

Belief, Credence, and Norms

Lara Buchak, UC Berkeley

ABSTRACT: There are currently two robust traditions in philosophy dealing with doxastic attitudes: the tradition that is concerned primarily with all-or-nothing *belief*, and the tradition that is concerned primarily with degree of belief or *credence*. This paper concerns the relationship between belief and credence for a rational agent, and is directed at those who may have hoped that the notion of belief can either be reduced to credence or eliminated altogether when characterizing the norms governing ideally rational agents. It presents a puzzle which lends support to two theses. First, that there is no formal reduction of a rational agent's beliefs to her credences, because belief and credence are each responsive to different features of a body of evidence. Second, that if our traditional understanding of our practices of holding each other responsible is correct, then belief has a distinctive role to play, even for ideally rational agents, that cannot be played by credence. The question of which avenues remain for the credence-only theorist is considered.

1. Introduction

Full belief (hereafter, just “belief”) is a familiar attitude: it is the attitude that the folk talk about, and it has been a subject of epistemology since epistemology began. Partial or degressed belief (hereafter, “credence”), on the other hand, is a semi-technical notion that has come to the forefront of epistemology more recently. Although the idea that probability features into epistemology traces back to at least John Locke, Frank Ramsey was the first to formalize the idea that beliefs come in precise degrees that can be measured by betting behavior.¹ Since then, credences have been closely associated with preferences about gambles. Some have proposed that disposing an agent to take certain bets is merely part of the functional role of credences, whereas others have proposed that the link is definitional: one's credence in p is the amount of money one is willing to pay in ordinary circumstances for a bet that yields \$1 if p obtains and \$0 if not.

A particular kind of belief will be important in the ensuing discussion: belief in propositions of the form *there is a chance c that p* . For example: “there is a 50% chance that the coin will land heads”; “there is a 99% chance that my lottery ticket will lose”; “there is a very low chance that this table will spontaneously combust.” These *objective-chance propositions* are not necessarily claims about what the chances are according to our best theory of physics. Rather, they are claims about the chance (frequency, propensity, etc.) of an event relative to an implied fixed background: the bias of the coin or the number of lottery tickets, but not a complete physical description of the workings of the coin-toss or ticket-picking

¹ Ramsey (1926).

mechanism. These propositions aren't merely reports of credences: when I tell you that the coin has a 50% chance of landing heads, I am not reporting a fact about my mental state or my evidence but a (purported) fact about the coin. Full belief in a chance- c -that- p proposition will ordinarily be accompanied by credence $\text{cr}(p) = c$. However, we will see that the fact that an agent believes a chance- c -that- p proposition for a particular c (even a very high c) doesn't necessarily mean that she believes p .

There is another important kind of chance-belief: belief in an *epistemic-chance* propositions. For example, one might believe that there is an 80% chance that a particular broken bone will heal without surgery or that there's only a small chance one's co-worker will make it into work on time. This is not a belief about objective chance relative to some background: there are no "chance mechanisms" of the kind involved in the coin-flip operating here. Rather, it is a belief about the relationship between one's evidence and the world.²

How do belief and credence each correspond to how an agent sees the world? When an agent believes p , she in some sense rules out worlds in which *not- p* holds. The truth of *not- p* is incompatible with the attitude she holds towards p (though it is not incompatible with her holding that attitude, since she may be mistaken). On the other hand, having a particular credence in p , at least if it is not 0 or 1, does not rule out either the p -worlds or the *not- p* -worlds. The truth of *not- p* is compatible with the attitude the agent holds towards p , even when she assigns a very high (not-1) credence to p . One way to see this is to notice that if two agents, one who believes p and the other who assigns a high credence to p , each learn *not- p* , then the former, but not the latter, takes himself to have been incorrect.³ Thus, belief that p involves an on-off commitment to p in a way that credence doesn't.

Believing an objective-chance-proposition amounts to representing the world as one in which the relevant event has the relevant objective chance. Like ordinary propositional beliefs, objective-chance propositional beliefs rule out worlds in which the chance-proposition is false. When I believe *that p has an r chance of obtaining*, what I believe is incompatible with p having a different chance of obtaining relative to the implied background. For example, consider the belief that a coin has an 80% chance of landing heads. The chance-proposition is true just in case 80% is the actual bias of the coin (the objective chance, relative to the implied background), and so believing it rules out worlds in which this is false. The coin in fact landing heads and in fact landing tails are each compatible with the chance-proposition, and so the possibilities that are left open are worlds in which there is an 80% chance of heads and the coin lands heads, and worlds in which there is an 80% chance of heads and the coin lands tails. Thus, belief in

² Here is another example to illustrate the difference between objective and epistemic chance. One may believe that a coin has an objective chance of either 80% or 20% of landing heads, and have symmetric evidence with respect to each bias; as a result, one can believe that the coin has an epistemic chance of 50% of landing heads, but one will not believe that it has an objective chance of 50% of landing heads.

³ See, e.g., Fantl and McGrath (2010: 141).

chance-propositions about p does not rule out worlds in which *not-p* obtains, but it does rule out some worlds.

Similarly, beliefs in epistemic-chance propositions rule out the world being some way. Whereas the belief that there is an 80% chance of the coin landing heads rules out certain hypotheses about the coin, the belief that there is an 80% chance of a bone healing without surgery rules out certain hypotheses about the character of the agent's evidence in relation to the world. For example, it rules out the hypothesis that agent's total evidence (the x-ray appearing a certain way, the frequency of broken bones healing in the population) strongly indicates that the bone will not heal without surgery. Does credence rule out worlds in a similar manner? I won't take a stand on this, but note that to the extent one thinks of credence as playing the same role as belief in an epistemic chance proposition, one will think that the attitude one takes by having $cr(p) = c$ rules out, for example, worlds in which c is an inappropriate credence for one to have. But to the extent that one thinks of credence as a state formed unreflectively (as a state, for example, that some animals can be represented as having) or that doesn't represent anything in the world, one need not think that the attitude rules out particular worlds.

There are two different, but related, questions concerning the relationship between belief and credence: how the mental states of belief and credence relate to each other, and how the normative states of rational belief and rational credence relate to each other. I am concerned here with the latter question: for an agent doing what she ought, how do the credences she has relate to what she believes? Naturally, the answer to this question bears on the question of how the mental states relate to each other, but I will not directly address this question here. Rationality is to be taken in the "reasonableness" sense of rationality rather than the "coherence" sense: rational agents are agents who form the beliefs and credences they ought to form given their evidence and, if it is relevant, the situation they find themselves in. I will as far as possible avoid relying on particular normative epistemic notions (justification, warrant, and the like) that each entail reasonableness, since my main argument will rely on data about what a reasonable person ought to believe in certain cases rather than why they ought to believe it. As for credences, whether there are requirements beyond coherence is a matter of debate among formal epistemologists. However, the only fact I will make use of in my argument is a relatively uncontroversial one: if it is part of one's evidence that the frequency of truths in the reference class to which p belongs is x , and if there is no narrower or competing reference class for which one has evidence, then it is at least rationally permissible to set $cr(p) = x$.

There is another debate in epistemology which I seek largely to cross-cut: the debate about whether belief-like or knowledge-like states underlie action.⁴ It may at first appear that anyone interested

⁴ Two prominent theories that claim rational agents act only on what they know (rather than only on what they believe) include those of Fantl and McGrath (2002) and Hawthorne and Stanley (2011). See also the debate about

in discussing the role of credence and belief in mental life cannot avoid taking a stand here, since while full belief has its epistemic counterpart in knowledge, it appears there is no corresponding epistemic notion associated with credence. After all, belief can be true or false – a belief that p is true just in case p is true – but a credence cannot, and since knowledge is factive, only beliefs can constitute knowledge. However, there is a degreed sense in which credences can be assessed by their truth: a credence can be closer or farther from the truth. For example, if p is true, then $cr(p) = 0.99$ is closer to the truth than $cr(p) = 0.1$.⁵ Furthermore, Moss (2013) has argued that credences can constitute knowledge because they can satisfy the factivity criterion when it is properly understood. Therefore, I tentatively accept that in addition to the ordinary notion of knowledge associated with belief, which we can call belief-knowledge, there is a notion of knowledge associated with credence, which we can call credence-knowledge, that plays a role parallel to the role that belief-knowledge is supposed to play with respect to belief. Thus, while I will be concerned with arguing that notions in the “belief package” and notions in the “credence package” each play particular justificatory roles, I won’t take a stand on whether these roles are played by belief and credence or belief-knowledge and credence-knowledge.

I will present a puzzle that, I will argue, lends support to two theses. The first is that (assuming the view known as the Certainty View is false) a rational agent’s having a particular credal state does not entail that she has a particular belief state, even within a given context and set of stakes: belief cannot be reduced to credence. The second is that the notion of belief is ineliminable from our moral practices of holding each other responsible: we cannot construct the norms associated with these practices using credences alone. Thus, unless there is a way of resisting the puzzle, we have either to revise these practices or accept two epistemic notions that don’t fit together well. (I will not in this paper consider a third route: trying to reduce credence to belief or to do without credence altogether.)

