dennett 1998, Mildenberger 2018) claim that an assumption of coherence of the speaker, or at least of classical consistency, is what we use to understand and predict a person's behavior.

Though I am concerned with an audience's assumptions about a speaker's logic, I depart from these authors in several ways. One is that the role of assumptions about a speaker's logic that I am concerned with is not identifying meanings, or determining whether to attribute beliefs at all, and is only indirectly to do with predicting or explaining behavior. Instead, assuming that we understand a speaker's language, I ask how we know which propositions the speaker believes, given the poverty of the stimulus of his assertions. Second, all of these authors assume that probabilistic coherence, or at least classical consistency, is required for rationality[3], and for interpretation. Though I assume that we should be rational I shall not assume that probabilistic coherence is required for this, but instead try to offer some support for this set of formal requirements over alternatives I will discuss. Third, I do not grant the claim that assumption of one particular logic or other is required for interpretation of a speaker's beliefs, or the claim that a speaker and hearer must share a logic, but only the weaker claim that a hearer needs to know which logic the speaker follows, in order to fill out a model of the speaker's beliefs from the subset revealed directly by his sincere assertions.

2. Truthfulness constrains the management of our beliefs.

[3] Mildenberger's (2018) approach is most similar to what I argue here, in that he finds a reason to be rational (which is assumed to mean coherent in the sense that one does not hold beliefs that could not all be true) in the value of being understandable to others and to oneself. He is responding to a contemporary challenge posed by the arguments of Kolodny (2005) and Broome (2013) to find a reason to be rational in this sense. I find Mildenberger's arguments persuasive, but they don't defend coherence over alternative formal requirements on rationality that I discuss here. My questions originate against the backdrop of Bayesian defenses of probabilistic coherence as a constraint on the relations among beliefs. I am focused here on how far the moral norm of truthfulness can deliver an obligation to be probabilistically coherent, and in general on efficiency in communication of one's belief states by means of assertions.

Whether we as interpreters actually make the assumption that speakers are probabilistically coherent is a quasi-empirical question that has not received enough attention. But anecdotal evidence abounds that we make the assumption that speakers conform to key portions of classical deductive logic. A famous example occurs in Lewis Carroll's dialogue "What the Tortoise Said to Achilles" (Carroll 1895), where the Tortoise goads Achilles into trying to provide an argument to convince the Tortoise to make a simple modus ponens inference from premises he has accepted. This is a Sisyphean task if the Tortoise doesn't already have the disposition to do it, which he purportedly doesn't. The dialogue is usually taken to show that no non-circular argument can be given to justify modus ponens, but Achilles' shock and frustration at the Tortoise's repeated refusal to draw the logically implied conclusion, and our sympathy for him, also illustrate our default assumption as hearers that speakers believe the things that follow obviously from their beliefs.

Ralph Waldo Emerson recognized the consequences of inconsistency for being understood by others in his famous (1841) diatribe on the matter:

> A foolish consistency is the hobgoblin of little minds, adored by little statesmen and philosophers and divines. With consistency a great soul has simply nothing to do. He may as well concern himself with his shadow on the wall. … — 'Ah, so you shall be sure to be misunderstood.' — Is it so bad, then, to be misunderstood? Pythagoras was misunderstood, and Socrates, and Jesus, and Luther, and Copernicus, and Galileo, and Newton, and every pure and wise spirit that ever took flesh. To be great is to be misunderstood.[4] (Emerson 2011)

Though he acknowledges the role of consistency in being understood, Emerson is focused on the consequences for himself, on his prospects of greatness and his lack of concern for whether he gets understanding from others. Maybe we can improve our chances of greatness by contradicting ourselves, but being in the audience of such a great one might not be so fulfilling. In life among ordinary mortals we do tend to appreciate consistency from others.

For example a candidate making inconsistent assertions in a campaign for public office is a news story – at least, it used to be – because the point of the campaign is for the audience to get evidence for deciding who will lead the government in the future. For this they need to know in what direction the candidates' decisions can be expected to tend, and this in turn depends, to some extent, on what her beliefs are. If the candidate

[4] I take it that he intended the "is" of predication rather than identity.

professes both p and not-p, which belief is going to guide her behavior?[5] Which, we might say, does she really believe? The need to anticipate behavior is one reason why a world leader publicly contradicting himself is a geopolitical problem. If the president of the United States says that NATO is obsolete, and also says that NATO is a bulwark of international peace and security, should allies and enemies think that the United States will abide by Article V of the treaty, the clause that promises mutual defense?

There can be strategic advantage to being unpredictable, and thus to making contradictory assertions, but if the advantage is supposed to be that others don't know how you will act but you do, then that inconsistency had better not exist at the level of your beliefs. It had better be that you really only believe one of the two propositions. This makes the strategic advantage of inconsistent assertions conflict with a moral norm, because from deliberately asserting contradictory propositions one of which you don't believe, it follows that the latter assertion is a lie. Whether to endeavor to make your beliefs be understandable from your assertions doesn't only depend on whether you aim to be a great one, or feel like conceding advantage to your audience. The moral norm of truthfulness places on assertion a moral obligation to take a reasonable amount of care to make your assertions such as to accurately convey your beliefs. This, I will argue, implies an obligation as a speaker to have your beliefs conform to the logic your audience expects, or else to make provision for its not doing so.

