Statistical Resentment, or: What's Wrong with Acting, Blaming, and Believing on the Basis of Statistics Alone

David Enoch and Levi Spectre*

If a colleague has been speaking ill of you behind your back, you may resent them for so doing. Furthermore, such resentment may be appropriate, or called-for, or fitting, and certainly morally (and otherwise) permissible. But for such resentment to be fully appropriate, it is not sufficient that that colleague has in fact been badmouthing you. It's also necessary for you to have adequate evidence to that extent. Even if there is an objectivist sense in which it's appropriate to resent them if and only if they have been badmouthing you, there are also more subjective notions that accommodate the intuition that if you resent them without sufficient evidence, you are being morally irresponsible, and this even if as things happen to turn out, they were badmouthing you. With such, more subjective forms of evaluation in mind, then, we can ask: What evidence suffices for justified resentment? We can ask, that is, questions in the evidence law of resentment (and of morality more generally).

We do not have an answer to this general question. But this shouldn't prevent us from recognizing paradigmatic cases of sufficient evidence, and paradigmatic cases of insufficient evidence. That you've overheard something, that a reliable mutual colleague tells you what has been said behind your back, a confession – all of these, in the right circumstances, could constitute sufficient (if fallible) evidence to justify resentment. That an unreliable frenemy – whom you know is envious of your colleague for your relationship – says they've been badmouthing you is not sufficient evidence to justify resentment.

Well, then, how about the following: In your social milieu, or in your workplace, almost everyone speaks ill of each other behind their back. The thought that you could resent someone for badmouthing you based merely on such statistical evidence – however strong – seems ridiculous. Even if you *can* resent on purely statistical evidence – even if you *do* – clearly, this is not the kind of evidence that can justify resentment (The Central Resentment Case¹). But why? What is it about statistical evidence of this kind – good, probative, sufficiently strong statistical evidence – that precludes it from being sufficient for resentment, even where fallible, probabilistically equivalent non-statistical evidence *is* sufficient for justified resentment? This is the (largely neglected²) problem of statistical resentment. Our aim in this paper is to clearly raise the problem, and make some progress towards solving it. As will emerge, the progress will

^{*} For comments on previous versions we thank Dan Baras, Jamie Fritz, Georgi Gardiner, Clayton LittleJohn, Berislav Marušić, Oded Na'aman, Ittay Nissan-Rozen, Erik J Olsson, Jim Pryor, Daniel Telech, Gideon Yaffe, and audiences at the Eastern APA, Cornell, Dartmouth, Oxford, The Hebrew University, Brandeis, King's College, Hamburg, and Bochum.

¹ Here and below, we use small caps for names of cases.

² It is not entirely neglected. See, for instance, Enoch, Spectre, and Fisher (2012, 222-3), Buchak (2014), Enoch (2018), Pundik (2020).

be limited, but, we will argue, there are rather strict constraints on how good and comprehensive an answer to this question can be. And on the way interesting lessons will emerge on related cases, and on the relations between the epistemic and the practical and moral more generally.

In section 1, we offer a host of structurally similar examples, in the hope, first, that they are of independent interest, and second, that they help generate adequacy constraints for a solution to the problem of statistical resentment, thereby also giving reasons to reject alternative suggestions that can be extracted from the literature. In section 2 we discuss impurist epistemological treatments of statistical resentment, that is, epistemological stories that incorporate moral considerations. In section 3 we focus on knowledge. In section 4 we discuss three more epistemologically conservative accounts of the inappropriateness of statistical resentment conservative in the sense that they try to offer an explanation of the problem in terms of standard, not-too-controversial or topic-specific epistemic principles. These attempts, we argue, do succeed in going some way towards solving the puzzle of statistical resentment, but they fall short of the solution we may have hoped for. In the concluding section 5 we suggest that the way forward proceeds via reduced expectations - perhaps we should settle for a somewhat messy, less than fully unified, account.

1. **Examples**

We already have in front of us THE CENTRAL RESENTMENT CASE. In this section there are going to be many more examples. Before getting to these further examples, we want to explicitly explain how we see the role of this multiplicity of examples. Our main task remains that of explaining the suspicious nature of statistical resentment. Other examples are thus introduced primarily as ways of checking possible explanations for the appropriate level of generality. At least ideally, we are after an explanation of the suspicious nature of statistical resentment that explains all and only sufficiently similar phenomena. So it's important to look at (purportedly) similar phenomena, and see how if at all different competing explanations apply to them. This means that for each example it will be important to ask not just what the intuitive verdict about it is, but also to what extent it seems similar in the relevant respects to THE CENTRAL RESENT-MENT CASE. We don't have a deep, metaphysical account here that establishes the unity of the different examples. Rather, we tentatively rely on intuitions about their similarity, together with the following plausible, rather minimal methodological principle: if two theories are equally good in other respects, and one of them offers a unified account of what seem like unified phenomena and the other doesn't, this is an advantage of the former. Thus, the role of the many examples that follow is primarily instrumental, to help and explain why statistical resentment is so suspicious. But, once these other examples are on the table, they also become intrinsically interesting, and we hope to gain insight about them as well.

You may have a worry pulling in the opposite direction, though: You may doubt the focus on resentment, asking what's so special about it compared to many other attitudes and responses. But resentment is special in our context, because – perhaps somewhat roughly – it lies in a unique way at the intersection of the epistemic and the practical. Resentment is not - certainly not merely - a belief or some other cognitive attitude. So while epistemic norms are clearly relevant to resentment, it's not as if *only* epistemic norms are relevant. And while resentment is clearly morally infused, it is not clearly and directly *practical*: Noting the appropriateness of resentment (in a given context) does not yet settle the question what actions if any are called for³. As we will see, this makes resentment an especially interesting, perhaps especially theoretically puzzling case.

Just in case you're not with us on the intuitive verdict we put forward in some of the examples, let us make the following two points. First, what we will typically ask you to accept about the different cases is not so much the bottom line (appropriate or inappropriate), but a comparative one: In The Central Resentment Case, for instance, what is crucial for our discussion is merely that statistical resentment is *more* problematic than resentment based on similar direct evidence (say, the testimony of a reliable colleague), even when there is no clear probabilistic difference between the two. Often, such comparative judgments will be extremely intuitively compelling, even when the non-comparative bottom line less so. Second, we invite you to see what constraints for an account of statistical resentment are generated by your own intuitive responses to the examples. In order to facilitate this, we supply – at the end of this section – a table summarizing the significance of the different cases.

With these preliminaries out of the way, then, consider the following examples.

It's natural to compare THE CENTRAL RESENTMENT CASE with legal cases. Basing a conviction on purely statistical evidence – say, convicting you of crashing the stadium's gates based on just the fact that many more people attended than bought tickets (GATECRASHERS)⁴ – seems just as unacceptable as statistical resentment is. And issuing a finding in torts against the Blue Bus Company when we know a bus wrongfully caused harm, simply because the Blue Bus Company owns, say 70% of the buses in town seems unacceptable, even if the law has no problem assigning liability based on a 70%-reliable eye witness (BLUE BUS COMPANY)⁵.

It's also natural to compare such practical cases – in morality or in law – to purely epistemic ones. In (one version of) the lottery paradox (LOTTERY), we are reluctant to attribute knowledge (that the ticket you just bought won't win), even when the relevant belief is supported by overwhelmingly strong statistical evidence (the lottery's odds), and even when we would be happy to attribute knowledge based on not-merely-statistical, but probabilistically equivalent or weaker, evidence. The reluctance to attribute knowledge in such cases intuitively seems very closely related to what's wrong with the resentment case and the legal cases, though of course

⁵ For this and related cases, see Enoch, Spectre, and Fisher (2012), and the references there.

³ We focus on resentment rather than on blame in order to avoid the practical implications that are sometimes associates with blaming. With resentment, we can postpone all practical questions and focus just on the appropriateness of resentment itself.

Buchak (2014) seems to mean by blaming only or almost only the reactive attitudes. In this respect, then, her discussion is directly relevant to ours. In others, though, not so much: Buchak is less concerned with explaining what is problematic with statistical resentment or blaming, and more in taking as data the fact that we don't assign blame proportionally to credence in blameworthiness as evidence for the indispensability of beliefs in a framework that already utilizes credences.

⁴ Originally from Cohen (1977).

how closely related is a matter of controversy, a point that will be relevant below⁶.

Some practical cases seem very different from the moral and legal cases, in that relying on statistical evidence in them seems entirely unproblematic (indeed, required). A physician attempting to diagnose a patient, for instance, should rely on statistical evidence (PHYSICIAN)⁷, and if that's all she's got, on naked statistical evidence as well. She would be negligent not to8. So the problem with statistical evidence does not extend throughout the practical domain in a similar way. PHYSICIAN also shows that in some contexts it's perfectly ok to treat people (at least also) as things, to use empirical methods to study them and to decide how to treat them, and so on9.

The problem with statistical resentment is not uniquely moral (with the legal somehow riding piggyback on its close cousin, the moral), as can again be seen from reflecting on cleanly epistemic lottery cases. Nor is the problem essentially related to the fact that in statistically resenting someone we would be wronging them¹⁰. To see this, think of positive moral cases. So suppose that among analytic philosophers it's very common for people to show the kind of courage needed in order to speak truth to power (TRUTH TO POWER). Can you then praise us – two specific analytic philosophers – for our willingness to speak truth to power based on solely this statistical evidence? This seems ridiculous¹¹. More generally, statistical praise and pride seem just as problematic as statistical blame and resentment, and intuitively problematic in the same way, for the same reasons. Still, it would be a huge stretch to say that in praising us (based on naked statistical evidence) you are wronging us. So the problem with statistical resentment

For critical discussions of these theories of ours, see Blome-Tillmann (2015), Pardo (2018), Smith (2017), Gardiner (2018a) and Littlejohn (2017). We hope to discuss these criticisms elsewhere. For pushback against Pardo (2018), and an overview of our account, see Enoch and Spectre (2019).

⁶ In previous work (Enoch, Spectre and Fisher (2012); Enoch and Fisher (2015)) we argue that the epistemic condition known as Sensitivity nicely explains the intuitions in all of these cases. We also insist, however, that it cannot justify the legal suspicions against statistical evidence, because, in a slogan, the law shouldn't care about epistemology. Instead, we offer a Sensitivity-related incentive story. It's not clear, however that that incentive story – or any other one – can be applied to THE CENTRAL RESENTMENT CASE. More work needs to be done, then, in search of a more unified account. Hence this paper.

⁷ This is a common example. See, for instance, Pundik (2008, 303), Buchak (2014, 305).