2. Assumptions and the State of Play

Let me begin by outlining where the project of reducing belief to credence currently stands. One initial thought, of course, is that (at least for a rational agent) to believe p is to assign $cr(p) = 1$: to believe something is to be certain of it. This view, the Certainty View, is naturally suggested by the description of credences as “partial beliefs.” For the purposes of this paper, however, I will bracket this view and simply assume it is false. This assumption represents a mainstream view, though of course not everyone will be on board with it (I will later discuss whether accepting the Certainty View can solve the problem I

norms of assertion, where Douven (2006) argues that the norm of assertion is not (as the consensus view holds) “assert only what you know” but rather “assert only what is rationally credible to you,” where what is rationally credible to one is what we can or could rationally believe.

⁵ Scoring rule arguments for probabilism have made use of the idea that it is an epistemic virtue to have $cr(p)$ closer to 1 in worlds in which p is true and closer to 0 in worlds in which p is false. See Pettigrew (2011) for an overview of these arguments.

raise for the eliminativist about belief).⁶ If belief cannot be reduced to credence 1, then there are two initially promising proposals. The first is the Threshold View: there is a threshold t such that a rational agent believes p if and only if $cr(p) \geq t$. Of course, t may be somewhat vague. The second is the Modified Threshold View: that credence above a threshold, where the threshold is relative to the *context* or *stakes* involved, is necessary and sufficient for belief for a rational agent.⁷

The (unmodified) Threshold View has met with problems in the form of familiar paradoxes such as the Lottery Paradox and the Preface Paradox.⁸ I will concentrate on the former.⁹ Consider a candidate threshold t . Now consider a fair lottery with n tickets, where $n > 1/(1 - t)$. For a rational agent, propositions of the form “Ticket m is not the winning ticket” are all rationally given credence above the threshold, as is the proposition “Some ticket will be the winning ticket.” On the Threshold View, this implies that it is rational to believe all of these propositions, even though they jointly contradict. Therefore, under the assumption that one ought to believe the conjunction of what one believes (the “conjunction principle”), a rational agent ought to believe a contradiction. Furthermore, Douven and Williamson (2006) show that “defeasible threshold views” – views that say that there is a presumptive threshold credence for belief but that belief can be defeated in the presence of some specified condition that is purely formal in nature – run into modified versions of the lottery paradox. However, not everyone is willing to accept that the Lottery Paradox defeats the Threshold View. The “Lockean,” for example, holds that the Threshold View is correct, and denies the conjunction principle: he thus allows that a rational agent can hold each of the “lottery beliefs” without holding their contradictory conjunction. In any case, we might worry about resting a claim about the relationship between rational belief and rational credence on a case which seems independently to present a puzzle for the belief package.

I will argue that neither kind of threshold view can be correct, using pairs of cases that have the same stakes. The initial pair of cases is familiar to legal scholars, although these cases have also been discussed somewhat in the epistemology literature. The pairs of cases all have a structure in which it is clear what the rational credences are. Furthermore, they illustrate a point that goes beyond refuting the

⁶ Authors generally focus on objecting to the claim that $cr(p) = 1$ is *necessary* for belief. Nonetheless, the claim that $cr(p) = 1$ is *sufficient* for belief has also met with challenges: see Maher (1993) and Hájek (ms.). For an alternative picture on which full beliefs have maximal credence, see van Fraassen (1995), who takes conditional credences to be basic and full beliefs to be derived from them.

⁷ That justified belief requires credence over a threshold, which is relative to the stakes involved, is motivated in Fantl and McGrath (2002) by consideration of the phenomenon of “pragmatic encroachment.” One kind of Modified Threshold View is what Schroeder and Ross (2012) call Pragmatic Credal Reductivism, spelled out in Weatherson (2005), and (under one interpretation) Fantl and McGrath (2010). (See also Harsanyi (1985) for a view of this type.) This view also fits with the general spirit of Hawthorne (2004) and Stanley (2005), although they both formulate their views in terms of epistemic probability rather than subjective probability or credence.

⁸ A more recent argument against Threshold Views that I won’t discuss, but that is worth examining, is Jane Friedman’s (forthcoming) argument from the rationality of suspending judgment on high-credence propositions.

⁹ Kyburg (1961).

threshold views: that there can be no purely formal reduction of belief to credence. Finally, consideration of these cases will help uncover the domain in which belief plays an essential role.

3. No Formal Reduction

We begin with a famous court case, the classic example of what is known as “the problem of naked statistical evidence” in legal scholarship.¹⁰ Here is the court case in broad outline. We will examine the hypothetical version that is usually presented in legal scholarship, the “Blue Bus Case,” which abstracts from the non-critical details of the actual case.¹¹ As Fred Schauer presents it:

“Suppose it is late at night...and an individual’s car is hit by a bus. This individual cannot identify the bus, but she can establish that it is a blue bus, and she can prove as well that 80 percent of the blue buses in the city are operated by the Blue Bus Company, that 20 percent are operated by the Red Bus Company, and that there are no buses in the vicinity except those operated by one of these two companies. Moreover, each of the other elements of the case – negligence, causation, and, especially, the fact and the extent of the injury – is either stipulated or established to a virtual certainty.”(81-82)

In civil cases, the standard of proof is that the plaintiff must prove her case “by a preponderance of the evidence.” This is usually taken to mean “by a balance of the probabilities” (Schauer notes that that is the phrase used in English law), which we might think means $cr(p) > 0.5$, where p is the proposition the plaintiff is trying to establish.¹² However, in the actual case, and “as the overwhelming majority of courts would conclude,” according to Schauer, the plaintiff cannot win the lawsuit, because the evidence that the plaintiff was hit by a Blue Bus is ‘merely statistical’. It is important to note that the statistical evidence is not *inadmissible*; rather, it is *insufficient* on its own.¹³

Given these facts, let us consider another hypothetical case, which we will call the “Green Bus Case”:

Suppose it is late at night, and an individual’s car is hit by a green bus. The two bus companies in the area, the Green Bus Company and the Yellow Bus Company, each operate 50 percent of the green busses. There is an eyewitness, who identifies the bus as belonging to the Green Bus Company (the two bus companies operate busses with distinctive shapes). It is night-time, and so her vision is not ideal: let us say she makes mistakes 25% of the time. All of the other elements of the case remain the same.

¹⁰ Central discussions of this case and others involving naked statistical evidence appear Nesson (1985); Cohen (1977); Thomson (1986); Colyvan, Regan, and Ferson (2001); and Redmayne (2008).

¹¹ Presentation based on Schauer, Chapter 3. See that chapter for further details of the actual case.

¹² But see Cohen (1977) for arguments (in addition to the one considered here) against the thesis that evidential standards can be cashed out in terms of credences or other “Pascalian” notions of probability.

¹³ See Cohen (1977: 82).

Given the standard of preponderance of the evidence, we could speculate that in this case, the plaintiff would win a suit against the Green Bus Company.¹⁴

The situations appear to license the following credences as rational: $cr(BB) = 0.8$; $cr(GB) = 0.75$, where **BB** stands for the claim that a bus belonging to the Blue Bus Company hit the woman in the first case, and **GB** stands for the claim that a bus belonging to the Green Bus Company hit the woman in the second case. However, only in the second case – the one with the lower credence – could the court judge that the plaintiff has won the suit. Let us use the language “a verdict that p is (or is not) licensed” to mean that a court ought (or ought not) to conclude that p . Here we have a case with the same stakes and context, in which $cr(GB) = 0.75$ does license a verdict that **GB**, but $cr(BB) = 0.8$ fails to license a verdict that **BB**. This is to say: threshold views of the relationship between *licensed court verdicts* and rational credence are false.

I don’t want to rest too much on the undoubtedly vexed relationship between it being licensed for a court to conclude that p on the basis of some evidence and it being rational for an epistemic agent to believe that p on the basis of that evidence. What is important about this example for our purposes is that the claims about belief analogous to those about licensed verdicts are intuitive in these cases. It seems clear that when we reflect on all the evidence available in the case, and reflect on what we ought to believe, we don’t have a clear (rational) belief about whether the Blue Bus hit the woman.¹⁵ But in the case of the Green Bus, we do. (If you are worried that 0.75 is never high enough for belief – that there is some *necessary* (possibly stakes-dependent) credence threshold that is higher – then vary the examples to increase both numbers above whatever threshold you think is high enough for the Green Bus Case, e.g. make 95% of the busses Blue Buses in the first case and make the eyewitness 90% reliable in the second.) So I want to tentatively conclude that rational beliefs about this case track the licensed verdicts.