The norm of truthfulness has broader scope than the norm of truth. It doesn't only forbid my deliberately asserting p when I don't believe p, also known as lying. It also enjoins me to avoid misleading, and so misleading about what my belief states are. What we convey by our assertions goes beyond what we assert, and as noted above it needs to do so if we are to communicate efficiently. If I reply to the murderer at the door that the person he seeks is not here, I might tell myself that I haven't lied because by "not here" I meant in this very spot.[6] But I will have been untruthful because I knew that, assuming the conversational maxim that one should say things that are

[5] Some, e.g., Dennett (1998), have thought that identifying a person's beliefs was necessary for predicting his behavior. Kristin Andrews (2012) has convincingly argued that much of a person' behavior can be and is predicted without knowing his beliefs, via simple heuristics such as that people usually do what they say they will do. But there are some behaviors whose prediction requires identification of beliefs – e.g., the enemies will change their encryption code if they believe the current one has been broken. And there are other purposes than prediction of behavior for which identifying a person's beliefs is crucial, for example, as Andrews discusses, understanding and explaining behavior.

[6] The idea that I could determine the meaning of a phrase by my intention is rejected by contemporary philosophers of language. I use this example because of the role it and others like it have played historically in discussions among ethicists and theologians of the moral status of misleading speech. A more linguistically sound example of misleading by implicature is: I ask where the nearest gas station is and you give me directions, but you neglect to tell me that the station is closed. (Grice 1975)

appropriately informative (Grice 1975, 41-58), the murderer would conclude that "not here" means not anywhere in this house.[7] The murderer already knew by inspection that his target was not right in front of him, and I knew that he knew that. Assumptions we make by default in the process of identifying a speaker's beliefs from her assertions – that is, in interpretation – are ripe for exploitation by a speaker for the purposes of misleading.[8,9]

A hearer can also be surprised, and misled, if a speaker's beliefs are not consistent or not coherent, as I claimed earlier. To illustrate this, begin with a case of conversational implicature. Suppose, in a classic example, you ask me how the lecture was and I reply that the speaker was prompt. Since promptness is not relevant to the question you've asked, and you know that I know the conversational maxim about informativeness, you conclude that I meant to inform you of something else, the most natural candidate being my low opinion of the speech. That would be a relevant content that you can assume I was too polite to say explicitly.

So far interpretation by implicature, but there is also entailment. You assume, since I have said the speaker was prompt, that I believe that he was prompt – because you have no reason to doubt my sincerity. And you assume that I don't also believe that the speaker has never been on time – because if I did I would be incoherent, being confident of two things that can't both be true. Consequently, you would not expect me to write a letter of reference that said this speaker has never been on time. If you can't assume that I am coherent, though, then those inferences are not on solid ground. If I am actually not coherent on this proposition then I will have misled you about my beliefs, and also my potential letter-writing, by asserting that the speaker was prompt without also telling you of my other, inconsistent belief state. Notice that I won't have lied. I won't have said anything I don't believe but only failed to say something I do

[7] My descriptions of hearers' expectations in terms of Gricean implicature via conversational maxims could be replaced by descriptions in terms of pragmatic cues, as many contemporary philosophers of language analyze these exchanges. The point remains that hearers have predictable expectations that affect the meaning and halo of implications they take from a speaker's utterances.

[8] In what I consider interpretation here, I assume that the hearer understands the speaker's language, at least the non-logical aspects. The epistemic challenge she has is to identify which of the propositions of that language the speaker believes, not what the speaker's assertions mean. She can take utterances as holdings true, but she can't figure out all relevant holdings true by means of utterances.

[9] For those who hear the echoes of Davidson here I will say that I don't need to assume his philosophy of language in order to claim that interpreters who share my language need to make assumptions about my logic, for the reasons above. In particular I don't need to assume that all speakers are interpreters, that having beliefs requires having the concept of belief, that having thoughts requires being a speaker, that a theory of truth for a language is a theory of meaning, or that interpretation is constitutive of meaning.

believe, and that I know you wouldn't think that I believe because of your assumption about my logic.

In case one thinks that it is not possible to believe both that he was prompt and that he has never been on time, notice that if incoherent beliefs aren't possible then we don't need to justify a requirement of coherence. One might protest that we can't possibly interpret someone as having inconsistent beliefs, even if he has, but if so then that is more grist for the mill that turns default interpretive assumptions into obligations for the speaker if he wishes to communicate truthfully.

Implicatures are cancellable. Entailments are not. In the implicature case, I could correct your knowing glance by saying "Oh, no, I didn't mean the lecture was terrible!" You'd have to allow it, though you might be annoyed. But suppose I write the reference letter saying this speaker has never been on time, since after all I believe that too. You will be more than annoyed if in response to your indignation I say "When I said he was prompt I didn't mean that he has not been late every time he speaks!" I didn't have to say that I didn't believe he has not been late every time in order to give you to believe that I didn't have that belief, because of the assumption of coherence that I knew you would make.