⁸ This is so, we believe, even if in order to achieve medical knowledge, or the best kind of medical knowledge, there's a need for more than just statistical evidence, like perhaps some access to the underlying causal mechanism. For an overview and some references, see Reiss and Ankeny (2016, section 5). The role of causal explanations in medical knowledge and practice is a delicate and contested matter. Compelling causal explanations are arguably responsible for the persistence of medical practices long after statistical evidence debunks their efficacy or shows them to be harmful or dangerous. See Parsad & Adam (2015). We cannot, of course, engage these complexities

⁹ We think that this suffices to reject Basu's (2019a) use of Sherlock Holmes as an example of someone who mistreats people merely by studying them scientifically. It's possible, of course, that she does not go for the most general version of this claim, in which case the point in the text is one she can accept.

¹⁰ For an emphasis on the need to accommodate the directionality of the relevant wronging – that it's wronging the relevant person – and for the (plausible) claim that the theory we put forward in our (2012) fails to satisfy this desideratum, see di Bello and O'Neill (2019).

¹¹ We're not sure what to say of *generic* admiration, as when we may admire nurses and fire-fighters for their often courageous and compassionate work. The example in the text is not about admiring the class, but specific members within it, based on purely statistical evidence. (There will be more of generics below).

is not (just) about wronging the resented.

You may think the problem is about *agency*. What's common to the resentment case, the legal cases, and the positive moral cases as well – and is absent from PHYSICIAN – is that the statistical evidence is brought to bear on exercises of agency. But an emphasis on agency would be too narrow. Suppose that statistical evidence shows that children of Asian descent are better at mathematics (MATH). Being better at mathematics is not an exercise of agency (it may be related to exercises of agency, but this is different). Still, believing of a specific child that she's good at math based on merely this statistical evidence seems problematic, in the standard statistical-resentment kind of way. So the problem is not about agency¹².

The distinction between people and (other) things seems relevant here (and we get back to it in discussing impurism, below). Compare believing of a specific child (of Asian descent) that she's good at math based on the statistical generalization that children of Asian descent are good at math, to believing of a specific Asian car that it is fuel-efficient based just on the statistical evidence that Asian cars tend to be fuel-efficient (CAR). We take it to be intuitively clear that there's a difference between the two cases, though evidentially they seem to be on a par¹³. Still, we take there to be no problem at all in buying the Asian car (because of its fuel-efficiency, based on statistical evidence alone), where presumably there would be a problem in basing action on the analogous evidence (or belief) about the math skills of a fellow student (say, choosing others for a joint assignment) based on the statistical evidence. If you're not convinced, think of such evidence as a tie-breaker. If you are considering which car to buy, and the relative advantages of two cars (as based on your evidence) are exactly balanced, except that one is an Asian car, and you have the statistical evidence about Asian cars being fuel efficient, there's nothing problematic about going for the Asian car for that tie-breaking reason¹⁴. Not so for the case of choosing a classmate for an assignment – relying on the statistical evidence there is, at the very least, problematic, even as a tiebreaker. So there may be a temptation to think that the distinction between people and things is only relevant here vis-à-vis actions (rather than beliefs). But this is not so: THE CENTRAL RESENTMENT CASE - resentment alone, which may or may not be appropriate even independently of any actions to which it may lead – was designed precisely to rule out this hypothesis.

You may think that people (unlike other things) are entitled to be treated as individuals, and that this explains the problem with statistical resentment. There's something to this thought, we're sure, but still, it's not sufficiently determinate to be helpful. In what sense are we not treating someone as an individual when we're relying on statistical evidence? After all, we must rely on *some* evidence. If you resent me for badmouthing you based on the testimony of reliable

¹² Which means – pace Pundik (2008) – it can't be about autonomy and certainly not about freedom of will.

¹³ There may be differences between them of the kind we discuss in section 4, below. For an emphasis on such stories, see Gardiner (2018b, 184-186).

¹⁴ This doesn't mean that you should believe, of a specific Asian car, that it is fuel-efficient based solely on the statistical evidence.

As a reviewer rightly noted, in some cases there may be further, not statistical evidence about Asian cars (say, that they have some relevant technological feature). Our point in the text is that preferring the Asian car makes sense even in the absence of such further evidence.

others, you don't conflate the evidence (their testimony) with what you're resenting me for (badmouthing you); you're just relying on the evidence available to you. Similarly, it seems, for the case of statistical evidence: You don't resent anyone for being a member of a class or some such; you resent them for badmouthing you, relying on the evidence available to you. How is this different, then? We seem to be back at square one: The intuition that people are entitled to be treated as individuals seems - without further details, at least - not to explain what's wrong with statistical resentment, but rather to repeat the intuition that something is. Moreover, at least many of the legal cases – civil cases like BLUE-BUS – or purely epistemic cases like LOTTERY don't seem to have anything to do with this kind of individuality.

Interestingly, there are cases of highly individual evidence that seem problematic in ways that are similar to those in which statistical resentment is problematic. Think in this context of PRE-RESENTMENT (or pre-pride)¹⁵, that is, resenting someone (or being proud of them) for something they haven't yet done, but will. The evidence you have that I will badmouth you need not be statistical in the sense we started with (say, something about our social milieu). It can be very specific things you know about *me* and *my* dispositions. Still, pre-resentment seems suspicious, and furthermore (we're pretty sure) in a similar way to statistical resentment. If so, the problem with statistical resentment is not primarily about individuals' right to be treated as individuals.16

In order to have a fuller picture of the status of statistical evidence, it is important to have in mind also a fairly wide variety of possible responses. So far, we've been focusing on mostly three: Belief, action, and a specific kind of attitude, resentment (and also its positive counterparts). But there are many other attitudes and emotions: We can ask about statistical admiration, statistical liking and disliking, statistical envy, statistical respect, and so on. All of these seem suspicious in ways very close to that of statistical resentment. This confirms the point above that what's problematic about statistical resentment is not morality-specific¹⁷. It also pushes in the direction of a different hypothesis – perhaps what's wrong with statistical resentment is not so much something that's common to that case and the legal cases of relying on statistical evidence, but rather something that's common to all attitudes and emotions, or to a very wide subset of them. Perhaps, for instance, while statistical evidence can support de dicto attitudes (say, preferring to hire those - whoever they are - who satisfy the description "did volunteer work in college", because of correlations between doing such work and success in a certain kind of work), it cannot support de re ones (like respecting her, or admiring him, or resenting you).

On the epistemic side too there are more possible responses than just beliefs. For one thing, there are credences, of which there will be much more later on. Also, think of being surprised

¹⁵ These are paraphrases on the pre-punishment literature. See, for instance, Smilansky (1994).

¹⁶ Notice that cases like PRE-RESENTMENT work better with causal accounts (see Thomson (1986)) according to which a necessary condition for E being evidence for H is that H plays a role in causing E. But such causal accounts have other problems facing them.

 $^{^{\}rm 17}$ Notice that some of these attitudes may be held vis-à-vis non-person things. So the problem doesn't seem to be person-specific either.

(SURPRISE). While the statistically-based belief that your lottery ticket won't win doesn't amount to knowledge – and according to some theorists, is irrational – still, if your ticket does win, it seems entirely appropriate to be surprised. And while the relations between believing that p (or believing that probably p, or having a high credence in p) and being surprised that not-p are not trivial, still the appropriateness of being surprised upon finding out that your ticket has won - when the only evidence you had that it wouldn't was statistical - shows something about how such statistical evidence, even if not of the kind that would otherwise have supported knowledge, can still rationally shape your expectations. Some of the other cases perhaps especially those involving people – seem different, though: Being surprised to find out, about your Asian-descent friend, that he's not good at math, seems at least shaky18. And being surprised that your friend did not in fact badmouth you - when the only evidence suggesting otherwise was statistical - seems clearly problematic in the statistical resentment kind of way. (At the very least: Being surprised in these cases is not just as unproblematic as it is if your lottery ticket wins.) Perhaps thinking of the appropriateness of being surprised, then, strengthens the thought that the problem with statistical evidence cannot be merely epistemic.

Another way in which different epistemic responses may behave differently concerns the content of the belief (when it is a belief). For we can ask not just about whether statistical evidence supports a judgment about a specific member of the relevant reference class, but judgments about the reference class as a whole. And here, in many cases there is no problem. Perhaps you can't know, on statistical evidence alone, that your lottery ticket won't win. But you can very well know on just that statistical evidence19 that your ticket is highly unlikely to win, or that the vast majority of tickets won't win. We think that the same is true across the board, but we want to register a doubt about this. One of the problems with many cases of objectionable racial profiling is the move from statistics to a specific person (or even a subgroup), beliefs (or credences) about them, and ways of treating them. But one may think that in some such cases the statistical evidence can't even support a belief in the statistical generalization: Isn't there something suspicious about the belief that, among Israeli Jews, Ashkenazi Jews are worse tippers, even if based on otherwise sufficient statistical evidence (ASHKENAZI TIPPERS)?²⁰ At

 $^{^{18}}$ Joel Kim Booster – a stand-up comedian of Asian descent who was adopted at a young age by a white American family - tells the story of how his adopting parents were disappointed to discover that he was not good at math (https://www.youtube.com/watch?reload=9&v=IzYJCqD117A&feature=youtu.be). Regardless of how accurate the story is, the inappropriateness of his parents' expectation and disappointment strengthens the point in the

¹⁹ Or on the basis of the statistics. Perhaps the very same statistics amounts to statistical evidence for the proposition that the ticket is a losing one, and a different kind of evidence for the proposition that the ticket is highly unlikely to win. (We thank Dan Baras for this point.) This difference doesn't make a difference for our point in the text, though.

²⁰ The example is loosely based on related rumors among those whose income depends on such things. For a similar example in the American context, see Basu (2019b), though she never explicitly distinguishes between the problem with the generalization and the problem with the belief about a specific person.

In the background are also partly linguistic issues that we can't address here: How are generics ("Dogs bark."; "Sharks are dangerous."; "Ashkenazis are worse tippers.") best understood? What evidence is appropriate for accepting them, or for rejecting them? But matters here are more complicated than merely the understanding of

least, isn't there something more problematic about such a belief compared to the belief that your lottery ticket is probably going to lose, even if the statistical evidence for both is equally weighty²¹? And – perhaps in some sense an intermediate stage between the statistical generalization and the proposition about a specific person – probabilistic beliefs about specific people (that this Ashkenazi or even this Ashkenazi table is highly unlikely to tip well) do seem problematic in a way that the analogous belief about the lottery ticket is not.