Here is another case, with the same form, that seems to prompt the same intuitions. You leave the seminar room to get a drink, and you come back to find that your iPhone has been stolen.¹⁶ There were only two people in the room, Jake and Barbara. You have no evidence about who stole the phone, and you don’t know either party very well, but you know (let’s say) that men are 10 times more likely to steal iPhones than women. I contend that this isn’t enough to make you rationally believe that Jake stole the phone. If you accused Jake, he could, it seems to me, rightly point out that you don’t have evidence that he in particular stole the phone. He could protest that you only know something about men in general

¹⁴ In any event, if “preponderance of the evidence” sometimes requires only that the claim is more probable than not, we could tweak the information given so that it would license the same credence as in this hypothetical case and also license a court verdict. Since the argument in this section only hinges on what we ought to *believe* in these cases, the complexities of the actual legal system are unimportant to the discussion here.

¹⁵ See also Thomson (1986), who argues that in the Blue Bus case, we don’t *know* whether the blue bus hit the woman.

¹⁶ I thank the students in Robert Audi’s graduate seminar at Notre Dame for suggesting this case.

or on average. But you should have a high credence that Jake stole the phone: if you had to place a bet with only monetary gain and loss at stake, it is clear that you should bet on Jake (given the statistics, you can expect to do better in general by betting on the man in in these kinds of cases: assuming there are an equal number of men and women in the population, then for every 11 cases of iPhone-stealing, 10 are perpetrated by men). On the other hand, if we modified the case so that you know that men and women are equally likely to steal, but a fairly but not perfectly reliable eyewitness (let's say, 90% reliable) tells you she saw Jake take it, it seems that you can rationally form the belief that Jake took it, even though you would have a lower credence in this case. A similar point holds if Jake has a guilty look and if guilty looks indicate strongly but not perfectly that the individual has perpetrated the crime in question.

Statistical evidence generally produces a rational belief in a chance- c -that- p proposition. It also presumably produces a rational credence of $cr(p) = c$. But what is interesting about statistical evidence is that it is often by itself not enough to produce a belief that p , even when c is very high. Admittedly, it will be hard to say what counts as merely statistical evidence, and I am leaving open whether statistical evidence can in some cases produce belief: I only claim that in many cases it cannot, even though it produces a higher credence than a rational agent will have in other cases in which she does believe. In at least some instances, belief is not fixed by credence, even in combination with stakes and context.

That bare statistical evidence cannot produce belief is a common enough position in the literature. The Blue Bus case has been discussed extensively in the legal literature, and to a certain extent in the epistemology literature.¹⁷ Furthermore, in the epistemological literature, Thomson (1986), Kaplan (1996), and Nelkin (2000) have each proposed to solve the Lottery Paradox by claiming that purely statistical evidence should not produce belief.¹⁸ There is disagreement about why exactly these cases don't give rise to a verdict or to rational belief.¹⁹ But most scholars seem to focus on the fact that beliefs formed on the basis of statistical evidence, if true, are correct as a matter of luck, and moreover, that the believer knows this (this makes them different from, say, Gettier cases). For example, as Thomson and Nelkin both point out, beliefs formed on the basis of statistical evidence are unsafe: crucially, they are not causally connected to the truth of the proposition. But the belief in the chance- c -of- p proposition *can be* safe – or, more generally, correct not as a matter of luck – and so need not run afoul of rationality. Furthermore, the relevant credences are not going to run afoul of rationality. If one's credence in p is

¹⁷ For references to the legal scholarship, see footnote 10. Discussions that focus on both legal and epistemological issues include Thomson (1986) and, more recently, Enoch et al (ms.). These do not explicitly focus on credence.

¹⁸ These accounts have come under fire. See, for example, Douven's (2000) reply to Nelkin. Douven's reply is specifically aimed at Nelkin's claim that the "One False Belief" accounts of Bonjour (1985) and Ryan (1996) cannot handle an additional case she proposes. The cases here, however, have a different structure than Nelkin's cases.

¹⁹ For an outline of the disagreement about why they don't give rise to a guilty verdict in the legal case, see Redmayne (2008).

based only on statistical evidence, then one's credence exactly matches the frequency in the relevant class.

What we've seen is that a certain kind of evidential basis can give rise to a justified high credence without giving rise to a justified belief, whereas other kinds of evidential bases can give rise to a justified lower (but still high) credence and yet also give rise to justified belief. What is important about the cases here, and has not historically been the focus of the literature (primarily because the literature on statistical evidence has focused on what makes a belief justified rather than on the relationship between belief and credence), is that (1) the statistical-evidence cases here can be paired with non-statistical evidence cases that have the same stakes and context; and (2) it is clear what the rational (or at least rationally permissible) credence is in these cases. Thus, the argument here against the Threshold View rests on few auxiliary assumptions about rational belief (it does not, for example, assume the conjunction principle) and contains fewer "escape routes" in the form of allowing the threshold to change in response to other facts about the agent's situation. Again, I want to be clear that I don't have a general thesis about the role of statistical evidence in belief-formation. Clearly, statistical evidence, when paired with other kind of evidence, can figure into rational belief-formation: for example, evidence that the fingerprints found at a crime scene are a statistical match with those of the defendant, in combination with some evidence suggesting that she had motive to commit the crime, can lead to both a verdict and a belief in her guilt, when motive alone would not. Furthermore, it is possible that there are some cases in which statistical evidence on its own *can* give rise to belief. I am making the rather modest point that in at least *some* cases of *bare* statistical evidence, the evidence fails to produce a rational belief but does produce a rational high credence: higher than the credence in analogous cases in which the evidence does give rise to belief.

Why not try to build in the type of evidence into a reduction of belief to credence? The problem is that we aren't going to be able to read off the type of evidence from purely formal features of one's credal state. Granted, when the statistical evidence is about objective chance, one will have a high credence, if not credence 1, in a chance-*c-of-p* proposition. But consider again the iPhone theft cases, in which the statistical evidence is clearly not about objective chance. In the first case, you know that Jake is a man and that men are more likely to steal. In the second case, you know that Jake looks guilty and that people are more likely to look guilty after they've stolen. In both cases, you have a high conditional credence that Jake stole, given, alternately, that Jake is a man and that Jake looks guilty.²⁰ But only in the

²⁰ I'm leaving open how we want to represent the statistical evidence in the credal framework, as $cr(p(Js) = 0.9 \mid Jm) \approx 1$, or as $cr(Js \mid Jm) = 0.9$. The latter seems more straightforward, but if we want to interpret statistical evidence as being evidence about the epistemic probabilities, we might want to employ the former. As for the suggestion that believing or having a high credence in an epistemic-chance proposition blocks outright belief, this won't work because epistemic-chance propositions are not believed only in response to statistical evidence: presumably one also

second case do we think you ought to believe that Jake stole. A plausible explanation of this is that the counterfactual “if Jake hadn’t stolen, Jake wouldn’t look guilty” is true if Jake did in fact steal, but the counterfactual “if Jake hadn’t stolen, Jake wouldn’t be a man” is false regardless of whether Jake stole. Or, alternatively, that if Jake is guilty, then his guilty look is caused by his guilt but his being a man is not. And there need be no formal differences in credences between the cases. The crucial point is that one can’t in general read the difference between causation and correlation off of a probability function; one needs to intervene in the world in order to establish a causal relationship.²¹

Even though there won’t be a “local” difference in credence in the cases, one might wonder whether there will be a “global” or “holistic” difference, a difference in credences related to the target credences. For example, one might hypothesize that credences based on statistical evidence are less resilient than credences based on non-statistical evidence.²² One might think that in most cases in which you have an extremely high credence, most pieces of evidence that you might get will not lower your credence very much, but that in the lottery case, for example, the announcement of the winner has the potential to drastically change your credence. Similarly, one might think that a second eyewitness will make less of a difference to the Green Bus case than a first eyewitness would make to the Blue Bus case. Cashed out formally, one might hypothesize that there will be a difference in the probabilities of BB and GB conditional on other relevant evidence. The problem with this response is that there won’t be a difference between these conditional credences when the new evidence is independent of both the old eyewitness and the statistical evidence. Consider in each case the effect of an independent eyewitness, with reliability 0.75, who states that the bus belonged to the other company. In the Blue Bus case, the rational agent’s credence on the new evidence will be $cr(BB) = 0.57$, and in the Green Bus case, her credence on the new evidence will be $cr(GB) = 0.5$.²³ And if it is true that the wrong causal direction is why the evidence that should produce a high credence should not produce a belief, then this point generalizes. The statistical indistinguishability of causation from correlation (in non-intervention

believes that there is a high epistemic chance Jake stole in the “guilty look” case – that is just what it means to believe the guilty look is evidence of Jake’s guilt in this case.


²¹ See Spites, Glymour, Scheines (1993). There are a few exceptions to this general claim but they are not relevant to the present case. Perhaps an objector could claim there will be a difference in one’s credences in the relevant counterfactuals. But I doubt that an agent needs to formulate a credal opinion about counterfactuals in order to count as rational. Alternatively, one could try to add more structure to credence functions. If one wants to take these escape route, it will be an interesting upshot of the argument here that rational agents need to have much more complex credences than is ordinarily supposed.

²² I thank Brian Weatherson and Roger White for raising this point.