On the scale of untruthfulness that extends from misleading to outright lying, misleading by inconsistent or incoherent beliefs stands somewhere around misleading by implicature, and the lack of cancellability suggests it is worse – the hearer has more right to the inference he makes. Ethicists disagree about how morally significant the distinction between misleading and lying is, but most today think it is not so significant as to make misleading always blameless. (Adler 1997, Williams 2002, Saul 2012) Unless misleading is never wrong, the norm of truthfulness disallows misleading by logical implication in the same way as it does misleading by implicature. This way might be stated, roughly: prima facie it's wrong to mislead, though there may be circumstances that would justify it, such as when you want to and have a right to keep something private and have no other way of doing so, or when being truthful would be hurtful with nothing gained.

Lack of truthfulness is uncooperative. Truthtelling – asserting only what we believe – is part of cooperativeness, and so, I'm saying, is avoiding misleading about the related belief states one doesn't make assertions about. This in turn means that having a logic that governs our beliefs isn't only a way to gird ourselves against adversaries – like bookies – trying to take advantage of us, but also is a part of cooperating with others.

How can a moral norm impose a constraint on our beliefs when believing isn't a directly voluntary act? I choose whether or not to make assertions, but how can I be held responsible for your gaining a false belief from me if that is ultimately caused by something I do not directly choose, namely having beliefs in both p and not-p? I might not even be aware of my inconsistent beliefs. Doesn't deception have to be deliberate? Some use the term "deception" only for deliberate acts. However misleading can be intentional or inadvertent.[10] Even when inadvertent it's something one has done, or left undone, that causes the audience to have a false belief. In the cases where I'm taking it we are responsible for the interpretation the audience makes this is, in part, because it has in common with the more familiar deliberate kind of misleading that the audience's interpretation was foreseeable and likely.

That beliefs are not voluntary acts does make it puzzling how there could be norms on them, but that has not prevented people from thinking there are norms on belief, ways beliefs ought to be formed and relate to each other. We take ourselves to be responsible for our beliefs in a general way, in part because they are under our control in a general way. We are able and obligated to voluntarily direct our attention to them when needed and to develop habits of doing so, and of sorting them in epistemically sound ways. One of the circumstances in which attention is needed is when it is foreseeable that an audience is likely to be misled by our epistemic messiness. Making an assertion is like inviting someone into your house of beliefs. For the party host who is about to have guests, it's time to tidy up. Not letting messes pile up makes one ready for the surprise visitor.

By analogy, you might be misleading because the look on your face when you are concentrating hard on a speaker's argument seems to others a scowl of disapproval. This is more or less involuntary but on learning about it, and not wanting to miscommunicate, you could take steps to train yourself to adjust your facial expression during times of concentration in public, perhaps with the help of a friend willing to elbow you when the expression falls back to its old habit.

How can a speaker be responsible for the inferences that others make? Isn't that unfair? It is not unfair if we think of assertion as coming with a duty of care for those who are listening. The responsibility we have for foreseeable and likely inferences others make from our assertions is analogous to liability for negligence. If I knowingly leave my

[10] "…[N]ew information had become available which convinced Rudd she had inadvertently misled parliament – and she had therefore phoned the prime minister on Sunday to tender her resignation." *The Guardian*, 29 April 2018.

swimming pool in a state of disrepair that isn't obvious, then (in the US) I'm typically liable for the injuries if one of my invited guests slips and falls.

US law around misleading advertisement is suggestive on this topic. It is a complicated question in US law when and whether an advertiser is responsible for the inferences a consumer makes from its speech (advertisements), but 1) it's clear that legal responsibility is not restricted to false assertions but includes some inferences on the part of the consumer, and 2), even if a consumer interpretation is based on a reasoning error – rather than, say, an implicature or valid entailment – the advertiser may be held responsible. One of the key factors is foreseeability; if it was foreseeable that a person practicing a reasonable degree of diligence would likely make that mistaken inference then responsibility may follow. (Shiffrin 2018). The legal doctrine may seem puzzling because it attributes responsibility without finding fault or an intention to deceive, but it makes sense if we think of the wrong of deception in terms of neglect of a duty of care, rather than in terms of intentional manipulation of the hearer for the speaker's own ends, though it also is the latter in some cases.

Just to belabor the point, surely we would think I was responsible for my hearer above inferring that I would not write a letter saying the speaker has never been on time. Or consider the case of David Lewis (1982), who reported once discovering he had inconsistent beliefs that went unnoticed because he hadn't thought about them all together. He believed that Nassau Street runs east-west, the railway line runs north-south, and the street and railway line are roughly parallel. Imagine someone with his inconsistent beliefs who was also, obviously quite unlike Lewis, an epistemic slob. When we ask for directions he is too lazy to get his mental house in order and tells us to walk along the railway because, he thinks, it's parallel to Nassau street, not bothering to notice that he has beliefs inconsistent with that, not all of which can be true; too bad for us that he happened to share with us the one that wasn't. If we knew how this misdirection occurred we would think he was responsible for it.