Finally, let us mention another family of cases, those where the relevant statistical evidence is DNA evidence (DNA). It's important to emphasize that DNA evidence - even "cold-hit" DNA evidence, that is, a DNA match that is achieved without any other basis for suspicion, that is, the DNA version of naked statistical evidence – is often treated by the law as less suspicious than statistical evidence of the kind present in GATECRASHERS and BLUE BUS COMPANY. More importantly, it seems to be intuitively different as well, though it's proved remarkably hard to show why. How does DNA evidence interact with thoughts about resentment? Some cases of DNA-based resentment are every bit as suspicious as other cases of statistical resentment. If scientists find the badmouthing-gene, and if 99% of people with that gene badmouth (all) their friends, still it's inappropriate for you to resent me for badmouthing you based just on me having this gene. Still other cases of DNA-based resentment may be less problematic. If a DNA sample was obtained at a scene of a crime, and if you're a match, then even if there's no other reason to suspect you, it doesn't seem ludicrous to say that it would be acceptable to resent you for committing the relevant crime (and maybe to punish you as well). It's not clear to us what to say of such a case: Most epistemological stories have a hard time explaining this data²².

Here, then, to conclude this section, is a table summarizing the cases mentioned so far (there will be a small number of further cases, and further variations on cases, below, as they are needed to highlight problems with the explanations considered):

generics. Presumably, there's nothing suspicious about the generalization that non-drinkers are worse tippers. Different generics behave differently here.

²¹ Things may get messy here, because, as a reviewer pointed out, it's natural to be suspicious of the statistical evidence about tipping in a way that it's not natural to be suspicious of the statistical evidence in LOTTERY. This is true, but it doesn't do away with the intuitive difference between these two cases (once such distorting factors are stipulated away), or if need be other cases. Consider, for instance, Moss's (2018b) at least somewhat problematic example of believing of someone that they are probably straight.

The reviewer also pointed out – rightly – that the relevant intuitions may be context-sensitive in all sorts of ways. But all we need for our purposes here is that there are *some* – perhaps perfectly natural – contexts in which such differences as between ASHKENAZI TIPPERS and LOTTERY are manifested. Suppose that a waiter tells a colleague: "Can we switch tables? I probably won't win the lottery and so I need more than minimal tips." There need be nothing wrong with the premise or the conclusion. But if it's clear that what explains the utterance is that sitting at the relevant table are Jews of Ashkenazi decent (and not, say, just the different sizes of the tables), then the inference starts to look problematic in a statistical-resentment reminding way.

²² The combined story we suggested in Enoch, Spectre and Fisher (2012) and Enoch and Fisher (2015) – in terms of a Sensitivity requirement and the relevant incentive structure - actually does a remarkably good job with DNA cases.

Name	Brief Description (NSE = naked statistical evidence.)	Tentative Verdict	Like resent- ment?	Suggested Upshot
THE CEN- TRAL RE- SENTMENT CASE	Resentment, based on NSE.	-	(Same.)	Not just about action. Not merely about belief.
GATECRASH- ERS	Criminal conviction, based on NSE.	-	Yes.	The problem is not grounded merely in attitudes; applies also institutionally.
BLUE BUS COMPANY	Tort liability, based on NSE.	-	Yes.	The problem is not just about conviction and punishment.
LOTTERY	Belief, based on NSE.	- ?	Yes.	The problem, restricted to knowledge, is not confined to the practical. (Whether the belief can be rational is a different matter.)
PHYSICIAN	Diagnosis, based on NSE.	+	No.	Not all practical cases are problematic.
TRUTH TO POWER	Admiration, based on NSE.	-	Yes.	The problem does not depend on wronging the relevant person.
Матн	Belief about competence, based on NSE.	-	Yes.	The problem is not grounded in agency.
CAR	Belief about fuel efficiency, based on NSE.	+ ?	No.	The problem is sensitive to the distinction between persons and things.
PRE-RESENT- MENT	Resenting for an action not yet performed, based on non-statistical evidence.	-	Yes.	The problem isn't just about statistical evidence, or about being treated as an individual. And the problem persists even when belief is justified, and when knowledge is arguably present.
SURPRISE	Being surprised (about any of the above going against statistics).	Different answers for differ- ent cases.	Different answers for different cases.	Epistemically, not just about beliefs. And not just about epistemology.
PROBABILIS- TICALLY HEDGED BE- LIEFS	Beliefs with probabilistically qualified content (e.g. "that probably p").	Different answers for differ- ent cases.	Different answers for different cases.	Relevance of not purely epistemic considerations also to such beliefs.
GENERICS	Beliefs in generics based on NSE.	Different answers for differ- ent cases.	Different answers for different cases.	The problem not restricted to unqualified beliefs de re. Not purely epistemic considerations are relevant.
DNA	Beliefs and atti- tudes based on cold-hit DNA evi- dence.	Different answers for differ- ent cases.	Different answers for different cases.	?

2. <u>Impurism</u>

The case of statistical resentment lies at the intersection of moral and epistemic considerations. So it is natural to wonder whether thoughts about the impure nature of epistemology²³ – about the relevance of pragmatic and moral considerations in evaluating beliefs – can help with diagnosing the problem with statistical resentment. We discuss this line of thought by first briefly introducing impurism (2.1), then showing how it may be relevant for statistical resentment (2.2). Then, in section 2.3, we argue that such impurist strategies will not work – we voice some general suspicions about impurism, but we focus on ways in which *even if* some impurist line can be made to work, still, it is unlikely to help much with the problem of statistical resentment.

2.1. Introducing Impurism

In epistemology, you're an impurist if, roughly, you think that non-evidential considerations – practical ones of sorts – constitutively affect the status of beliefs as (epistemically) rational, or as counting as knowledge. The idea first became popular through thoughts of how pragmatic stakes affect, for instance, our willingness to endorse certain knowledge attributions, so that the same level of evidential support suffices for knowledge when the stakes are low, but not when the stakes become significantly higher²⁴. If this is so, then epistemic status does not supervene on evidence, and the pragmatic (here, in terms of stakes) encroaches on the epistemic.

If thought about high and low stakes motivates pragmatic encroachment, it seems to equally motivate moral encroachment²⁵. Perhaps, for instance, in some contexts some body of evidence (say, memory of where you placed which sandwich) suffices for knowledge or for justified belief that the one on the left doesn't contain peanut butter, but if your guest is dangerously allergic to peanuts, perhaps the same body of evidence no longer suffices. The high *moral* stakes involved seem to affect what's needed for the desired epistemic status (justified belief, or knowledge) to be satisfied. If so, the moral too encroaches on the epistemic²⁶.

Other ways of motivating pragmatic encroachment also seem to equally apply to the moral domain. For instance, many²⁷ think there is a close connection between *knowledge* and *action*. Perhaps roughly, the thought is that one is entitled to treat as a premise in one's practical reasoning all and only those propositions one knows to hold. If this knowledge-action link (or some other close one) holds, then all the pragmatic considerations that are relevant to one's practical reasoning are potentially relevant to whether or not one knows. But moral considerations too are relevant, of course, to one's practical reasoning, and so a strong knowledge-action

²³ Other names in the literature for impurist views include *Subject Sensitive Invariantism*, *pragmatic encroachment*, and *Interest-Relative Invariantism*. "Invraintism" implies a rejection of contextualism and relativism. For our purposes "Impurism" carries no such rejection. All of the claims we'll be making can hold across and within contexts (or assessment points).

²⁴ See, for instance, DeRose's (1992) discussion of what have come to be called bank cases. See also Lewis (1996), whose Belief Rule is stakes-dependent.

²⁵ See, for instance, Fritz (2017).

²⁶ The example is from Ross and Schroeder (2014).

²⁷ E. g. Hawthorne (2004), Stanley (2005), Moss (2018b).

link, if it is plausible, and if it supports pragmatic encroachment, also supports moral encroachment.

Or consider the thought that whether or not a belief amounts to knowledge depends on whether the believer is in a position to rule out all salient alternatives based on her evidence. If pragmatic or moral considerations affect which alternatives are salient, this entails that the pragmatic – and the moral – encroach on the epistemic²⁸.

There are also ways of motivating moral encroachment that are not parasitic in any way on the motivations for pragmatic encroachment. Marušić (2015), for instance, thinks we should believe that we will fulfill our promises and act on our intentions or resolutions, independently of the evidence available to us. And Stroud (2006) thinks - or is at least willing to entertain the thought - that the moral significance of friendship places constraints on the appropriate response to evidence (that one's friend has behaved shamefully).

Lastly, consider that some beliefs (and perhaps other doxastic attitudes) may themselves be morally wrong²⁹. Indeed, perhaps, if you know that I'm an Ashkenazi Jew, and you proceed to believe that I'm a poor tipper, this belief already wrongs me, it already has a moral status. If so, moral considerations apply directly to beliefs. This isn't impurism just yet, but impurism is now very close: If you want to defend a view according to which epistemology does not require wrongdoing (or wrongbelieving), you must allow the moral considerations to affect epistemic status as well. You must, that is, accept moral encroachment.

For these reasons, then, and possibly for others as well, it's not surprising to see a rise in the attention moral encroachment – perhaps alongside other forms of impurism – has been getting recently30.

2.2. Impurism and Statistical Resentment

Impurism - in particular, moral encroachment - is not usually motivated by thoughts of statistical resentment. Still, once the view is out there, as it where, applying it to our case seems straightforward:

Resenting someone for purportedly doing something they did not in fact do is serious moral business. The stakes, then, are not low. So what is needed for knowing that the relevant colleague badmouthed you is *especially* strong evidence, and perhaps the statistical evidence is just not good enough. Or perhaps the stakes dictate that the needed evidence has to be of a certain $kind^{31}$.

Similarly, perhaps there's sufficient practical reason not to treat your colleague in ways that

²⁸ Moss (2018b).

²⁹ The relation between doxastic wrongdoing and impurism takes central stage for Basu (2019a, 2019b). See also Schroeder (2018) and Basu and Schroeder (2019). For the record, we reject doxastic wrongdoing, though we don't rely on this rejection here. See our "There Is No Such Thing as Doxastic Wrongdoing" (ms.).

³⁰ See Pace (2011), Enoch (2016; 2017), Moss (2018a; 2018b). See also Gardiner (2018b), and the references there (mostly in the first three sections).