²³ In the Blue Bus case, where E is the new eyewitness’s testimony and S is the statistical evidence, $cr(BB | E \& S) = cr(E | BB)cr(BB | S) / [cr(E | BB \& S)cr(BB | S) + cr(E | \sim BB \& S)cr(\sim BB | S)] = (0.25)(0.8) / [(0.25)(0.8) + (0.75)(0.2)] = 0.2 / 0.35 \approx 0.57$. In the Green Bus case, where E is the new eyewitness’s testimony and O is the old eyewitness’s testimony, $cr(GB | E \& O) = cr(E | GB \& O)cr(GB | O) / [cr(E | GB \& O)cr(GB | O) + cr(E | \sim GB \& O)cr(\sim GB | O)] = (0.25)(0.75) / [(0.25)(0.75) + (0.75)(0.25)] = 0.5$.

settings) means that taking *all* the formal properties of a credence function into account – even the global ones – won’t be enough to distinguish between causation and correlation.

What these cases bring out is that **rational credence and rational belief are sensitive to different features of evidence**. So while a given body of evidence will *usually* support a belief just in case it supports a high credence, there is no necessary connection between the two. The statistical cases show that credences don’t distinguish between certain facts about our evidence in the way that belief does. What this suggests is the following picture: at the “base level,” we have a body of evidence, which separately determines rational credence and rational belief. Since evidence that supports a high credence is often evidence that supports belief, there is generally a connection between the two. But the in-general connection is not intrinsic: it occurs because of the way both credence and belief are related to evidence, not because of the way they are related to each other.



Consideration of the fact that two different evidential bases can be such that the one produces a higher credence in p and no belief that p , and the other a lower credence in p but belief that p , also allows us to question an initially plausible sounding tenet about the relationship between credence and belief: if one believes p , and one’s credence in p increases, then one continues to believe p . The statistical cases provide an easy example. Consider Kelly, a rational agent who is participating in a game show where she might win a prize. She has a very high credence (and belief) that the winner is determined by another contestant’s choice, and she has a very high credence (and belief) that the contestant hates her, so she has $cr(\text{WON’T WIN}) = 0.95$. Let’s say that she also believes she won’t win the prize. She then discovers that the winner is determined by a fair 100-ticket lottery. Her credence increases to $cr(\text{WON’T WIN}) = 0.99$, but she no longer believes that she won’t win; rather, she believes that she will almost certainly not win. If you are torn about this case, consider an analogous case involving judgment about a person’s guilt, e.g., you learn that Jake didn’t have a guilty look on his face (just a bad reaction to cold medicine) and simultaneously learn that men are more likely to steal. I submit that your credence in Jake’s guilt will increase, but you will lose your belief. The principle that belief is stable in response to an increase in credence (we might say, that belief is “monotonic” with respect to credence) is generally true. However,

the fact that it is sometimes false shows that its appeal might be explained not by a tight relationship between credence and belief, but by the fact that in most ordinary cases, evidence that leads to an increase in credence also preserves belief.

Belief cannot be read off the purely formal properties of a credal state, even if we take into account stakes and context. However, as I will argue in the remainder of this paper, belief is ineliminable from our best theories about the norms associated with holding each other responsible.

4. Belief and Blame

Given that belief is not reducible to credence, we might hope that we can do away with the notion of belief entirely by precisifying the principles in which it plays a role, or by relegating its role in the mental life of non-ideal agents, e.g., as a heuristic. However, as I will argue in this section, it turns out that we need belief, and its accompanying epistemology, precisely because there is a domain in which our norms involving belief are sensitive to the kinds of evidential connections that belief tracks but credence doesn't.²⁴

Let us consider the context in which the idea of credence was developed, and the norm in which it is well-suited to play a role: that of decision theory. Initially, decision theory was developed to characterize how one should bet in explicit betting contexts where the payoff of a bet depends on an objective-chance mechanism, such as the roll of a dice or the arrangement of a deck of cards. The norm of decision theory in its initial form, as developed by Pascal, was that one ought to choose, among the available actions, the action that maximizes expected monetary value, given the objective probabilities involved.²⁵ That is, when facing a choice among lotteries of the form $L = \{ \$x_1, p_1; \$x_2, p_2; \dots; \$x_n, p_n \}$, where L yields $\$x_i$ with probability p_i , one ought to choose the one with the highest value of $EV(L) = \sum_{i=1}^n p_i x_i$.

²⁴ Theories that seek to eliminate belief altogether include Jeffrey (1970) and Christensen (2004). The latter argues that the notion of binary belief is useful, though “may not in the end capture any important aspect of rationality”(ix). Theories in which belief and credence play different roles in the same domain include the “reasoning disposition account” of Ross and Schroeder (2012). Theories in which credence and belief play the same role but occupy a different discourse include that of Frankish (2009). Sturgeon (2008) is a difficult theory to categorize, since he thinks that everyday evidence does not always rationalize sharp credence, and fuzzy confidence of a certain sort is identical with belief, but I tentatively place his theory in the category of theories in which credence and belief play a role in the same domain. Two theories that do recognize different primary roles for credence and belief are Mark Kaplan’s (1996) Assertion View and Patrick Maher’s (1993) notion of “acceptances.” Both Kaplan and Maher claim that our ordinary notion of belief is not coherent, and each propose to replace it by a notion that shares many of the features of belief and does much of the same work. (Therefore, there is a sense in which these theories are eliminativist.) These theories are not reductionist in the sense that they don’t reduce assertions or acceptances to credence, but they are reductionist in that they reduce the *rationality* of assertions or acceptances to facts about the agent’s credences plus something else: for example, according to Maher, one rationally accepts a proposition if doing so maximizes expected “cognitive” utility. I think these theories are on the right track in their recognition of two very different kinds of activity, one which involves credence and one which involves something else.

²⁵ See Fermat and Pascal (1654).

Decision theory in its modern form is the result of several modifications to this norm. First, the injunction to maximize expected monetary value was replaced by the injunction to maximize expected utility, where utility is a function of money.²⁶ Next, it was proposed that utility is a *subjective* function of money; furthermore, the domain of the utility function was expanded to include any consequence whatsoever, not just monetary consequences.²⁷ Finally, the domain to which the norm applied was expanded: instead of just pronouncing on how one ought to choose between lotteries with objective probabilities, it could pronounce on how one ought to choose between any acts whatsoever: the objective probability function, which assigns values to events that have objective probability, was replaced by a subjective probability or credence function, which assigns a value to any event whatsoever.²⁸ So, the norm of decision theory in its modern form is: when choosing among acts of the form $g = \{E_1, x_1; E_2, x_2; \dots; E_n, x_n\}$, where g yields x_i in event E_i , choose the act with the highest value of $EU(g) = \sum_{i=1}^n p(E_i) u(x_i)$. For example, when one is deciding whether to bring an umbrella to work or leave it at home, one considers the utility of getting wet, of staying dry while not carrying an umbrella, and of staying dry while carrying the umbrella, as well as one's credence in rain and not-rain.²⁹ (Some have proposed further modifications of this norm, but these are irrelevant for our purposes.)

Thus, a norm that was originally developed in the context of betting came to be applicable to all actions, as actions are seen as bets on what the world is like. An important thing to note about this norm is that it captures the structure of the considerations involved in instrumental or means-ends reasoning. We might, colloquially, think of being instrumentally rational as taking the means to one's ends. This idea presents instrumental rationality as applying to an agent who wants some particular end and can achieve that end through a particular means. But the situation of actual agents is more complex. In typical cases, an agent faces a choice among means that lead to *different* competing ends, which he values to different degrees. And, in typical cases, none of the means available to the agent will lead with *certainty* some particular end. So the agent's judgment about what to do – and the norm that describes what he ought to do – must be sensitive both to judgment about which ends he cares about, and how much, and to the likely result of each of his possible actions. Thus, the norm of decision theory can be stated:

²⁶ This was proposed independently by Daniel Bernouilli (1738) and Gabriel Cramer (see Bernouilli 1738: 33).

²⁷ See, e.g., Ramsey (1926) and von Neumann and Morgenstern (1944).

²⁸ See, e.g., Ramsey (1926) and Savage (1954).

²⁹ I note that on a view that has become fairly standard in decision theory (the constructivist view), one does not have a utility function independent of one's preferences, so the norm of decision theory is technically to have preferences that obey particular axioms, from which it will follow that you are representable as preferring the act with the highest expected utility value. So a rational individual will not and cannot explicitly consider the utility of various outcomes. See, e.g., Dreier (1996). The difference, however, won't matter for our purposes – we will primarily be considering whether the norm of decision theory can capture certain intuitive decisions concerning our moral practices of blame – and so I will continue to speak in terms that make the discussion less cumbersome.

ACTION NORM: Perform an act only if that act has higher expected utility than any of the other available acts, and perform any one of the acts that has the highest expected utility, given your credences in the events which bear on the utility of the acts.³⁰

This, of course, is a subjective norm. The objective norm “perform one of the acts that will in fact produce the highest utility” may also be important, but here the focus is subjective norms: what you ought to do given your epistemic state.