3. Scope of the Requirement

How much of a coherence requirement on beliefs can we get out of a truthfulness requirement on assertions? There are clear limitations. Even if the obligation not to lie is unconditional, it is generally agreed that the obligation not to mislead is not. In the excusing conditions for misleading in which, for example, the truth hurts with no benefit, or one has a right to privacy, the truthfulness norm also will not yield pressure to be probabilistically coherent. This is not surprising since I am drawing a norm of

coherence out of the conditions of communication, and in these cases one is taken not to have an obligation to communicate.

Even in the cases where truthfulness does obligate us not to mislead, this does not impose a requirement to be ideally probabilistically coherent. It will only support an obligation to be coherent over the range of obvious and contextually salient logical and probabilistic relationships. This is because our audiences are typically not much closer to being logically ominiscient than we are, so they are unlikely to project coherence on propositions that are not relevant in a given context, or too abstrusely related, and consequently are unlikely to be misled about one's beliefs about them. When one's audience changes, one's logical obligations change with it, which seems right and proper. In one respect this is a strength of the view. Human beings' logical abilities are limited because we have finite cognitive resources and limited cognitive interests.[11] It is useful to have an explanation for the limit on our obligations that goes beyond "ought implies can" and identifies the boundary more explicitly.

It is similar for inference.[12] Truthfulness does not obligate us to believe all of the logical consequences of our beliefs, but only, usually, the ones that are obvious; most of the consequences further out our audiences will be no more aware of than we are. For the same reason, the norm of coherence that the norm of truthfulness can support also doesn't require me to have degrees of belief about every proposition in the language, as the ideal probabilistic agent would have. The scope expands with the logical acuity of the audience.

But this is roughly the scope of coherence that I think we take to be required for a finite rational being like ourselves, the requirement for which I'm looking to expand our explanatory resources about beyond self-interest. If David Lewis hadn't noticed his inconsistent beliefs already, then if he were in a context where the disposition of Nassau Street relative to the railway line mattered to an audience – if he were giving directions – then he would notice the hitherto ignored incoherence among his beliefs and sort it out, and that is the scope of the coherence we expect in that context. The fact that the inconsistent beliefs had been quarantined in one's mind does not excuse one from the demand of consistency in a context where both of them are relevant.

[11] It is related to the minimal consistency requirement on Cherniak's (1986, 16-18) minimally rational agent, who must sometimes remove some of his inconsistencies, but here the inconsistencies the agent is obligated to sort out at a given time are specified as those relevant to interlocutors in his context at that time. To generate Cherniak's requirement, we could take the agent to count as his own interlocutor.

[12] Compare Cherniak (1986, 10-16); the requirement here is stronger only in specifying the set of inferences that must be made at a given time to those that are obvious and salient in the context.

To take stock, I've argued that hearers need to make some assumption or other about a speaker's logic in order to infer a model of her relevant beliefs. Having beliefs in violation of the logic that I know others are assuming I use can make my assertions misleading about my belief states. We might think that it follows from this that in a community where others assume I'm probabilistically coherent, truthfulness on my part requires my relevant beliefs to actually be probabilistically coherent.

This doesn't follow, but what does follow is enough for my purposes. As I see it we have three broad options as speakers if we want to avoid misleading audiences about our unspoken beliefs. One is to conform to the logic you know you are assumed to be following, another is to hold your peace – make no assertions at all, and a third is to post notice that you are not following the logic the audience expects (and possibly to educate them). If you make no assertions, then your assertions can't mislead others about your beliefs, and as far as truthfulness is concerned you may follow whatever logic you like, or none at all. None of us will be choosing this option all of the time, but even if we did, notice that it would be a choice not to communicate. It is no surprise that a constraint that follows from the conditions for communication would not apply where communication is not attempted.

Nevertheless this is a limitation of the account. The betting argument for probabilistic coherence would be able to forbid incoherent beliefs even on the part of one who has taken a vow of silence, because he may be betting even if he is not making assertions. The accuracy argument for probabilistic coherence would be able to constrain your private beliefs even if you never spoke, because your beliefs could be less accurate violating the probability axioms than they would be if you followed probability. However my aim is not to replace these arguments for probabilistic coherence, but to supplement them with a complementary view of the matter.