³¹ In the pragmatic encroachment literature, usually the issue is the *strength* of the evidence, not its kind (Moss (2018a) is a prominent exception). But perhaps this is just an oversight. For a similar suggestion (in the political context), see Enoch (2017, 152).

would have been justified had (you had good evidence that) they been badmouthing you, like not inviting them to the bar next time people are going. If so, and if there's a fairly strong link between what one is entitled to treat as a premise in one's practical reasoning and what one knows, perhaps this explains why you can't know, on the basis of the statistical evidence, that they badmouthed you.

If moral considerations partly determine which alternatives are salient, so that ruling them out is necessary for knowledge or epistemic justification, then perhaps because of such moral reasons, the possibility that your colleague is not a typical member of the relevant reference class is rendered salient because of the wrongness of false resentment (or some such). Perhaps this is why you don't know, or don't justifiably believe, based only on the statistical evidence, that they badmouthed you, and perhaps this in turn explains why statistical resentment is unacceptable.

And perhaps, if beliefs themselves sometimes morally wrong, you are morally wronging your colleague, perhaps because you're failing to treat him as an individual in some sense, by resenting him based on just the statistical evidence. And if such moral status can affect epistemic status, perhaps this is also why you don't know, and are not even justified in believing, that he has been badmouthing you.

2.3. Why Impurism is not the Way to Go (Here)

Still, going impurist is not the way to solve the problem of statistical resentment. One way to show this would be to reject impurism in its entirety. Now, we do think that impurism – especially in its more radical forms³² – faces serious, possibly devastating, challenges and difficulties³³, and some of these will emerge below. But we cannot here engage a full evaluation of impurism. Instead, we focus on the special ways in which impurism is unpromising as a strategy of coping with the problem of statistical resentment.

One problem is that it's not clear why, assuming impurism, statistically-based beliefs are special. Assume that the stakes really are high, for instance, and so that only especially good evidence can suffice for a justified belief of a specific colleague that they badmouthed you. Why think, though, that statistical evidence – at least if the statistics are sufficiently robust – may not be good enough even for this higher bar? We're not saying that this question cannot be answered; we're just emphasizing that this question should sound familiar to you. Wondering why it is that vis-à-vis resentment statistical evidence is more morally problematic than some other kinds of evidence is, after all, very close to the question why it is that statistical resentment is problematic, which was the question we started from. So it's hard to see how progress has been made. A similar point applies to Moss's (2018b) insisting on the need to be able to rule out the possibility that the relevant person is an atypical member of the relevant group as a necessary condition for knowledge, while resisting a similar condition in non-person cases.

 $^{^{32}}$ For one distinction between radical and moderate impurism, see Fritz (2020).

³³ For a critical discussion of moral encroachment, which we are largely sympathetic with, see Gardiner (2018b, section 5). See also the references there. For some examples of critical evaluation of pragmatic encroachment, see Reed (2010), Brown (2013), Anderson & Hawthorne (2018a; 2018b), and Jackson (2019).

Why is it, we can ask, that a-typicality is a salient possibility in the former case but not in the latter? Perhaps Moss can answer this question. What we're emphasizing here is, first, that she has to, and second, that when she does, it will be this *further* answer that is lifting all the weight. Answering this further question is *very* close to answering the question why it is that statistical evidence is so problematic in some cases and not in others. So it's hard to see what progress has been achieved by putting things in terms of the need to rule out relevant alternatives³⁴.

Similarly, while there may be a sense in which you're wronging your colleague if you believe on insufficient evidence that he has been badmouthing you, we need an answer to the question what it is that makes statistical evidence worse, in our context, than non-statistical evidence, and this means we're not sufficiently far from where we started.

Another problem is one of generality. Even if impurism helps with some cases of statistical resentment, it is not promising at all with regard to some of the other cases that exemplify similar phenomenology. For instance, while the stakes in PRE-PRIDE or TRUTH TO POWER may not be entirely low, they are certainly not as high as they are in statistical resentment cases, and yet statistical admiration doesn't seem more easily justifiable than statistical resentment. Even LOTTERY – if you agree that it exhibits a similar phenomenology to that of statistical resentment cases - counts against an impurist explanation, for it's not clear there are pragmatic stakes involved in a relevant way (although everything here is contested³⁵). Furthermore, some of the underlying motivations for impurism don't even apply to THE CENTRAL RESENTMENT CASE. Basu and Schroeder (2019), for instance, emphasize the ways in which our evidence, in many cases having to do with race, is likely to be tainted by past injustices, as a way of motivating moral encroachment. But the problem of statistical evidence – as exemplified in THE CENTRAL RESENTMENT CASE - need not be about race or any other category to which this line of reasoning applies.

The two objections just discussed take the impurist apparatus on board pretty much as is, and raise problems for applying it in order to solve the problem of statistical resentment. We now want to show how the details of that apparatus may be important here, and may raise even more problems. For this, we need to be explicit about what epistemic status precisely it is that is taken to be impure. All impurists, we take it, accept knowledge-impurism, that is, the claim that pragmatic and moral considerations – and not just evidence – may be relevant to whether or not someone knows something. Justification-impurism is the claim that whether or not one's belief is (epistemically) justified may depend also on pragmatic or moral considerations. Given the uncontroversially strong relations between knowledge and epistemic justification³⁶, it's very hard to see how someone could accept knowledge-impurism and reject justification-impurism.

³⁴ We briefly revisit Moss's theory in the next section.

³⁵ For a comprehensive impurist (Subject Sensitive Invariantist) account regarding lack of knowledge and lotteries, see Hawthorne (2004). For a critique of the notion of stakes (or stakes dependence) see Anderson and Hawthorne (2018a, 2018b).

³⁶ It is, of course, controversial what these relations precisely are. But it's not as controversial that knowledge entails justified belief.

For our purposes the crucial distinction is between both knowledge- and justification-impurism on one side, and *credence-impurism* on the other: Credence-impurists, if there are any³⁷, think that even the justified credence in a proposition may depend on non-evidential, pragmatic or moral, features. Credence-impurism is a *much* more radical view, and is arguably not supported by the motivations that underlie knowledge- and justification- impurism. Indeed, one natural way of understanding knowledge- and justification-impurism is as rejecting credence-impurism, but insisting that the bar for the credence needed (for knowledge or justification) varies with stakes. Perhaps unsurprisingly, then, the literature for the most part rejects credence-impurism³⁸.

Armed with this distinction, we now want to argue that credence-impurism is implausible; that rejecting credence-impurism destabilizes justification- and knowledge-impurism³⁹; and that the resulting impurism is *especially* implausible as a solution to the problem of statistical resentment.

First, then, the implausibility of credence-impurism. Suppose you know – you have conclusive evidence – that 20 of the 100 balls in an urn are red. You draw a ball at random. How high should your credence be (before looking) that the ball is red? The answer, clearly, is 0.2. It matters not what – pragmatically or morally – hangs on having selected a red ball⁴⁰. Given the assumption of randomness, any other credence will clearly be irrational.⁴¹ This is strong intuitive evidence against credence-impurism. And this applies to our resentment case as well, of course: Given conclusive evidence about the high percentage of colleagues badmouthing their colleagues in your work environment, it seems clear that your credence that a specific (randomly selected!) colleague has badmouthed you should be equally high⁴², independently of pragmatic or moral stakes. Now, perhaps a supporter of credence-impurism may insist that

³⁷ The closest we know of in the literature is Gao (2019). But even Gao, for the most part, endorses only what she calls "creedal pragmatism" as a *descriptive* thesis (about how people's credences in fact change in the face of a change in the stakes). For reasons we cannot get into here, we find her initial normative discussion (section 3) unpromising.

³⁸ For explicit rejections, see Buchak (2014, 303-4) and Ross and Schroeder (2014, 260). For an implicit one, see Brown (2014). For some discussion, see Gau (2019), Fritz (2020), Fritz and Jackson (forthcoming).

Moss's (2018b) view is special, in that she thinks that credences too are encroached on – but only in the sense that pragmatic considerations affect whether or not a given credence amounts to knowledge (as she insists some credences do). Moss too, however, stops short of accepting credence-impurism.

³⁹ Jackson (2019) relies on belief-credence dualism in order to reject impurism, but she does things in a very different way from ours.

⁴⁰ Despite Marušić's (2015) resisting talk of *stakes*, we think that the objection in the text, suitably modified, applies to his view as well.

⁴¹ Bolinger (2020) discusses the credence case, and mentions many real-life considerations – of the kind we proceed to discuss in section 4 (especially section 4.2) – that complicate matters. But she seems to agree that when no such complicating considerations apply, credences should just follow probabilities.

We can also rely on Lewis's Principal Principle that the conditional credence of event A given the objective chance that A is x, is x: cr(A|ch(A)=x)=x (assuming (ch(A)=x)>0 and that the background evidence is all about events before A takes place).

 $^{^{42}}$ At least if you have any credence at all. The argument in the text is not affected by the possibility of permissible credence-suspension.

there is something special about the URN case, that perhaps, even if in this case credence is immune to encroachment, in other cases it may not be. It's not an impossible line to take – but it's not easy either, and it's going to be especially hard to make a convincing case that such a restriction of credence-impurism is not objectionably ad-hoc.

So credence-impurism is highly implausible. But if this is so, this places serious pressure on knowledge- and justification-impurism as well. One way of developing this point – one that is too general for our purposes here (though we revisit it in the next section) – proceeds via some theoretically motivated relations between justified credences and justified beliefs⁴³. Here is another way, centered on statistical evidence cases. Impurist views all rely on a close connection between beliefs and action. This is clearest for those utilizing a knowledge-action link, which we'll focus on for the rest of this paragraph. On such views, recall, one is only entitled to act, roughly, on propositions one knows. So think again about LOTTERY: Perhaps one is not entitled to act on the proposition that one's ticket won't win, and perhaps this shows that one doesn't know that the ticket won't win. Still, given the rejection of credence-impurism, one is justified in having a very high credence in that proposition, and also – it seems – in believing that the ticket is highly unlikely to win. Indeed, one may know that (why wouldn't one?). And this means that one can be justified in acting on this, probabilistically hedged, proposition - not that the ticket won't win, but that it's extremely likely that it won't. And then, any (or almost any) action one would have been justified in performing based on practically reasoning from the proposition that the ticket won't win, one will be justified in performing based on practically reasoning from the proposition that the ticket won't in all likelihood win. Indeed, without sanctioning action based on the probabilistically hedged belief, the view is a non-starter: For instance, it would not allow your doctor to prescribe medicine even when the available statistical evidence is compelling. So an attempt to offer an impurist solution to the puzzle of statistical resentment must insist that morally loaded matters are relevantly different from (other) practical ones, so that while action based on probabilistically hedged beliefs can be perfectly justified - perhaps in bank cases, for instance - this is not the case of many morally loaded cases. It's not clear how the impurist can justify this asymmetry – at least not without saying much more, indeed, not without offering what is essentially a different solution to the puzzle of statistical resentment.