Decision theory, then, with credences, is extremely good at explaining justified actions in the domain of what we might call personal action: action when the only or primary relevant stakes are for the agent. It is able to capture the norm of personal action because it explains how both epistemic and value facts jointly contribute to a pronouncement about what one ought to do. Notice that this norm doesn’t itself mention beliefs.³¹

If we want to argue that credence can do all of the work that belief is supposed to do, we would need to show that in all of the contexts in which we seem to have a norm that mentions beliefs or belief-knowledge, we can formulate a norm that mentions only credences or credence-knowledge that recommends the right actions. For example, while it may be debatable what beliefs a rational agent has in the lottery cases, we arguably don’t need the concept of belief in these cases anyway, since we can explain all of the actions an individual should undertake with respect to the lottery just by citing her credences and the Action Norm.³² Given that the Action Norm is the main norm that uses credences, and it is supposed by its proponents to be very general, the natural thing to do is to try to explain all norms that purport to employ belief as particular applications of the Action Norm.³³ For example, one might explain the apparent norm that we ought to act on what we believe as follows. While ideally rational agents reason about what to do using credences, reasoning with belief leads to roughly the same practical upshots in a large range of cases, and given the cognitive costs of reasoning with credences as opposed to belief (or as opposed to the propositions that are the contents of belief), actual humans are better off using beliefs.

What I will show in this section is that there appears to be an important norm governing our current practices that involves belief. I will take for granted that the apparent norm really does govern our

³⁰ If there are no ties, this norm reduces to: “Perform an act if and only if that act has the highest expected utility among the available acts, given your credences in the events which bear on the utility of the acts.”

³¹ Though Ross and Schroeder (2012) argue that its application rests on belief: belief plays a role in setting up decision problems, about which we can then reason using credences.

³² This isn’t to say we can’t explain actions in lottery cases using beliefs: they can perfectly well be explained using beliefs about objective chances. The point is just that we can also explain them using credences.

³³ Additionally, although this isn’t the subject of this paper, the “belief-only” theorist who wants to eliminate the need for credences but doesn’t think they can be reduced to beliefs, would need a way to capture the Action Norm and the interaction between the epistemic and value facts, using just beliefs and chance-beliefs as epistemic facts – or would need to argue that this norm is not the correct norm for personal action.

practices, and I will try to say why this might be so, though I will note an avenue for resisting this. I will then show that this norm cannot be readily redescribed, using decision theory, as a norm involving credences. I will then consider what options are open to the defender of reduction.

Within our practices of holding each other morally responsible, having a reactive attitude – e.g., resentment, indignation, guilt, or gratitude – toward someone on account of her action is a prevalent, perhaps indispensable, way to hold her responsible for that action.³⁴ Whether to blame or praise someone via the reactive attitudes is an all-or-nothing decision based, so it seems, on what I *believe* (or know) about the facts concerning her and her action, such as whether she actually performed the act and whether that act was permissible. While reactive attitudes do come in degrees, the degree of blame I assign to a particular agent is based on the severity of the act, not on my credence that she in fact did it. If I have a 0.99 credence (and full belief) that you shoplifted a candy bar, I feel a small amount of indignation toward you, but if I have a 0.2 credence (and lack a full belief) that you stole from a hungry orphan, I withhold indignation altogether, even if the mathematical expectation of how much blame you deserve is higher in the latter case. Merely statistical evidence seems to play a similar role as in the legal cases: even if I know that 80% of teens shoplift, I ought not to believe of a particular teen that she has shoplifted and I ought not to condemn her for shoplifting. Again, I am called on to make a pronouncement about whether you did some act, and treat you accordingly.

So the norm associated with blame is, roughly speaking:

BLAME NORM: Blame someone if and only if you believe (or know) that she transgressed, and blame her in proportion to the severity of the transgression.³⁵

³⁴ I am presupposing the common view that having a reactive attitude toward someone is sufficient to praise or blame him, in the tradition of Strawson (1962). While I note that there is some disagreement with this claim, it is a natural view to take, as reactive attitudes appear to be an important component of our moral responsibility practices. Further, even if this turns out to be incorrect, the points I make here hold under any reasonable understanding of blame.

³⁵ A few caveats are necessary. First, on some views, there is a distinction between when we ought to blame someone and when we ought to find them blameworthy, and theorists adhering to these views may think that the above norm should concern when to judge someone blameworthy – whether to blame her will be a matter of whether some additional condition (e.g. concerning my relationship to her) is fulfilled. However, this distinction will not matter for the discussion here, since in all cases we may simply assume that the additional condition is met. Second, this norm might be thought of as defeasible, for example if the individual is believed to have transgressed but is judged not to be a member of the moral community. A final complication arises from the possibility that we believe an individual committed a transgression, but we are unsure about the exact badness of the transgression. How to in general handle examples in which the uncertainty is not about whether the agent performed the act but about the status of the act itself is beyond the scope of this paper. It is possible that these examples are best handled by introducing a decision-theoretic calculus into the assessment of the severity of the transgression itself, so that the norm is “Blame someone if and only if you believe (or know) that she transgressed, and blame her in proportion to the expected severity of the transgression.” Nonetheless, the main point is that the norm concerning blame has the general form given above: an epistemic component which must be satisfied if blame is to be apportioned at all, and a separate component describing the amount of blame to be apportioned.

Like the Action Norm, this is a subjective norm. There is also a corresponding objective norm: “Blame someone if and only if she transgressed, and blame her in proportion to the severity of the transgression.”

As I have pointed out, the corresponding norm that would be an application of decision theory appears to be false:

BLAME NORM, CREDENCE VERSION: Blame someone in accordance with the expectation of how severely she transgressed, given your credence that she transgressed and the severity of the transgression.

One might object that we do sometimes blame individuals in a more guarded way, precisely because we are not certain whether or not they meet the conditions required for blame. For example, you might blame a colleague for not showing up to a meeting, but temper your blame to the extent that you are not sure whether she has an excuse. If I am correct that blame essentially rests on belief rather than credence, then there are at least two ways to describe these kinds of cases. First, it might be that you blame her, but you have doubts about whether this is the right thing to do. Here, we might say that you believe she did something wrong, but you have second-order doubts about whether your belief is correct. Thus, you apply the norm, but doubt whether your application is correct. Alternatively, it might be that you don't blame her but are unwilling to definitively withhold blame. Here, we might say that you suspend judgment on whether she did something wrong or not, and we might revise the suggested blame norm to say “Blame someone if you believe she transgressed and don't blame her if you believe she didn't,” where not definitive instruction is given when neither of the conditions are met.

Regardless of the explanation for tempered judgment in any particular case, the test for whether these cases undermine my claim that the blame norm rests on belief rather than credence is as follows. Consider someone who definitively performed an act that merits a level of blame that is exactly equal to the *expectation* of blame in the tempered judgment case: for example, a colleague who arrives late to a meeting (a less bad offense) who you know has no excuse. If the attitude one takes towards the colleague in this situation is the exact same attitude one takes in the tempered judgment case, then this is a point in favor of resting blame on credences. But if the attitude is different – if, for example, you blame the absent colleague more than the late colleague, but take some second-order attitude that mitigates your blame in the former case – then that's an indication that blame has belief rather than credence as a necessary component.

As further support that practices of holding each other responsible are governed by on-off rather than partial attitudes, consider what happens in the courtroom. A jury is called on to offer a verdict – a verdict about whether a particular defendant is guilty or not formed only on the admissible evidence – and it is on that basis that the defendant is punished or not. We could imagine a legal system that punishes defendants on the basis of some partial attitude the jury forms in her guilt: the defendant gets two years if

the jury forms a credence (or partial verdict) of 0.9 in her guilt, four years if the jury forms a credence of 0.95, and so forth. Indeed, perhaps this system would maximize expected utility, when we take into account the value of punishing a guilty person and the disvalue of punishing an innocent person. Alternatively, in civil disputes over a sum of money, we could award the money in proportion to our credence in who it rightfully belongs to. As Nesson points out, this approach “addresses the concerns of the decision theorists so well that a question arises as to why our legal system is so firmly committed to the all-or-nothing rule.”³⁶ And it does not seem that the main objections to such a system are practical: although forming a precise credence would be too practically difficult, we could imagine there being more than two possible verdicts, e.g., definitely guilty, probably guilty, probably not guilty, definitely not guilty, each with an appropriate sentence.

The reason such a system is not adopted is that the jury is called upon not merely to assess the total strength of the evidence, but to render a yes or no verdict as to the defendant’s guilt: to take a stand about whether she is guilty; as Nesson puts it, to “make a statement about what happened.”³⁷ While it is possible that this stand needs to be informed by credences, and perhaps the evidence will justify a verdict only if it also justifies a credence above a threshold (though the ideas of “beyond a reasonable doubt” and so forth may also be explainable without reference to credence), the jury is called on to do something in addition. While the cases of jury verdicts and interpersonal blame aren’t directly analogous – arguably courtroom practices are shaped by merely practical considerations more than reactive attitudes are – they are both examples concerning a norm we have adopted for evaluating individuals, which rests on an on-off attitude, and in which there is some obvious alternative norm that rests on a partial attitude but that we are hesitant to adopt.