Another way of discharging the cooperative duty of avoiding misleading speech, while not having your beliefs conform to the logic your audience expects, is to make an explicit announcement that you don't follow the logic the hearer probably thinks you do. This would be the best course of action if you found yourself, for example, plopped into a community where everyone else was logically omniscient. (Philosophy audiences often seem this way.) You could preface every speech with a warning that you are only believe logical consequences that would be obvious to an earthling of average intelligence. The applicability of this sort of maneuver is not confined to science fiction or charitable audiences. A woman might be speaking to an audience that she knows assumes women are illogical, and this could be a source of misinterpretation of her unasserted beliefs. Truthfulness does not obligate her to conform to their assumption; if

she informs them of their mistake, then misinterpretations due to their assumption will not be down to her.[13]

Differences between the logic the audience assumes and the one the speaker uses may also be due to audience mistakes. Suppose, for example, that 90% of the population committed the base rate fallacy 90% of the time, which might be true in the actual world. Suppose that because most people commit the fallacy themselves they assume by default that others do the same (not under that description, of course). Surely one is not obligated as a speaker to have base-rate-fallacy-conforming beliefs in order to avoid misleading this audience about one's beliefs. But for anyone in our example who is one of the 10% of people who knows what the fallacy is, it will be foreseeable that a typical member of her audience will commit it on a typical occasion. Does the maneuver of posting notice discharge one's obligation to hearers who make honest mistakes systematically? Merely informing such an audience that I don't follow the logic they're expecting doesn't inform them of what I do follow, or what they can safely infer from my assertions. In the previous examples it did, because in informing the audience their assumption was wrong the speaker in the first case informed the logically omniscient hearers that he was using a restriction of their logic, and the speaker in the second case informed the sexist audience that she was using the logic they use.

Truthfulness obligates me as a speaker to make a reasonable effort to help my hearer, who in this case is likely to commit the base rate fallacy, to avoid getting false beliefs from my assertions, including beliefs about my salient beliefs. A typical occasion for commission of the base rate fallacy is in interpretation by a patient, or doctor, of a positive test result for a disease. A positive test result often makes one think a disease is likely, even in cases where the test result leaves the disease quite unlikely because of a very low incidence of the disease in the population, and a non-zero false positive rate for the test. Imagine one is a doctor who doesn't commit the fallacy and believes a patient's positive test result does not, all things considered, make it likely at all that the patient has the disease. But suppose one is sure that on hearing the test result this patient will commit the base rate fallacy and mistakenly conclude it is likely he has the disease.

[13] This says that informing them is sufficient to meet the demands of truthfulness, not that it is necessary. If an audience makes uncharitable assumptions about the speaker – is not being cooperative – then it is questionable whether she has an obligation to help them avoid false beliefs about her beliefs. She may have an interest in doing so, though, and it is not clear that merely announcing that one is logical would be effective at convincing an audience who believes the opposite. Being truthful doesn't imply that you will succeed in instilling true beliefs in an audience that believes you to be unreliable.

As a doctor one doesn't have the luxury to remain silent. Informing the patient of the test result and that one uses a different logic than he does is not sufficient to tell him whether he has the disease or even whether you think he does, only that it's not safe to conclude that you do think he has the disease. Informing him directly of one's belief that he doesn't have the disease, and that one uses a different logic than he does, is likely to cause suspicion of one's belief, especially if one has told him of the positive test result. This approach truthfully reports on your beliefs, but it is unlikely to lead the patient to a belief in what you think is true. As far as the problem involves logic the doctor does have the option of educating the patient about the mistake, to whatever level is tolerable. For example: Many will think a positive test result means it is likely you have the disease, but that's a probability mistake. But the thicker the explanation the less likely it is to be understood, and the sparser the explanation the more likely it may be to be greeted with suspicion. The most effective way of getting the patient to a belief in what you think is true of his disease status might be to report that the test result was negative, but of course this would be flatly lying. When reporting what you believe and reporting what you think will be most likely to bring about a true belief in your hearer come apart, it is unclear to me which course is more truthful.

To come back to our starting point, we can't always have explicit discussions about default interpretive assumptions, if we want to identify others' salient unspoken beliefs, even if it would be effective to do so, because we just don't have the time. As far as truthfulness is concerned it is possible to excuse ourselves from conforming to the logic others expect of us, but having to either teach your audience a different logic, or explicitly assert every belief of yours that your audience might be interested in – since not knowing what your logic is they won't be able to safely infer them – will be inefficient.

However even if we're satisfied that the most efficient arrangement is to have one default logic we all use and we all know we share, the truthfulness argument does not seem to single out probabilistic coherence in particular as the logic any of us ought to assume or conform to. When my hearers' assumptions about my logic obligate me to conform to it (or hold my peace, or post notice, or educate), that's because they assume that is my logic, and truthfulness to them requires me to take that into account, not because the logic is probabilistic coherence. Maybe the philosophers are right that people actually generally assume probabilistic coherence, but as far as this argument goes, it would obligate me to conform to a different logic had my hearers happened to have *assumed* a different logic.

4. Efficiency in Communication

This contingency looks less devastating in context. It is not well-known among philosophers, but Dutch Book arguments also don't single out classical probability as uniquely rational because the only possible way to avoid sure loss. There are Dutch Book arguments for intuitionistic probability (Weatherson 2003), and for paraconsistent probability (Bueno-Soler and Carnielli 2016). The Dutch Book strategy doesn't even single out probability, for there is a Dutch Book argument for sub-additive belief functions. (Kerkvliet and Meester 2017) Not singling out probabilistic coherence doesn't make Dutch Book arguments uninteresting, and it doesn't make my argument here uninteresting either. But the approach via efficiency of communication can be more specific.