So while impurist views utilizing the knowledge-action link have the resources to explain some of the central cases, still, given the relevance of probabilistically hedged beliefs, and given the rejection of credence-impurism, it becomes very hard to see how impurism can accommodate many of the other cases discussed above. After all, we have already seen that there is a difference between This child is probably good at math and this car is probably fuel efficient. And, as we've also seen, the impurist relying on the knowledge-action link has to accept the justifiability of actions based on probabilistically hedged beliefs, but makes an exception for

⁴³ Williamson (forthcoming-a) rejects any systematic relation between belief and credence: On his view, one may have a fully justified credence 1 that p, while it would nevertheless be irrational for one to believe that p. Even Williamson, though, doesn't deny that in such a case a probabilistically hedged belief - that it's highly likely that p - would be rational. The discussion of credences and of probabilistically hedged beliefs below shows why this suffices for our purposes.

morally charged cases. The impurist apparatus itself does not suffice to accommodate all this.

The relevance of credences as a challenge to Impurist solutions to the puzzle of statistical resentment becomes especially clear when thinking about the view according to which moral standards apply directly to beliefs, though the point generalizes to any Impurism suggestion here: Suppose, then, that believing that the Ashkenazy man at the next table is a poor tipper – based on just the statistical evidence – morally wrongs him. Can this explain the problem with statistical resentment? It's hard to see how, given the falsity of credence-impurism. After all, having a high credence in him being a poor tipper (after finding out he's of Ashkenazi descent), or indeed coming to believe that he's *probably* a bad tipper, seems intuitively just as offensive as forming the belief that he is⁴⁴. And this shows yet again that the prospects of an impurist view – *any* impurist view – about justification and knowledge successfully explaining the oddness of statistical resentment is undermined by the unacceptability of credence-impurism⁴⁵.

3. <u>Knowledge</u>

But perhaps it's all just about knowledge⁴⁶. Perhaps the rule for resentment, and admiration, and pride, is that one can only rely on what one knows. The problem with statistical resentment, on this suggestion, is that the statistical evidence does not ground knowledge (that the specific colleague you're thinking of has badmouthed you), and that without such knowledge, resentment is inappropriate. This line of thought, then, takes very seriously indeed the relation between The Central Resentment Case and Lottery, and offers a similar diagnosis for both: lack of knowledge. In this section we focus on a purist version of a knowledge-based account of statistical resentment (for we've already argued against the impurist suggestions), arguing that despite some advantages, at the end of the day a knowledge-account cannot suffice as a solution to the puzzle of statistical resentment.

One natural way to motivate a knowledge-based account in our context is to rely on knowledge being the norm for belief – so that a belief is only appropriate or even justified when it amounts to knowledge⁴⁷. This is a promising line to take in our context, because beliefs are arguably a constitutive part of each of our many examples. If beliefs are involved in all of the explananda, and knowledge is the norm for belief, it may be hoped that knowledge plays a

⁴⁴ For a related discussion – also re Basu and Schroeder – see Gardiner (2018b ,179, fn 28).

⁴⁵ Buchak (2014) takes it for granted that statistical evidence can unproblematically support credence and probabilistically hedged beliefs; she takes the unacceptability of basing blame on statistical evidence as evidence that beliefs are not best understood in terms of credences. Buchak is thus a kindred spirit – but note that, as we show throughout this paper, we don't think that putting things in terms of beliefs solves the puzzle of statistical resentment.

⁴⁶ For some discussions along these lines, see Blome-Tillman (2017), LittleJohn (2017), and Gardiner (forthcoming).

⁴⁷ Williamson (2000, 256) seems committed to the idea that belief in the loss of a lottery ticket is less than perfectly rational and is at least sympathetic to a knowledge norm of belief (see Hawthorne, Rothchild, and Spectre (2016) for discussion). He clearly endorses and defends a knowledge norm of belief in more recent work: Williamson (2011, 2013, 2017, forthcoming-b, forthcoming-a). This is not to suggest that Williamson is the only one to endorse a knowledge norm of belief. Several theorists – e.g. Moss (2018a, 2018b) – have accepted his (or similar) norms, some have advanced knowledge norms of belief independently. See, for instance, Bird (2007).

crucial role in explaining the full scope of what we are out to explain. In response, we want to briefly argue – utilizing some of the discussion of credences above – that a knowledge norm for belief is implausible, and that even putting such general doubts aside, such an understanding of the relation between beliefs and knowledge will not work in our context, because it cannot generate the desired generality.

Recall the lesson about credences from the previous section. It's very hard to deny that statistical evidence affects at least justified credences and probabilistically hedged beliefs (and knowledge), both in purely epistemic cases (as in LOTTERY), and in practically relevant ones. One family of problems for knowledge-based accounts starts with this observation, and then notes how difficult it is to make plausible the needed claims about beliefs and knowledge, given what we already know about credences. There seems to be strong linguistic evidence that beliefattributions behave much more like attributions of fairly low (though above .5) confidence levels or credences than they do like attributions of certainty, or something more resembling the epistemologist's, highly theory-laden, "all-out belief". There is, for instance, something at the very least awkward – perhaps Moore-paradoxical – in saying something like "John thinks that it is probably raining, but it's not that he thinks that it is."48. A natural explanation of the oddity is that thinking-p-probable is not that different from thinking-that-p, or indeed, believing-thatp⁴⁹. If beliefs are similar – perhaps identical – to thinking-sufficiently-likely, and if thinkingsufficiently-likely can unproblematically be justified based on statistical evidence (because credences certainly can, and because of the close relation between credences and thinking-likely), then beliefs too can be justified based on statistical evidence. Needless to say, such linguistic evidence may be doubted, but it seems at least to shift the burden⁵⁰ onto those introducing a very sharp divide between credences and beliefs, as all supporters of a knowledge-based account must⁵¹. In other words, knowledge does not seem to be the norm for credences, so if beliefs are very closely tied to credences, it's very hard to accept knowledge as the norm for belief⁵².

⁴⁸ Applied to the cases that interest us here we have sentences such as "John is confident that Mary is badmouthing him, but it's not that he believes that she is." This is, we think, Moore paradoxical. Also, "confident that" doesn't neg-raise: From "I'm not confident that p" it doesn't follow that "I'm confident that not-p". A central claim in Hawthorne, Rothschild and Spectre (2016) is that "believe that" does neg-raise, and that neg-raising is an indication of relatively weak epistemic standards. If so, "confident that" is more evidentially demanding than "believe that". And notice that there's reason to believe that "think that" (as in the text) and "believes that" behave alike for current purposes.

⁴⁹ We are relying here on Hawthorne, Rothschild, and Spectre (2016). See also Williamson's (forthcoming-a) response to the former. In Williamson (2017, n. 11, 160-170) he takes a more neutral attitude toward the idea that belief is weak and suggests that perhaps the necessary and sufficient condition for knowledge is rational sureness.

⁵⁰ Hawthorne, Rothchild, and Spectre (2016) argue that the burden of proof has been on the knowledge side all along. If indeed believing that your ticket won't win the lottery is unjustified, it's hard to explain the epistemological attention Kyburg's (1961) lottery paradox has been – and still is – getting.

 $^{^{51}}$ And not only them. Martin Smith's (2016) interesting discussion of *normic support* also invokes a sharp divide between beliefs and credences, and so he too is vulnerable to the point in the text.

⁵² Moss (2018-b) as noted earlier, stops short of endorsing a knowledge norm of credence levels. With regard to belief, though, she clearly wants to defend a knowledge norm. For an exchange on this point see Moss (2019),

Furthermore, the suggestion that beliefs can never be epistemically justified based on statistical evidence (a suggestion that follows from a knowledge norm for belief, together with the claim that statistical evidence cannot ground knowledge) seems to run in the face of fairly clear counterexamples. Do you believe, for instance, that your close friends have the same credit card four-digit PIN code as you do? We're pretty sure the answer is "no", and not just in the sense that you don't have that belief – you believe its negation⁵³. Furthermore, this belief of yours seems entirely justified (indeed, it may even, if true, amount to knowledge). But your only evidence for it is statistical, and (depending on how many friends we are talking about) not even that strong statistical evidence (it's much weaker, for instance, than the statistical evidence you have in LOTTERY).

For these reasons, then, the thought that beliefs cannot be justified based on just statistical evidence – central to accounts of resentment and related phenomena that endorse a knowledge norm for beliefs – is implausible. But, we now want to argue, even if knowledge-based accounts can somehow overcome this problem, they can still not offer a satisfactory solution to the problem of statistical resentment, for a reason that will – at this stage in the paper – sound familiar: They are not sufficiently general.

Suppose, then, that resentment is only in place when based on knowledge (of guilt, say, or of the relevant wrong, or in our case, of the fact that the relevant colleague badmouthed you)⁵⁴.

Even with this assumption in place, it remains entirely unclear what a knowledge-based account can say about DNA cases. Furthermore, a knowledge-based account cannot explain the legal cases, unless, in a slogan, the law should *care* about knowledge, prioritizing it at least sometimes above other things the law can secure (like deterrence of crimes). But this thought seems guilty of "knowledge-fetishism" – it seems unjustifiable to pay a price in other goods the law may be able to deliver just in order to secure some desired epistemic status (other than truth or accuracy) for legal findings⁵⁵.

Next, recall PHYSICIAN, where acting on statistical evidence (in diagnosing a patient) seems not only unproblematic, but required. In the face of such cases it becomes very hard to maintain both the knowledge-action link, and the claim that statistical evidence cannot establish knowledge. Given that the physician ought to treat the patient based on the statistical evidence, we must either accept statistical evidence here as sufficient for knowledge (that the patient has condition so-and-so), or accept that action is justified without knowledge (that the patient has condition so-and-so), or both. And this means that we must reject as over-general the

Williamson (forthcoming-a), and Rothschild (2020). Other views have been proposed along the Hawthorne, Rothchild, & Spectre (2016) line: Dorst (2019), Holguín (MS).

⁵³ Again see the discussion of "neg-raising" in Hawthorne, Rothschild and Spectre (2016). Note that this example is problematic on Smith's (2016) normic support view.

⁵⁴ This may come as good news for a Sensitivity-based account. If knowledge is needed for appropriate resentment, then, and if insensitive belief is (even just almost) never knowledge, then given that beliefs based on statistical evidence are insensitive, it follows that statistical resentment is never appropriate. And a Sensitivity-based account has the resources not just to assert that the relevant beliefs don't amount to knowledge, but also to offer a plausible explanation why this is so.