If I am right that norms involving attitudes like blame involve belief, then we can close off one kind of response to the statistical cases. One might have thought: every bit of evidence is ultimately statistical, and whether we think a subject is justified in her belief changes with the context, i.e., with whether the evidential uncertainty is described as a lottery. For example, when we are really made to focus on the fact that an eyewitness being 75% reliable means that it is merely a matter of chance whether she got it right in this case, we may no longer think the belief based on eyewitness testimony is justified.³⁸ But we do think there are context-independent facts about whether individuals ought to be blamed on the basis of the evidence. If the norms concerning the attitudes like blame involve belief, then there has to be some privileged description of the evidence with respect to which beliefs ought to be formed or not

³⁶ Nesson (1985: 1382).

³⁷ Nesson (1985: 1359).

³⁸ I thank Sarah Moss and David Christensen for raising this objection.

formed, and with respect to which reactive attitudes are appropriate. (We similarly take there to be context-independent facts about whether jury verdicts are justified.)

The norm concerning blame appears to rest on belief rather than on credence, at least according to our ordinary practices. However, there are a few responses available to the credence-only theorist. The first, of course, is the route this paper assumed away: to accept the Certainty View. Related to this strategy, one could argue that blame requires knowledge rather than belief, and that the Certainty View about the relationship between knowledge and credence – that knowledge entails or requires credence 1 – is correct. A second response is to argue that there is something mistaken about our ordinary practices. One might seek support for this view by pointing out that we do tend to place too much weight on testimony, and are susceptible to the base-rate fallacy. However, in order to solve the problem of statistical evidence for the pairs of cases at hand, one would have to argue either that we ought to blame a particular man for stealing or a particular teenager for shoplifting on the basis of the statistical evidence, or that we ought to withhold blame unless we have credence 1 – that is, unless there is no evidence we consider possible that could tell against an individual's guilt. Both routes would constitute a radical revision of our practices, and strike me as unattractive (indeed, the first strikes me as repugnant). A third response is to argue that decision theory, with credences, can handle all of the cases in question. I now consider this response in detail.

Even though credence is largely absent from legal sanctions, and from our ordinary practice of blame, might the defender of the credence-only taxonomy argue that what ultimately justifies our evaluations is credence, not belief? The argument could run as follows. First, the credence-only theorist could argue that to partially sanction an individual for, say, stealing from a hungry orphan is no worse for the individual than fully sanctioning her for that particular act, so there really are only two possible judgments: guilt and not guilty. This seems to me a difficult claim to establish, but perhaps not impossible. Next, the credence-only theorist could point out that if there really only are two options – sanction someone for a particular act and don't sanction someone for that act – then which of these acts maximizes expected "moral" utility will track an agent's credences, and so credences alone can explain which of these two acts an agent ought to adopt. Thus, the blame norm, employing credence, is:

BLAME NORM, REVISED CREDENCE VERSION: Blame someone, and blame her in proportion to the severity of the transgression, if doing so has a higher expected moral utility than not blaming her, given your credence that she transgressed and the moral utility of blaming/not blaming a guilty person for that transgression and blaming/not blaming an innocent person for that transgression; don't blame someone if doing so has a lower expected moral utility than blaming her; and do either, or follow some tie-breaking rule, if both actions have the same expected moral utility.

Indeed, this norm, when supplemented with the natural thought that it is worse to blame an innocent person than to withhold blame from a guilty person and that how much worse it is increases in magnitude with the severity of the act, will explain precisely why we require a higher credence to blame someone the more severe an act is.

The naked-statistical-evidence cases present a problem for this norm, however, because what the existence of these cases shows is that we can have the same credence and the same stakes in two different cases, but whether we blame in the two cases can be different. Now, the credence-only theorist could argue that the stakes are not the same in, say, the Blue Bus case and the Green Bus case, on the grounds that to judge a company (or person) guilty on the basis of merely statistical evidence itself has a negative utility. For example, she could argue that it is worse to wrongly convict a company on the basis of statistical evidence than it is to wrongly convict a company on the basis of eyewitness testimony. Or, wrongly convicting an individual on the basis of the reference class he belongs to (man, teenager, etc.) is worse than wrongly convicting an individual on the basis of eyewitness testimony. It is *unfair* to convict someone on that basis, because doing so disproportionately harms innocent individuals that, through no fault of their own, belong to the wrong reference class.³⁹ Just as it is wrong to convict on the basis of an illegal search, it is wrong to convict without direct evidence. This is a way of sidestepping the fact that belief but not credence tracks causal relationships between evidence and the hypothesis in question, and that these relationships matter to whether we ought to blame someone on the basis of the evidence, by relocating these relationships in the inputs of the utility function.

Nonetheless, this response will not work. If it is true that our unwillingness to convict or condemn on the basis of merely statistical evidence can be captured by the disutility of a false positive when there is no direct evidence, then this disutility can potentially be *outweighed*. If the statistical evidence yields a high enough credence, then the balance will tip towards convicting or condemning. But I submit that we never think it justified to blame an individual on the basis of merely statistical evidence: doing so is not merely bad, it is prohibited. Even if 99.9% of people in your reference class steal, I can't blame you for stealing on that basis alone.⁴⁰ And this is best explained by the fact that we need a *belief* in someone's guilt to blame her, and that merely statistical evidence cannot give rise to a belief in these cases.

Another way to argue that decision theory, with credences, can handle the intuitions in question – that we cannot blame an individual on the basis of statistical evidence alone but that we can sometimes

³⁹ See Colyvan et al. (2001) for a discussion of the relationship of reference classes to the use of bare statistics in the law.

⁴⁰ Perhaps condemning on the basis of merely statistical evidence has a high disutility even if the person is in fact guilty. If so, then the utility of correctly blaming someone for an act could in principle outweigh this disutility, but again, I submit that it does not.

blame an individual on the basis of evidence that doesn't give rise to credence 1 – is to argue that statistical evidence alone cannot give rise to a high *credence* under certain circumstances. For example, one might hold that when the proposition in question concerns the free choice of an individual and the evidence consists in the existence of an accidental correlation between belonging to a class to which that individual belongs and performing a particular act, then one should not form the credence in question. This proposal would be a radical revision of our theory of credence: indeed, some who work in the foundations of formal epistemology (e.g. Pollock (1990)) hold that frequency-within-the-smallest-known-reference-class judgments underlie all credences, not just credences which explicitly derive from facts about a reference class.

The primary objection to this proposal is that it severs the link between credence and rational betting behavior. For we should clearly prefer to bet on the Blue Bus Company's guilt than on the Green Bus Company's guilt, if the only stakes are monetary. The force of the example is not that we are at a loss to make any epistemic judgment at all about the Blue Bus Company because our evidence is merely statistical, but that we are unable to make the kind of epistemic judgment that is tied with legal and moral condemnation. The force of the example is that legal and moral condemnation are not fundamentally matters of betting. The same point holds in the case of the stolen iPhone: if only monetary gains and losses are at stake, we ought to bet that Jake stole the phone. (I expect that this statement might be met with mild discomfort, and here my suspicion about why. A bet on someone's guilt can never be cleanly separated from a judgment about him: it harms him in the same way that moral condemnation harms him. Thus, it is hard to imagine a situation in which only monetary gains and losses are at stake.)

Indeed, that it is rational to bet on p given naked statistical evidence but not rational to form a reactive attitude that rests on p given the same evidence explains why we are somewhat torn in cases in which there seem to be both personal stakes and moral evaluation involved in the very same action. For example, consider a shopkeeper deciding whether to keep an extra watchful eye on some teenager in his store. On the one hand, it seems as if this action really might maximize his expected utility if the costs of shoplifting are sufficiently high, even taking into account his concern for her. On the other hand, it seems as if keeping a watchful eye on her is akin to judging or asserting that she is not trustworthy (not, note, asserting that his credence is high that she's a shoplifter).⁴¹ In these cases, I submit that we sometimes feel torn about what to do, and I submit that this is precisely because the blame norm coupled with a lack of belief, and the norm about personal action, coupled with a high credence, give conflicting recommendations.

⁴¹ Note that this situation and the situation in the previous paragraph do not have the same structure. In the Blue Bus case, there are two different actions in one situation (rendering a verdict and betting), and in the shoplifting case, there is a single action (keeping an eye on the teenager) that is subject to two different norms: the action norm and the blame norm.

One might be tempted by the examples here to think that when we have evidence that supports a high credence but does not support belief, then judgments about what to do from the point of view of self-interest track what the high credence would imply, and judgments about what to do from the point of view of morality track what the lack of belief would imply. However, there are clearly cases in which merely statistical information is relevant to a moral or other-interested goal, in the sense that using it would have positive consequences for others. Let us consider two such cases, one in which we think that the statistical evidence ought to be used in the judgment about what to do, and one in which it ought not.