Efficiency of communication is the goal that makes truthfulness relevant to which logic one should obey; if we had no concern about efficiency then we could be truthful by professing our relevant beliefs and disbeliefs one by one and having audiences refrain from any inferences about other beliefs we might hold. The ability to efficiently convey a bevy of beliefs while only asserting a fraction of them is one of the things that logic and assumptions about a speaker's logic give us. Some logics, I will argue, are more efficient than others at this task, and probabilistic coherence is the most efficient of them.

In communication, truthfulness has a dual which I'll call "communicativeness". We could think of truthfulness as having one's utterance and its cloud of inferrables reflect one's associated beliefs accurately. Communicativeness refers to how much of your set of beliefs is revealed by a given utterance. The utterances "p" and "p & q" can both be completely truthful, but, as long as p and q are not redundant, "p & q" is more communicative than "p". The properties of truthfulness and communicativeness could be thought of on analogy with soundness and completeness of a logic. A sound logic proves *only* truths. A complete logic proves *all* truths. A truthful utterance conveys only things the speaker believes. A more communicative utterance conveys more of what the speaker believes.

Uttering a greater number of propositions – in other words, saying more – is a way to increase communicativeness, but it is not the only way. I can publicly commit myself to a level of confidence in one proposition p, and convey more, or less, about my beliefs depending on the logic that I follow and that my interpreter expects me to follow, as I will now argue.

The central axiom of probabilistic coherence that distinguishes it from competitors is finite additivity. If p_1, p_2, …, p_n are exclusive, then $P(p_1 \vee p_2 \vee \ldots \vee p_n) = P(p_1) + P(p_2) + \ldots + P(p_n)$. This is what gives us generalizations of the law of non-contradiction and law of excluded middle, via its implication: $P(p) + P(-p) = 1$. In this way of thinking of what makes a belief system probabilistic, the alternative logics are super-additive and sub-additive, with $B(p) + B(-p) \geq 1$ and $B(p) + B(-p) \leq 1$ respectively, with "B" for degree of belief. (If some subset of the person's belief states is super-additive and the other sub-additive, as we might find in someone who appears intuitively to have no logic at all, then the points below will apply to those subsets.) I will argue that the additive option is the most communicative.

To understand these equations concretely, consider a few observations about what we actually assume. Examples suggest that as a matter of contingent fact, as interpreters we don't quite expect additivity. We expect people not to be super-additive, except in odd cases like liar paradoxical sentences. That is, we expect them not to be over-committed over the logical possibility space; when someone says he thinks p and not-p are equally likely we don't think his confidence might be 2/3 in each. And our expectations are largely fulfilled by speakers for the simple cases.

However we don't automatically assume that a speaker is not sub-additive. If a speaker says she's not confident that the defendant is innocent, then even if we take this to mean that her confidence in innocence is less than 50% I don't think we automatically infer that she is more than 50% confident of his guilt. We accept the possibility that she is not confident of either of the only two possible verdicts, probably because she doesn't think she has enough evidence. This behavior can be made sense of, which is why sub-additive belief functions, even if not ideally rational, are not exactly irrational. We can make sense of sub-additive believers as long as we know that is allowed in their logic. This is one reason why, as I indicated above, I don't think that charity in interpretation requires an assumption of probabilistic coherence.[14]

We as interpreters may not expect exactly probabilistic coherence, but I do think it is the most communicative logic that hearers and speakers could use to interact. There is a sense in which it would be ideal for communication if we were additive believers and speakers and thus were able to expect additivity when we play the role of interpreters. Consider a simple assertion p. By asserting p one will convey more or less information

[14] Another reason is that when we are not doing radical translation – starting with zero knowledge of a language – we can use a shared natural language to discuss which logic we are using or should use. That the interpreter must assume the speaker has the same logic as she does seems to be motivated only by the radicality of the original examples.

about one's beliefs depending on what logic one adheres to and whether one's audience expects one to adhere to that logic. Suppose in every case I'm about to describe, the audience knows what the logic of the speaker is. Even so, there is a difference in what they can infer from the same assertion, depending on what that logic is.

Consider first the case where the logic is super-additive: we allow ourselves and others inconsistent beliefs on commonplace propositions, and, putting that in graded terms, we are allowing ourselves to be over-committed, with $B(p) + B(-p) > 1$. If so then an assertion of p does not rule out a belief in not-p. From an assertion of p the hearer cannot infer lack of belief in not-p. From even a specifically named confidence in p the hearer cannot infer any specific level of confidence in not-p, only that $1 – B(p) \leq B(-p) \leq 1$. If the hearer wants to know whether the speaker believes not-p, or how much, she will have to ask a question. If the speaker wants to convey that information he will have to make another assertion, e.g., "I don't believe not-p" or "My confidence in not-p is 80%". Analogously, if we all allow ourselves sub-additivity, then the assertion that one is not confident in p does not bring with it the information that one is at least 50% confident in not-p. Maybe that's fine, but it does mean that if the hearer of one's statement of lack of confidence in p wants to know one's confidence in not-p, then she will have to ask, and one will have to make another assertion.