⁵⁵ See Enoch, Spectre and Fisher (2012).

knowledge-based attempt to diagnose what's wrong with statistical resentment. True, as noted above, the physician may have *other* relevant knowledge, like the knowledge that the patient *probably* has condition so-and-so. And this allows the knowledge-theorist to save the knowledge-action link (though with regard to probabilistically hedged propositions, and with the generalization problem we are about to discuss). But – regardless of the merits of such a move more generally – it can't help knowledge-based explanations of the problem with statistical resentment. This is so, because while in PHYSICIAN the probabilistically hedged belief (that the patient probably has condition so-and-so) was unproblematically justified based on the statistical evidence, in The Central Resentment Case, as we've seen, the probabilistically hedged belief (that your colleague has probably badmouthed you), when based on statistical evidence, *is* problematic, indeed, almost as problematic as the unhedged belief.

Also, within purism, it is very hard to distinguish between MATH and CARS. And we've already rejected impurist versions of a knowledge account when rejecting impurism as a solution to the statistical resentment puzzle, in the previous section.

Or consider the case of SURPRISE. As you may recall from section 2, the appropriateness of being surprised when you find out that your lottery ticket has after all won shows how statistical evidence, even if not sufficient for knowledge, suffices for shaping our expectations. Still, being surprised to find out, about a specific colleague, that they have not in fact been badmouthing you – when the opposite expectation had been shaped just by the statistical evidence – remains problematic. Indeed, we think it is problematic in the same way statistical resentment in this case would be. This cannot be explained by a knowledge-based account. In this way too, then, the knowledge-based account is not sufficiently general to explain everything that has to be explained here. Lastly, consider what seems to us a potentially clear-cut counterexample to knowledge accounts, whatever their details. PRE-RESENTMENT and pre-pride cases are problematic, indeed problematic in the statistical resentment kind of way. But – unless one is willing to endorse wholesale skepticism about knowledge and the rationality of beliefs about people's future actions, which seems excessive - nothing about knowledge can explain why. Arguably, we can sometimes know that someone will badmouth us. Even when we do, though, (pre-)resenting them is out of place. So nothing about knowledge explains the problematic nature of statistical resentment. And notice that this counterexample – again, dialectically conclusive, as far as we can see, against everyone but future-skeptics - applies to knowledge accounts whatever their details, purist or impurist⁵⁶.

4. <u>Conservative Epistemological Explanations: Some Progress</u>

Having rejected impurist attempts at making progress on statistical resentment, as well as purist

⁵⁶ So in particular, this reinforces the objection from the previous section against Moss's (2018b) impurist, relevant-alternatives version of a knowledge-based account. Note also that while in some contexts, relevant alternatives views are associated with contextualist views, this is not so in our context, where all of the relevant factors (the statistical and the non-statistical evidence, in the different cases) are those believers are sensitive to, not knowledge attributors. So the relevant kind of sensitivity is subject-sensitivity, not context-sensitivity.

knowledge-based accounts⁵⁷, in this section we pursue perhaps the most conservative route: We consider to what extent there is a way of accommodating the relevant intuitions about the specific cases – or anyway, most of them – without accepting too radical revisions of orthodox epistemology. We discuss three attempts in this spirit: one in terms of the phenomenon of screening off, one in terms of worries about representativeness and randomness, and one in terms of concerns about how it is that one got to have all and only the evidence one has. Our bottom line is going to be that such conservative explanations go some way towards solving the problem of statistical resentment, at least with regard to many of the relevant cases, but the solution they offer leaves some things to be desired⁵⁸.

4.1. Screening Off

Clean cases are easier to stipulate than to find in the real world. For the resentment case to be sufficiently clean, we stipulated that in it resentment is based on statistical evidence alone. But this is not how the world works. Usually – perhaps even always – you have so much more information about the relevant colleague. So even purely epistemically – before doing ethics – it's unacceptable for you to ignore all this other evidence. And this other evidence may change things significantly here, with regard to the appropriateness of resentment.

True, but not enough. The intuition against statistical evidence seems to be stronger than that. It's not just that there's something suspicious about resentment entirely based on statistical evidence. Rather, it's that this kind of statistical evidence shouldn't play *any part* in supporting resentment, at least in many cases. Suppose you have a joint appointment, in a law school and a philosophy department, and suppose that badmouthing colleagues is very common in the law school, but not in the philosophy department. Suppose further you have two colleagues, one from the law school and one from the philosophy department, and that the relevant individual, non-statistical evidence about them is similar, and not on its own sufficient for justified resentment (JOINT APPOINTMENT). The thought that you could resent your law school colleague but not your philosophy department colleague because of the difference the statistical evidence makes – while taking into account all the other evidence as well – seems just as problematic as the case of resentment based on naked statistical evidence. In order to accommodate this intuition we need to do more, then, than merely insist on taking into account all relevant evidence.

A natural thought to have here is to think about the *interaction* between the statistical and the non-statistical evidence. For in typical resentment cases, the non-statistical evidence is ever so much *thicker*, so much richer in information, compared to the rather shallow statistical evidence. It is very plausible to think that any information included in the statistical evidence is

⁵⁷ Throughout – including in the context of discussing knowledge – we've been putting things in terms of beliefs. Bolinger (2020) discusses, in a similar context, *acceptance*, an attitude that combines epistemic and pragmatic features. In the text we assume that resentment, and perhaps related phenomena as well, go with belief rather than with acceptance (to the extent that the two are distinct). One reason Bolinger's acceptance won't do here is that resentment cases need not have any further practical implications. Moreover, going for acceptance does not immune one from Moorean results. The following sentence, we submit, is Mooreanly incoherent: "He believes she'll win but it's not that he accepts it."

⁵⁸ Much of the discussion in this section is in the general spirit of Gardiner's (2018b, mostly in section 6) discussion. But at the end of the day, Gardiner seems less tentative about the success of this project than we are.

already included in the much richer non-statistical evidence. If so, the non-statistical evidence you have about the person you may resent *screens off* the statistical evidence, and so the statistical evidence should be ignored, once the non-statistical evidence is taken into account.

Intuitively, one piece of evidence (fully) screens-off another when the former already includes any information that's a part of the latter. More precisely, if we're considering theory T, and if E1 and E2 are pieces of evidence each of which confirms T (that is, P(T|E1)>P(T), and P(T|E2)>P(T)) then E1 fully screens off E2 only if P(T|E1&E2)=P(T|E1). In such cases, having factored in the screening off evidence E1, it would be a mistake to then change credence in T based on the screened-off E2. Whatever information (say, in corroborating T) E2 conveys was already taken into account in factoring in E1. Now taking E2 as a reason to further increase credence in T will amount to double counting.

Returning to resentment, then. It seems plausible that at least in the majority of resentment cases, the more personal, non-statistical evidence you have about the relevant person doesn't just add to the statistical evidence. Rather, being so much richer in information and sensitivity to specific details, it screens off the statistical evidence. Given everything you know about your law school colleague, you shouldn't now factor in the statistical evidence as well. Whatever information the statistical evidence may be thought to disclose, it's now been fully incorporated in the non-statistical evidence. So having taken into account all other evidence, you should now ignore the statistical evidence entirely, with regard to both your law school colleague and your philosophy department colleague. And this, recall, was the desired result.

It is hard to deny that screening off is a major part of the picture here. Resentment is often personal, and the other attitudes and emotions that are not comfortably based on statistical evidence are equally personal, it seems, at least in paradigmatic cases. And because they are personal they occur in an information-rich environment. The richer the environment in details about the relevant individual, the more likely it is that the statistical evidence is screened off in this way, and so should be ignored⁵⁹. The screening-off point also succeeds in accommodating some of the distinctions in our examples. For instance, if you've been in the same math class with me for a while, you have a lot of information about my math skills, and then the fact that I'm of Asian descent doesn't give you any information not already included in the evidence you've considered, and should therefore be ignored. But you are unlikely to have such screening-off information about an Asian car you're considering buying⁶⁰. In a similar vein, the statistical evidence is not screened off by any more specific evidence in the typical case of a physician diagnosing you (and when it is, then of course the physician should ignore it!). Furthermore, this way of explaining what's problematic with statistical resentment is refreshingly simple – it doesn't employ problematic moral theses, it doesn't rely on intricate relations between the practical and the theoretical, it doesn't even rely on any too-fancy or controversial epistemological suggestion. It's all about taking into account, in a probabilistically respectable way, all and only relevant evidence.

⁵⁹ This observation coheres nicely also with the suggestion that the fact that these attitudes are usually de re makes a difference here.

⁶⁰ Gardiner (2018b, 185-6) emphasizes a similar point.

Still, tempting though this line of thought is, it doesn't give us the full account of the problem with statistical resentment. We offer three (related) reasons for rejecting it as an attempt at such a full account.

First, resentment is not *always* or *necessarily* personal in the sense of presupposing some fairly thick, information-rich relationship. If a complete stranger cuts me in line (STRANGER) it seems that some mild resentment may be appropriate. Nothing more is needed – no shared history, or intimacy of any kind. Still, it would have been inappropriate for me to resent him or her on purely statistical evidence, evidence that is not (as we've just stipulated the case) screened off by any other evidence. So thoughts about screening off do not have the needed scope – they don't explain what's problematic with statistical evidence across all the relevant cases.

Second, while what's been said about the real world and how in it statistical evidence will often be screened off is true, we can still insist on the relevance of hypothetical cases. There's no incoherence in thought experiments about cases with no such screening-off effects, and if the only explanation we have of what's wrong with statistical resentment is grounded in the phenomenon of screening off, then about such cases we'll have to conclude that statistical resentment in them is entirely appropriate. This is an unwelcome result.

Which brings us to the third point: it just doesn't seem plausible that this is *what's* problematic about statistical resentment. Intuitively it's not the (very common, and important) probabilistic errors that explain the problem with statistical resentment. This deep problem – that this is not the kind of thing in virtue of which statistical resentment is inappropriate – gives rise to the symptoms discussed in the previous two paragraphs.

4.2. Randomness

From the fact that, say, 90% of your colleagues speak ill of each other behind their back, together with the fact that *this* person is a colleague of yours, it doesn't follow – in any sense of probability – that there's a probability of 90% that this person has badmouthed you (or others). For that to follow we need another premise, namely, that she or he are a *typical* member of the relevant set⁶¹. And with some of the cases we've been discussing, this may not be the case. With yet others, it may not be known whether or not this is the case.