Let us consider some group R , where membership in this group is determined by some innate characteristic (such as one's race or the social class of one's parents). First, let us assume that being in some group R is correlated with having a certain non-fatal and non-harmful medical condition, the tests for which are expensive and painful: correlated in the sense that the conditional probability, p_1 , of an individual having the condition given that she is in group R is much higher than the conditional probability, p_2 , of an individual having the condition given that she is not in group R . Now consider a doctor deciding whether to administer one of two drugs to a patient in group R . Condivan works better for people who have the condition and Nocondine works better for people who lack the condition. (No other relevant facts are known about the patient.) As long as p_1 and p_2 are such that using Condivan has a higher expected utility for the patient than using Nocondine given probability p_1 of the patient having the condition, and using Nocondine has a higher expected utility for the patient than using Condivan given probability p_2 of the patient having the condition, prescribing Condivan for a patient of group R will have better consequences given the doctor's credences, because it would maximize the patient's expected utility, and prescribing Nocondine will similarly have better consequences given the doctor's credences, if the patient is not in group R .

Next, let us assume that we have statistical evidence correlating being in group R with impaired driving on a particular stretch of highway, in the sense that the conditional probability, r_1 , of an individual being impaired given that she is driving on that highway and is in group R is much higher than the conditional probability, r_2 , of an individual being impaired given that she is driving on that particular highway and not in group R . And consider a policeman at a checkpoint deciding whether to stop someone and check to see if she is impaired, where pulling someone over takes time and energy and prevents the policeman from paying attention to other driving violations. As long as r_1 and r_2 are such that stopping a particular individual has a higher expected utility than not stopping that individual given probability r_1 of her being impaired, and not stopping an individual has a higher expected utility than stopping that individual given the probability r_2 of her being impaired, stopping an individual in group R about whom no other information is known will have higher expected utility given the policeman's credences, where utility is a measure only of consequences for drivers on the road; and not stopping an

individual not in group R about whom no other information is known will similarly have higher expected utility given the policeman's credences.⁴²

In both cases, there are non-self-interested benefits to be had by acting on the statistical information. Nonetheless, we intuitively think that the doctor ought to take the patient's group characteristics into account but the policeman ought not to take it into account: the doctor ought to act on his credences and the policeman ought not to. The difference between the two cases is that the doctor's action does not even implicitly involve a character judgment, but the policeman's action does: he cannot disproportionately stop people in group R without making an implication that violates the blame norm.⁴³

The difference between these two cases suggests the following conjecture: the natural home of credence is in consequentialist norms, and the natural home of belief – and the domain in which we cannot eliminate belief in favor of credence – is in deontological norms.⁴⁴ If this conjecture is correct, then there may be one more potential escape route for the eliminativist about belief, one that revises our ordinary judgments, but perhaps not radically. Colyvan et al (2010) consider whether standard decision theory can handle the three different kinds of moral theories: consequentialist, deontological, and virtue theory. Unsurprisingly, they point out that consequentialist moral theories fit very well with standard decision theory. They go on to note that deontological theories don't fit as nicely, because in order for decision theory to capture the fact that particular acts are prohibited or required, these acts must have a utility value of negative infinity or infinity, respectively. And assigning infinite utilities carries with it a host of problems: aside from violating the axioms of standard decision theory,⁴⁵ it gives rise to what Colyvan et al call the "swamping problem": any act which has a tiny probability of leading to the satisfaction of an obligation and no probability of leading to the violation of a prohibition will have infinite expected utility and so will be ranked indifferent to an act that will certainly satisfy an obligation (and analogously for any act with a tiny probability of leading to the violation of a prohibition).⁴⁶ However, Colyvan et al think it plausible that the judgments of the deontologist can be captured by assigning 'prohibited' acts very low but finite utility, and 'required' acts very high but finite utility. Their strategy can be employed in our case of blame: we can assign a high, but not infinite, negative utility to blaming someone on the basis of bare statistical evidence. This would imply that the benefits of acting on naked statistical evidence can sometimes outweigh the disadvantages, and this appears to be a revision of our ordinary judgments, but perhaps this is not such a radical revision. (I won't say anything about the

⁴² We could add that stopping an impaired individual would have much better expected consequences *for her*, given the likelihood of injuring oneself while impaired.

⁴³ This is true even if the standard for stopping someone is "reasonable suspicion" rather than outright belief: reasonable suspicion cannot be cashed out in terms of credences either. I thank Jennifer Nagel for raising this point.

⁴⁴ I thank Matt Smith for this suggestion.

⁴⁵ In particular, the "continuity axiom."

⁴⁶ An analogous problem is discussed in Hájek (2003).

case of legal judgment here, because it might be that the legal and moral judgment cases are sufficiently different so that what is sometimes but rarely outweighed in the moral judgment case still ought to be actively prohibited in the legal judgment case because of the necessity of applying simple, general, understandable rules in the latter.)

Relatedly, one could argue that what has disutility isn't the act of condemning on the basis of merely statistical evidence, but the act of thinking in crude maximizing terms in the first place. For example, Nesson claims that the reason we do not adopt a legal system in which awards are proportionate to credences is because this would make the behavioral norm exemplified by the law one of making crude risk calculations rather than one of taking care and ensuring safety. Similarly, we might think that there is something objectionable about treating our interactions with others in the same way we treat moves in a poker game. So, this objection would say, the ultimate norm that justifies our practice of blaming really is best cashed out in terms of maximizing expected utility given credences – but we ought to think of ourselves as doing something else, because thinking this way also maximizes expected utility. According to this response, while the badness of making crude calculations in our interactions with others could in principle be outweighed (otherwise its badness could not be captured by a utility function in the maximization structure), prohibiting these calculations (and instead using a heuristic involving belief) has better effects overall than allowing them, and this explains why they are prohibited. This response again amounts to a revision of our intuitions – if there was a way to use the information only in cases in which the statistical evidence points to a high enough probability of guilt and the utility of condemning a guilty person is high enough, then the response implies that we ought to use it – but again only a revision of intuitions that are perhaps less central.

The proponent of this type of strategy would have to contend with several worries. First, as mentioned, this strategy still appears to require a revision of some of our judgments, and whether it can be successful will ultimately depend on how important it is to preserve those judgments. Second, as Colyvan et al point out, modeling deontological moral theories within decision theory obscures their explanation for why actions are right and wrong: this is to say, while decision theory can capture the judgments that a deontological theory makes, it obscures the deontologist's reasons for making these judgments. Similarly, we might worry that capturing the blame norm using this type of strategy obscures the link between blaming someone and the reasons for doing so. The obvious justification for the blame norm is that when one believes that a person committed a transgression, one is taking a stance on their guilt, and this is intimately tied with blame. However, it is not as clear why having a high credence in someone's guilt should be tied to blaming her. Finally, the proponent of this strategy would still have to explain why degree of condemnation is proportionate to the badness of a transgression but is not proportionate to the credence that the individual committed that transgression.

For the sake of completeness, I should mention one final avenue of escape for the eliminativist about belief. I have been assuming that the theorist who wants to use credences to reconstruct the norm associated with our practices of blame ought to use decision theory. However, one might attempt to use credences but couple them with something other than the norm of standard decision theory. This strategy appears unpromising upon first glance because if the source of the problem doesn't appear to be the limits of the tools of decision theory, but rather the fact that credences can't track features of evidence that are important to whether we ought to blame someone on the basis of the evidence. Still, but it is an avenue that remains open.

5. Conclusion

There are cases in which the only evidence we have is bare statistical evidence. In at least some of these cases, bare statistical evidence seems to not justify a belief that p even though it does justify a high credence in p . For these cases, we can think of parallel cases with the same stakes and context and in which the potential belief has the same content, but in which the evidence does justify belief despite justifying a lower credence than in the statistical cases. This phenomenon poses a problem for both the Threshold View and the Modified Threshold View of the relationship between belief and credence. Furthermore, if the explanation for this phenomenon concerns the causal relationship between the hypothesis and the evidence, then there can't be any formal reduction of belief to credence, because the difference between causation and correlation can't be read off a credence function, even given its global features.

The impossibility of reducing belief to credence wouldn't be problematic if we could eliminate belief from our taxonomy altogether and show that credence can do all of the work that belief appears to do. However, here cases of bare statistical evidence present a further problem. The norm associated with our practices of blame appears to employ belief (or knowledge) rather than credence. Because we cannot condemn someone when we merely have a high credence in her guilt, where this credence is formed on the basis of statistical evidence that doesn't give rise to belief, the prospects for reconstructing the blame norm in terms of credence are dim. However, there still remain some avenues for the eliminativist about belief. First, she could adopt the Certainty View, either about the relationship between belief and credence or about the relationship between belief-knowledge and credence, and argue that the blame norm can be cashed out in terms of credence 1 in the individual's guilt. Alternatively, she could argue that we ought to revise our judgments about cases in which the evidence we have is merely statistical: either revise them radically (we are licensed to condemn on the basis of bare statistical evidence) or less radically (condemning on the basis of bare statistical evidence is not prohibited but has high negative

utility). Finally, she can argue for a theory alternative to standard decision theory that still rests on credences rather than beliefs.