All of that could have been avoided. Extra questions and answers, or at least extra assertions, are needed if interpreters know of the speaker's logic that it is or might be sub-additive or super-additive, and these extra conversational steps are not needed if they know she is additive. Assuming that p and not-p are mutually exclusive, from an additive person's assertion that p we can infer that she doesn't believe not-p and from her reported lack of confidence in p we can be sure she is at least 50% confident in not-p. Similarly with entailment and evidential support. If we can assume probabilistic coherence, then we don't have to go around checking every speaker at every step for whether he's a tortoise.

So, unlike Dutch Book arguments, the criterion of communicativeness is able to show a superiority of probabilistic coherence over sub-additive and super-additive belief functions. It also shows an advantage of classical probability over paraconsistent probability, in the following way. Paraconsistent probability is additive, not super-additive, because for inconsistent propositions one takes p and not-p to be non-exclusive. However this doesn't make it as communicative as classical probability. For one thing, not all propositions have this non-exclusive relation with their negations, and an extra tag must be used to indicate which do; the tag is an extra thing one must

express in order to communicate accurately. Also the additivity that is preserved does not preserve the efficiency of classical probability in that the non-exclusivity of p and its negation means that for an inconsistent p, a stated confidence in p will not make a confidence for not-p automatically inferrable.

Conclusion

Our utterances convey more about our beliefs than they explicitly say, and they need to do this if we are to communicate effectively in finite time. They do this by means of default assumptions that audiences make, about for example a speaker's logic of belief. If a speaker knows that an audience assumes his context-relevant belief states adhere to a logic, then the truthfulness norm obligates him to actually have his relevant belief states – spoken and unspoken – adhere to that logic, or else to post notice that they don't, possibly with instruction about what they do adhere to.[15] The option of posting notice will always yield less efficient communication than speaker and audience would have had with the speaker simply following the logic the audience assumes.

Thus far the argument did not single out probabilistic coherence over rival logics, but further considerations about efficiency of communication do. A finitely additive logic of degrees of belief will be more communicative than a super-additive or sub-additive logic, because it will allow us to say more with less. Also a logic of classical probability has an efficiency edge over paraconsistent probability in that though the latter has finite additivity, the degree of belief in an inconsistent proposition doesn't determine that of its negation. For finite beings like ourselves, then, the norm of truthfulness and the aim of communicativeness together provide some defense and explanation of the use of (classical) probabilistic coherence as a logic of belief.

It is welcome but surprising that a pre-existing moral norm would give us leverage in constraining our logic of belief. One might wonder whether the role of the truthfulness norm in this argument gives support to the currently popular thesis of moral encroachment, the thesis that the epistemic status of a belief depends in part on moral factors. (Pace 2011, Gardiner 2018, Moss 2018) Usually this is cashed out as the idea that if a belief might wrong a person, then more evidence is needed to justify holding that belief. Taken so, the phenomenon is not exhibited in the situations I'm concerned with. A belief in p could be the occasion or subject matter of the wrong of misleading someone about one's beliefs, if, say, in addition to an unspoken belief in p you have a belief in not-p that you profess, and you know that your audience assumes you are

[15] If he chooses the option of refraining from speech, then he is not a speaker.

consistent. But what wrongs a person in this case would be the speech-act of asserting not-p – which misleads concerning the belief in p – not the belief in p itself. Moreover, this wrong wouldn't show that the belief in p was not justified. Belief in p might have been strongly supported by one's evidence. If so, then that evidence would likely make it wrong for you to have believed not-p, but that wrong would be thoroughly epistemic.

That the moral norm of truthfulness constrains us to adhere to the logic our audience expects, or else give notice or hold our peace, doesn't show any particular belief or type of belief to be unjustified, or raise the standard of evidence for a justified belief. The direction of normative flow is from a moral norm to epistemic constraints on the relations of one's beliefs, but truthfulness alone doesn't determine which logic to use. I take the lack of conflict between the moral and epistemic demands here to be happy, but not generalizable. That is, for all I know there are cases where our moral and epistemic obligations conflict.[16] In that case I would not automatically infer that the moral trumps the epistemic, or vice versa.

The consideration here that determines which particular logic is preferable is pragmatic, namely, efficiency of communication. But this phenomenon is not pragmatic encroachment on the epistemic realm. In the thesis of pragmatic encroachment (Stanley 2005), how much evidence it takes for a belief to be justified or to be knowledge depends on the stakes in the context, but the role that efficiency of communication plays here is not to affect the standards of evidence for any belief. The role of the pragmatic factor here is rather the one that Carnap (1950) gave to pragmatics in settling external questions about which language to use. All substantive questions were internal questions, questions formulated within the chosen language, and settled by appeal to evidence, that is, settled epistemically.

Bibliography

Adler, Jonathan (1997), "Lying, Deceiving, or Falsely Implicating," *Journal of Philosophy* 94 (9): 435-452.