Perhaps in some cases, one is entitled to take the absence of evidence of a-typicality as sufficient here. Perhaps, for instance, if one is considering buying a specific Asian car, and there's no evidence that it's in any way an atypical Asian car, one is entitled to apply the statistics to it, and to believe that it's (likely to be) energy-efficient. But first, even this is contestable. Perhaps what's always needed for the inference to be justified is some positive evidence that the specific case was picked randomly from the relevant reference class. If the specific car was picked randomly from the set of Asian cars (or from all cars) then such inference may be licensed. If, however, the car was non-randomly chosen for display by the car dealership, this may change things radically here (perhaps the dealership wants to sell the less sought-after cars first). And

⁶¹ Or perhaps, if we're working with some subjective or epistemic notion of probability, what's needed is that you *know* or *reasonably believe* that he is typical.

if one doesn't know whether the car was chosen randomly from the relevant reference class – well, then things become really unclear⁶². All of these complications arise for the case of statistical resentment as well. But what if, say, the relevant colleague was "chosen" for statistical resentment precisely because in his case there's no non-statistical evidence available, one way or another?⁶³ This may change the relevant probabilities significantly. And what if you don't know whether the colleague was picked randomly? Then all the complications just mentioned apply, and it's very hard to know how to proceed.

Second, even if randomness has a kind of default-status in some cases, it's not clear it has this status in the cases of interest to us. Recall the ways in which persons seem special here. While you may not know that your lottery ticket is not a winning one, you can unproblematically justifiably believe and also know that it's highly unlikely to win. But in many of the person-involving cases, things are different. There's something suspicious about believing that the child of Asian descent is probably good at math⁶⁴, and there's *certainly* something suspicious about similar beliefs (even with the "probably" qualifier) that apply a negative but statistically grounded stereotype to an individual (about a specific Ashkenazi Jew, for instance, that he's a bad tipper). And this seems to indicate that even if there's no problem in assuming (because we know that the mechanism for picking a winner is random) that a specific lottery ticket has a very low probability of winning, there very well may be a problem in assuming that a specific person is a typical member of the relevant reference class (or that your access to them is sufficiently random-like).

In a way, we've come close to square one: Asking why it is that there's something suspicious about even the probabilistic belief in some cases but not others seems very close to asking why it is that statistical evidence is problematic (in the cases in which it is).

Regardless, the points from the previous subsection apply, and we can afford to be quick on them now: Even if thoughts about typicality and randomness can explain some of the cases, they can't explain all, certainly not cleaned-up hypothetical cases, and most importantly – the fairly robust intuitions we started with are not just that statistical resentment is problematic, but also that it's not for *such* reasons that it's problematic.

4.3. Why Do You Even Have the Evidence that You Have?

Sometimes the problem with a belief – think of certain well-supported arrogant beliefs, for instance – is best seen not when focusing on the specific moment when the belief (or other epistemic attitude) is held, but rather considering the entire epistemic process that led to the belief and that sustains it. We can ask questions about how we got here, how you came to have all and only the evidence you now have, and we can search for a flaw there. Perhaps, then, if you resent a colleague believing that they badmouthed you, based on purely statistical evidence, we may wonder about how it is that you've ended up here, in this evidential state. How is it that you've

⁶² See Karlander and Spectre (2010), especially the appendix.

 $^{^{63}}$ See here Posner's (1999, 1509) similar suggestion about the legal treatment of statistical evidence.

 $^{^{64}}$ Or, if Moss (2018b) is right, in believing about someone about whose sexuality you know nothing in particular, that he is probably straight.

come to have the statistical evidence for that belief? Why have you even wondered, of a specific co-worker (and perhaps not of others?) whether they badmouthed you? Getting the statistical evidence may not be easy. Why did you even seek it? And why did you stop there? After all, often it won't be that hard to obtain further evidence (say, for starters, to ask them). The combination – perhaps especially in the case of people you know – of obtaining the statistical evidence about them, and not making further inquiries seeking non-statistical evidence (or even further statistics) – this combination is at the very least suspicious⁶⁵. Perhaps this is what explains, then, the problem with statistical resentment.

Clearly, this kind of story has an explanatory role to play in accounting for the oddity of statistical resentment. At least in some cases, resenting based purely on statistical evidence will be problematic at least partly due to reasons having to do with worries about how it is that you came to have all and only the evidence you do in fact have. And this kind of story may explain some of the other cases we've been discussing: Perhaps, for instance, there is something problematic if one stops inquiry and settles on the belief that a specific child of Asian descent is good at math based on just the statistical evidence (and how did we get that statistical evidence in the first place?), but perhaps there's nothing problematic in settling for the statistical evidence in support of the claim that Asian cars are fuel-efficient. If so, this may⁶⁶ nicely explain why there's nothing problematic with the latter belief, but there is with the former. Still, it's also clear that this can't be the full story of statistical resentment. There are cases – there are certainly hypothetical cases, and this is good enough - where we can stipulate that there are no problems in how you got to have the evidence you have, and in those cases too statistical resentment still seems problematic. Perhaps, for instance, you came across the statistical evidence about how common it is for your colleagues to badmouth each other without seeking it (maybe it was there in some report you were required to read about department climate), and maybe this happened while you're away on vacation somewhere without internet connection. In such a case, there is no flaw in how you came to possess all and only the evidence you have - you had to read the report, and you can't for now ask follow-ups or seek more evidence. Still, resenting a specific colleague based on the statistical evidence seems just as problematic in this case⁶⁷. And arguably

⁶⁵ What if the statistics are very strong? Wouldn't that make it rational not to further investigate the issue? Actually, in many cases strong statistical evidence is *unstable*, in the sense that it's fairly easy to come by evidence that will outweigh it, and then, stopping inquiry may not be justified. For instance, the very strong statistical evidence you have that your lottery ticket won't win is easily outweighed by the phone call from the lottery company saying that you did win. (The denominator of the Bayes factor is very small because a mistaken call *to you* is extremely unlikely forcing the posterior probability way up.) Certainly, you shouldn't stop inquiry and refuse to take their call. See here Buchak's talk of resilience (2014, 294). And for discussions of stability and revisability in our context, see Schroeder (2018b), Gardiner (2018b, 189), and Bolinger (forthcoming). Resilience (or Stability) is developed as an overall account of rational belief in Leitgeb (2014). For a good and less detailed discussion of resilience see Joyce (2005) and the references there. Leitgeb doesn't argue that lottery beliefs are irrational in all contexts; in fact, his view is designed to vindicate their rationality at least sometimes.

⁶⁶ Or the explanatory order may proceed in the opposite direction. We're genuinely not sure.

⁶⁷ A reviewer pointed out that the symmetry we're assuming in the text between the evidence-related practices leading up to the belief and those that follow (say, a decision to stop one's inquiry) cannot be taken for granted. This may be so, but we believe that the point in the text, with sufficiently creatively stipulated hypothetical cases, holds even if this symmetry assumption is relaxed.

- but not uncontroversially⁶⁸ - there are non-resentment cases, perhaps some lottery cases for instance, that seem to exhibit the same kind of oddity, where a similar explanation does not seem plausible. And even in the cases where there are such problems, they don't seem to exhaust the problem with the relevant statistical resentment.

Still, this story, emphasizing how it is that the believer got to have all and only the evidence that they have, has a considerable advantage over impurist and knowledge-based accounts (and to an extent, also over the other conservative solutions discussed earlier in this section): this explanation stands out because it doesn't criticize the credence level or the belief itself, at least not directly. Rather, it criticizes the believing agent, for violating norms of the ethics of evidence gathering⁶⁹. And because these are practical norms, they are pragmatically sensitive to many of the points raised throughout (like the differential price of false beliefs, perhaps even considerations of incentives, and so on). So even if, as we've argued, relative to a body of evidence the rationally permissible credence is fixed, and even if raising one's credence that one's colleague badmouthed one is just as problematic as believing that they did, we still have here a partial explanation of what's (often) wrong with statistical resentment.

5. Partial Solution or an Impossibility Result: In Lieu of a Conclusion

The prospects for a fully satisfactory solution to the puzzle of statistical resentment seem bleak⁷⁰. In particular, the emphasis on the relevance of credences resulted in what may be

⁷⁰ Here's an updated table of the new lessons – since the previous table – learned from cases, both new ones and ones that were already introduced in section 1:

Name	Additional lessons			
GATECRASHERS	Counts against a knowledge-based account (because of knowledge-fetishism).			
BLUE BUS COM- PANY	Counts against a knowledge-based account (because of knowledge-fetishism).			
PHYSICIAN	Action based on statistical evidence clearly justified – so either such evidence suffices for knowledge, or knowledge is not needed. Either way, a problem for a knowledge-based account.			
Truth to Power	Low stakes; so encroachment is not the solution.			
MATH AND CAR	Hard to distinguish on purist knowledge-based accounts.			
PRE-RESENTMENT	Counter-example to knowledge-based accounts.			
SURPRISE	Knowledge-based accounts can't distinguish between SURPRISE LOTTERY and SURPRISE RESENTMENT.			
DNA	A challenge to knowledge-based accounts.			
Urn	A counterexample to credence-impurism.			
ASHKENAZY TIP- PER: CREDENCE	The moral objection applies just as powerfully to the credence case; given the implausibility of credence impurism, this shows impurism will not solve our problem.			
PROBABLY ASHKE- NAZY TIPPER	Still objectionable. Again, against impurism.			

⁶⁸ One of us is more convinced about this than the other.

⁶⁹ For a closely related point (in the context of criticizing Stroud's discussion of epistemic partiality) see Enoch (2016, 31-3).

thought of as almost an impossibility proof: Relative to the evidence, the credence value is fixed. So given some (incriminating) statistical evidence, we have to accept either the raising of credence in blame (which seems just as objectionable as statistical resentment in the case of beliefs), or an irrationally low credence. Impurist and knowledge-based solutions cannot, we think, find a way out. The conservative considerations from the previous section do better, but fall short of offering a fully general, fully satisfying solution.

So we think it's important to reduce expectations about uniformity and generality, and to be willing to bite some bullets.