What is it about blame that makes it subject to evaluation by a different type of norm than typical individual actions are? One suggestion is that reactive attitudes cognitively commit one to certain beliefs.⁴⁷ For example, perhaps reactive attitudes associated with blame, like resentment and indignation, are partially constituted by representing the world as being such that their targets are culpable for the act. If this is right, then it might be that many emotions, not just reactive attitudes, take belief – as opposed to credence – as a basis for their warrant: for example, it might be that to fear something is partially constituted by representing it as dangerous, so that fear is appropriate when and only when you justifiably believe something dangerous, and the warrant of fear cannot be put in terms of facts about credence.⁴⁸

Here is a final thought about where the conflict seems to lie, and how we might further characterize the domains in which credence seems to play a natural role and the domains in which belief seems to play a natural role. Both the norm of betting and the blame norm include both an epistemic condition and a condition about value. However, in the norm governing which bets to take, these conditions make a combined contribution to the instrumental value of a bet, rather than separate contributions. And insofar as individual actions are like betting – as decision theory assumes – this point can be generalized. In the decision-theoretic picture of instrumental rationality, one does not first settle on a single goal and then pick the act that has the highest probability of achieving that goal; nor does one first commit to a single picture of what the world is like and then pick the act that is best in that world. Rather, the value of all of the possible outcomes of all of the acts, as well as the probabilities of all of the possible states, figure into the procedure for choosing an act. The norm is of the form “If [condition concerning both credence in various states and the utility of acts in various states], then [action]”: the antecedent cannot be separated into two independent conditions, one epistemic and one involving value.

Unlike the norm of decision theory, however, the norm involved in our practices of blame is a separated norm: it is of the form “If [condition concerning belief in a particular state] and [condition concerning the value of the act when the world is in that particular state], then [action].” The procedure for choosing whether and how much to blame someone involves separately settling on what the world is like and determining the amount of blame that is appropriate when the world is like that. Thus, the blame norm is composed of two independent judgments. In deciding whether to blame, we settle on one possibility (e.g. that the individual is guilty) and the others don’t play a role in our judgment, but in deciding which bets to take, even unlikely possibilities play some role in our judgment.

⁴⁷ I thank Jonathan Weisberg for this suggestion.

⁴⁸ I thank Jada Twedt Strabbing for this suggestion.

The kinds of norms in which credences fit well and those in which beliefs fit well appear to be of a different form: credences fit well into norms in which the epistemic and value components are integrated, and beliefs into norms in which these components are separated. Perhaps, then, the standards for belief and credence – how each must be responsive to the evidence and what general coherence principles they must obey – arise from this difference in the kinds of norms they figure into.

Acknowledgements

I would like to thank Matt Benton, Lindsey Crawford, Julien Dutant, Jane Friedman, Julian Jonker, Matthew Lee, Jennifer Nagel, Ted Poston, Jada Twedt Strabbing, and Jonathan Weisberg for comments on earlier versions of this paper. I would also like to thank Sherri Roush's graduate seminar, Robert Audi's graduate seminar, the philosophers at Leeds, and participants in Epistemology Above the Arctic Circle (sponsored by the Oslo Center for the Study of Mind and Nature) and the Harvard Workshop on Belief for helpful discussions of earlier versions of this paper.

Bibliography

- Bernoulli, D. (1738), 'Exposition of a New Theory on the Measurement of Risk', in *Econometrica* 22(1) (1954): 23-36.
- Bonjour, Laurence (1985). The Structure of Empirical Knowledge. Cambridge: Harvard University Press.
- Christensen, David (2004). Putting Logic in its Place: Formal Constraints on Rational Belief. Oxford: Oxford University Press.
- Cohen, L. Jonathan (1977). The Probable and the Provable. Oxford: Oxford University Press.
- Colyvan, Mark, Damien Cox, and Katie Steele (2010). "Modelling the Moral Dimension of Decisions." *Noûs* 44(3): 503-529.
- Colyvan, Mark, Helen M. Regan, and Scott Ferson (2001). "Is it a Crime to Belong to a Reference Class?" *Journal of Political Philosophy* 9(2): 168-181.
- Douven, Igor (2003). "Nelkin on the Lottery Paradox." *Philosophical Review* 112(3): 395-404.
- Douven, Igor (2006). "Assertion, Knowledge, and Rational Credibility." *Philosophical Review* 115(4): 449-485.
- Douven, Igor and Timothy Williamson (2006). "Generalizing the Lottery Paradox." *British Journal for the Philosophy of Science* 7 (4): 755-779.
- Dreier, James (1996). "Rational Preference: Decision Theory as a Theory of Practical Rationality." *Theory and Decision* 40: 249-276.

- Enoch, David, Talia Fisher, and Levi Spectre (ms.). "Statistical Evidence, Sensitivity, and the Legal Value of Knowledge."
- Fantl, Jeremy and Matthew McGrath (2010). "Belief." Chapter 5 of *Knowledge in an Uncertain World*, Oxford University Press.
- Fantl, Jeremy and Matthew McGrath (2002). "Evidence, Pragmatics, and Justification." *The Philosophical Review* 111(1): 67-94.
- Fermat, P., and B. Pascal (1654), letters, collected as 'Fermat and Pascal on Probability' and translated from the French by Professor Vera Sanford, in D. Smith (ed.), *A Source Book in Mathematics* (McGraw-Hill Book Co., 1929): 546-565.
- van Fraassen, Bas C. (1995). "Fine-Grained Opinion, Probability, and the Logic of Full Belief." *Journal of Philosophical Logic* 24: 349-377.
- Frankish, Keith (2009). "Partial Belief and Flat-Out Belief." In *Degrees of Belief*, eds. Franz Huber and Christoph Schmidt-Petri. Springer.
- Friedman, Jane (forthcoming). "Rational Agnosticism and Degrees of Belief." *Oxford Studies in Epistemology* 4.
- Hájek, Alan (2003). "Waging War on Pascal's Wager." *Philosophical Review* 112(1): 27-56.
- Hájek, Alan (ms.). "Staying Regular."
- Harsanyi, John C. (1985). "Acceptance of Empirical Statements: A Bayesian Theory Without Cognitive Utilities." *Theory and Decision* 18 (1):1-30.
- Hawthorne, John (2004). "Knowledge and Lotteries." Oxford: Oxford University Press.
- Hawthorne, John and Jason Stanley (2011). "Knowledge and Action." *Journal of Philosophy*.
- Kaplan (1996). *Decision Theory as Philosophy*. Cambridge, UK: Cambridge University Press.
- Kyburg (1961). *Probability and the Logic of Rational Belief*. Wesleyan University Press.
- Maher, Patrick (1993). "The Concept of Acceptance." Chapter 6 of *Betting on Theories*. *Cambridge Studies in Probability, Induction, and Decision Theory*. Cambridge, UK: Cambridge University Press.
- Moss, Sarah (2013). "Epistemology Formalized." *Philosophical Review* 122(1): 1-43.
- Nelkin, Dana (2000). "The Lottery Paradox, Knowledge, and Rationality." *Philosophical Review* 109(3): 373-409.
- Nesson, Charles (1985). "The Evidence or the Event? On Judicial Proof and the Acceptability of Verdicts." *Harvard Law Review* 98(7): 1357-1392.
- von Neumann, John and Oskar Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.

- Pettigrew, Richard (2011). "Epistemic Utility Arguments for Probabilism." *Stanford Encyclopedia of Philosophy*.
- Ramsey, Frank (1926). "Truth and Probability." Reprinted in The Foundations of Mathematics and other Logical Essays, ed. R.B. Braithwaite, 1950. New York: The Humanities Press.
- Redmayne, Mike (2008). "Exploring the Proof Paradoxes." *Legal Theory* 14: 281-309.
- Ross, Jacob and Mark Schroeder (2012). "Belief, Credence, and Pragmatic Encroachment." *Philosophy and Phenomenological Research*.
- Ryan, Sharon (1991). "The Preface Paradox." *Philosophical Studies* 64: 293-307.
- Savage, Leonard (1954/1972). The Foundations of Statistics. Dover. First published in 1954 (John Wiley and Sons, Inc.).
- Schauer, Fred (2003). Profiles, Probabilities, and Stereotyping. Belknap Press.
- Spites, Peter, Clark Glymour, and Richard Scheines (1993). Causation, Prediction, and Search. New York: Springer-Verlag.
- Stanley, Jason (2005). Knowledge and Practical Interests. Oxford: Oxford University Press.
- Strawson, Peter F. (1962). "Freedom and Resentment." *Proceedings of the British Academy* 48 (1962): 1-25.
- Sturgeon, Scott (2008). "Reason and the Grain of Belief." *Nous* 42(1): 139-165.
- Thomson, Judith Jarvis (1986). "Liability and Individualized Evidence." *Law and Contemporary Problems* 49(3): 199-219.
- Weatherson, Brian (2005). "Can we do without pragmatic encroachment?" *Philosophical Perspectives* 19: 417-443.