[16] Stroud (2006) provides an argument for such a case, in which the evidence points to something terrible about a friend and loyalty requires me not to believe that thing. I'm not convinced, since we are much better informed about what a friend is likely to do than we are about a random stranger, and a person we know bad things about is unlikely to remain a friend, so not believing what the evidence points to about our friend as easily as we would about a stranger seems to be epistemically justified. But if I were convinced of the conflict then I would conclude that in following the moral obligation we would fail in our epistemic obligation, and vice versa, not that the moral constrains the epistemic.

Andrews, Kristin (2012). *Do Apes Read Minds? Toward a New Folk Psychology*. Cambridge, MA: MIT Press.

Broome, John (2013). *Rationality Through Reasoning*. Chichester: Wiley Blackwell.

Bueno-Soler, Juliana and Walter Carnielli (2016), "Paraconsistent Probabilities: Consistency, Contradictions and Bayes' Theorem," *Entropy* 18: 325; doi:10.3390/e18090325

Carnap, Rudolf (1950), "Empiricism, Semantics, and Ontology," *Revue Internationale de Philosophie* 4: 20-40.

Carroll, Lewis (1895), "What the Tortoise said to Achilles," *Mind*, New Series 4 (14): 278-280.

Cherniak, Christopher (1986). *Minimal Rationality*. Cambridge, MA: MIT Press.

Davidson, Donald (1973), "Radical Interpretation," *Dialectica* (27): 313-328.

Dennett, Daniel (1998). *The Intentional Stance*. Cambridge, MA: MIT Press.

Emerson, Ralph Waldo (2011), "Self-Reliance (1841)," in *Self-Reliance and Other Essays*. CreateSpace Independent Publishing Platform, 21-39.

Gardiner, Georgina (2018), "Evidentialism and Moral Encroachment," in K. McCain (ed.), *Believing in Accordance with the Evidence*, Synthese Library 398, 169-195. https://doi.org/10.1007/978-3-319-95993-1_11

Glymour, Clark (1980). *Theory and Evidence*. Princeton: Princeton University Press.

Grice, Paul (1975), "Logic and conversation," in Cole, P.; Morgan, J. (eds.). *Syntax and semantics. 3: Speech acts*. New York: Academic Press, 41–58.

Hàjek, Alan (2008), "Arguments for – or against? – Probabilism," *British Journal for the Philosophy of Science* (59): 793–819.

Joyce, James (2009): "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief," in F. Hüber and C. Schmidt-Petri (eds), *Degrees of Belief*. Springer.

Kant, Immanuel (1963). *Lectures on Ethics*, L. Infield, trans. New York: Harper.

Kerkvliet, Timber and Ronald Meester (2017), "A Behavioral Interpretation of Belief Functions," *Journal of Theoretical Probability*. https://doi.org/10.1007/s10959-017-0776-y

Kolodny, Niko (2005), "Why Be Rational?" *Mind; A Quarterly Review of Psychology and Philosophy* 114 (455): 509-563.

Lewis, David (1983), "Radical Interpretation," reprinted in *Philosophical Papers*, Volume 1. New York: Oxford University Press.

-------- (1982), "Logic for Equivocators," *Noûs* 16 (3): 431-441.

Mayo-Wilson, Conor (2018), "Epistemic Closure in Science," *Philosophical Review* 127 (1): 73-114.

McKinnon, Rachel (2015). *The Norms of Assertion: Truth, Lies, and Warrant*. New York: Palgrave MacMillan.

Mildenberger, Carl David (2018), "A Reason to be Rational," *Inquiry*, DOI: 10.1080/0020174X.2018.1470570

Moss, Sarah (2018), "Moral Encroachment," *Proceedings of the Aristotelian Society*, Vol. cxviii, Part 2. doi: 10.1093/arisoc/aoy007

Pace, Michael (2011), "The epistemic value of moral considerations: Justification, moral encroachment, and James' 'will to believe", *Noûs* 45(2): 239–268.

Priest, Graham, Koji Tanaka, and Zach Weber, "Paraconsistent Logic", *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2017/entries/logic-paraconsistent/>.

Quine, W.V.O. (1960). *Word and Object*. Cambridge, MA: MIT Press.

Roush, Sherrilyn (2017), "Closure Failure and Scientific Inquiry," *Res Philosophica* 94 (2): 275–299. http://dx.doi.org/10.11612/resphil.1537

Saul, Jennifer (2012). *Lying*, *Misleading*, *and What is Said*. Oxford: Oxford University Press.

Shiffrin, Seana Valentine (2018), "Deceptive Advertising and Taking Responsibility for Others," in Anne Barnhill, Tyler Doggett, and Mark Budolfson (eds.), *The Oxford Handbook of Food Ethics*. Oxford: Oxford University Press, 470-493.

Stanley, Jason (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.

Stroud, Sarah (2006), "Epistemic partiality in friendship," *Ethics*, 116 (3), 498–524.

The Guardian, "Amber Rudd resigns hours after Guardian publishes deportation targets letter", 29 April 2018.

Weatherson, Brian (2003), "From Classical to Intuitionistic Probability," *Notre Dame Journal of Formal Logic* 44 (2): 111-123.

Williams, Bernard (2002). *Truth and Truthfulness*: *An Essay in Genealogy*. Princeton: Princeton University Press.