Regarding uniformity, we stand by our starting point - that, say, there seems to be something very similar going on in LOTTERY, BLUE BUS COMPANY, THE CENTRAL RESENTMENT CASE, TRUTH TO POWER, and a few other of our cases. In other words, other things being equal, a theory that offers a unified story about all of these is better for so doing than a theory that doesn't. But other things – we've come to see – are not equal. If no unified account can plausibly work, we should take that as strong evidence that the landscape is not quite as unified as we may have thought. And so we should be willing to accept partial theories - theories that get some of the cases in ways that seem importantly right, but whose scope should be limited so as to allow other theories to shoulder their part of the burden in other cases. This doesn't mean that nothing about uniformity will be at all important. We can still hold on to especially strong uniformity intuitions - perhaps, for instance, we should be more confident that the oddity of statistical resentment and statistical pride get a similar explanation than we should be that statistical resentment and legal cases regarding statistical evidence should be accounted for by a similar explanation⁷¹. But this is the kind of discussion we should now be having, following the skeptical results so far. It's the discussion about how much and what kind of uniformity we want to maintain, and how high an intuitive price we're willing to pay for it in other desiderata.

And of course, it's not impossible that there will be some remaining cases of statistical resentment with no good explanation (unified or otherwise) of their inappropriateness. And for those, perhaps we should be willing to bite the bullet, and declare them (surprisingly) acceptable after all. Perhaps – we're really not sure – some DNA cases are of this kind⁷². Bullet-biting

CREDIT CARD PIN	Some beliefs based on purely statistical evidence seem justified and may amount to knowledge. A problem for the knowledge-based account.
JOINT APPOINT- MENT	Resentment is problematic (in some cases) even when based on statistical evidence together with other evidence.
STRANGER	Statistical resentment is problematic even in the absence of rich information, so the problem is not just that of screening off.

⁷¹ Notice that once the full generality and uniformity ambition is deserted, the story we initially told about the legal case of statistical resentment may be vindicated after all. And indeed, perhaps then incentive stories can do work elsewhere as well, not as a general explanation, but as a part of a messy one.

Whether a similarly conciliatory spirit can save something from other explanations we rejected along the way – say, those in terms of moral encroachment, or in terms of the knowledge account – will depend on the fate of the more principled arguments we used against them.

⁷² Let us again stress that an account that combines a role for Sensitivity and a role for incentives does especially well – better than all alternative accounts, as far as we can tell – in explaining and vindicating intuitions about DNA cases.

- despite its impressive pedigree - is never happy. It would have been better had we been able to avoid it. But if the arguments in this paper work, there are powerful reasons – having to do with the relation between the moral and the epistemic, between beliefs and credences, and indeed with intuitive judgments about some other, more central cases – to settle for explanations (like the ones in the previous section) that cannot accommodate all of our intuitive judgments in this context. Some gentle bullet-biting, then, seems unavoidable.

Such a complicated, not-fully-unified, somewhat messy story of the inappropriateness of statistical resentment may not be what we've hoped for when we embarked on the project. But it may be – and for now, we tentatively conclude that it is – the best that can here be achieved.

Bibliography

- Anderson, Charity & Hawthorne, John (2018a), "Pragmatic Encroachment and Closure", in Pragmatic Encroachment in Epistemology, Brian Kim & Matthew McGrath (eds.), (Routledge).
 - (2018b), "Knowledge, Practical Adequacy, and Stakes," Oxford Studies in Epistemology v.5.
- Basu, Rima (2019a) "What We Epistemically Owe to Each Other", *Philosophical Studies* 176(4): 915-931.
 - (2019b), "The Wrongs of Racist Beliefs", Philosophical Studies 176(9): 2497-2515.
- Basu, Rima and Schroeder, Mark (2019), "Doxastic Wronging", in Pragmatic Encroachment in *Epistemology*, Brian Kim & Matthew McGrath (eds.), (Routledge: 181-205).
- Bird, A. (2007), "Justified judging" *Philosophy and Phenomenological Research* 74: 81–110.
- Blome-Tillmann, Michael. (2017) "More Likely Than Not' Knowledge First and the Role of Bare Statistical Evidence in Courts of Law." In Knowledge First. Carter, Gordon, and Jarvis (eds.), (Oxford University Press: 278–292).
- Blome-Tillmann, Michael (2015), "Sensitivity, Causality, and Statistical Evidence in Courts of Law", Thought: A Journal of Philosophy 4 (2):102-112.
- Renée Jorgensen Bolinger (2020), "The rational impermissibility of accepting (some) racial generalizations", Synthese, 197, 2415-2431.
 - (forthcoming), "Explaining the Justificatory Asymmetry between Statistical & Individualized Evidence", forthcoming in Truth and Trials: Dilemmas at the Intersection of Epistemology and Philosophy of Law, Edited by Zachary Hoskins and Jon Robson (University of Nottingham).
- Brown, Jessica (2013), "Impurism, Practical Reasoning, and the Threshold Problem", Noûs 47 (1):179-192.

- Buchak, Lara (2014), "Belief, Credence, and Norms", Philosophical Studies 169: 285-311.
- Cohen, Jonathan L. (1977), The Probable and the Provable (Oxford: Clarendon Press).
- DeRose, K. (1992), "Contextualism and Knowledge Attributions", *Philosophy and Phenomenological Research* 52: 913–29.
- Di Bello, Marcello and O'Neill, Collin (2019), "Profile Evidence, Fairness and the Risks of Mistaken Verdicts", *Ethics* 130, 147-178.
- Dorst, K. (2019), "Lockeans Maximize Expected Accuracy," Mind 128(509):175-211.
- Enoch, David (2018), "In Defense of Procedural Rights (or Anyway, Procedural Duties): A Response to Wellman", *Legal Theory* 24(1): 40-49.
 - (2017), "Political Philosophy and Epistemology: The Case of Public Reason", Oxford Studies in Political Philosophy 3: 132-165.
 - (2016), "What's Wrong with Paternalism? Autonomy, Belief, and Action", *Proceedings of the Aristotelian Society* 116: 21-48.
- Enoch, David and Fisher, Talia (2015), "Sense and Sensitivity: Epistemic and Instrumental Approaches to Statistical Evidence", *The Stanford Law Review* 67: 557-611
- David Enoch, Levi Spectre, and Talia Fisher (2012), "Statistical Evidence, Sensitivity, and the Legal Value of Knowledge", *Philosophy and Public Affairs* 40, 197-224.
- Enoch, David and Spectre, Levi (2019), "Sensitivity, Safety, and the Law: A Reply to Pardo", *Legal Theory* 25, 178-199.
 - (Ms.), "There Is No Such Thing as Doxastic Wrongdoing".
- Fantl, Jeremy & McGrath, Matthew (2009), Knowledge in an Uncertain World, Oxford: OUP.
 - (2007), "On Pragmatic Encroachment in Epistemology", *Philosophy and Phenomenological Research* 75:3, 558–89.
 - (2002), "Evidence, Pragmatics, and Justification", *The Philosophical Review* 111:1, 67-94.
- Fritz, James (2017), "Pragmatic Encroachment and Moral Encroachment", *Pacific Philosophical Quarterly* 98, 643-661.
 - (2020), "Moral Encroachment and Reasons for the Wrong Kind", *Philosophical Studies* 177, 3051-3070.
- Fritz, James and Jackson, Elizabeth (2020), "Belief, Credence, and Moral Encroachment", forthcoming in *Synthese*.
- Gao, Jie (2019), "Credal Pragmatism", Philosophical Studies 176: 1595-1617.
- Gardiner, Georgi (Forthcoming), "Legal Evidence and Knowledge", in *The Routledge Hand-book of the Philosophy of Evidence*, eds. Maria Lasonen-Aarnio and Clayton Littlejohn. Routledge. Available at https://drive.google.com/file/d/1cXwUHqLrMBgVL_2UJjbk-

DefxE5iizjL/view

- (2018a), "Legal Burdens of Proof and Statistical Evidence", in David Coady and James Chase (eds.) The Routledge Handbook of Applied Epistemology, 179-195.
- (2018b) "Evidentialism and Moral Encroachment", in Believing in Accordance with the Evidence, McCain Kevin (ed.), (Springer Verlag: 169-195).
- Hawthorne, John (2003), Knowledge and lotteries, (Oxford: Oxford University Press).
- John Hawthorne, Daniel Rothschild, and Levi Spectre (2016), "Belief Is Weak", Philosophical Studies 173, 1393-1404.
- Holguín, B. (Ms.), "Thinking, guessing, and believing", (2020), NYU.
- Elisabeth Jackson (2019), "How Belief-Credence Dualism Explains Away Pragmatic Encroachment", The Philosophical Quarterly 69, 511-533.
- Jacobson, Daniel (2011), "Fitting Attitude theories of Value", The Stanford Encyclopedia of Philosophy, available here: https://plato.stanford.edu/entries/fitting-attitude-theories/
- Joyce, James (2005), "How Degrees of Belief Reflect Evidence", Philosophical Perspectives 19(1):153-179.
- Karlander, Karl and Spectre, Levi (2010), "Sleeping Beauty Meets Monday", Synthese 174: 397-412.
- Kyburg, H. E. (1961), Probability and the logic of rational belief, Middletown: Wesleyan University Press.
- Leitgeb, Hannes (2014), "The Stability Theory of Belief. Philosophical Review 123(2)131-171.
- Lewis, D. (1996), "Elusive Knowledge", Australasian Journal of Philosophy 74: 549–67.
- Littlejohn, Clayton. (2017), "Truth, Knowledge, and the Standard of Proof in Criminal Law", Synthese: awaiting an issue.
- Marušić, Berislav (2015), Evidence and Agency: Norms of Belief for Promising and Resolving (Oxford: Oxford University Press).
- McCain, K., and Poston, T. (2014), "Why explanatoriness is evidentially irrelevant", Thought 3: 145-153.
- Marušić, Berislav and White, Stephen (2018), "How can Beliefs Wrong? A Strawsonian Epistemology", Philosophical Topics 46 (1): 97-114.
- Moss, Sarah (2018a), Probabilistic Knowledge, (Oxford: Oxford University Press).
 - (2018b), "Moral Encroachment", Proceedings of the Aristotelian Society 118 (2):177-205.
 - (2019) "Full Belief and Loose Speech", Philosophical Review 128 (3):255-291.
- Nozick, Robert (1981), Philosophical Explanations, (Cambridge, Mass.: Harvard University

Press).

- Pace, M. (2011), "The epistemic value of moral considerations: Justification, moral encroachment, and James' 'Will To Believe'", *Noûs*, 45 (2): 239–268.
- Pardo, Michael S. (2018), "Safety vs. sensitivity: Possible worlds and the law of evidence", *Legal Theory* 24 (1): 50-75.
- Parsad, Vinyak K. & Cifu, Adam S. (2015), *Ending Medical Reversal improving outcomes, saving lives*. Baltimore: John Hopkins University Press.
- Posner, Richard A. (1999), "An Economic Approach to the Law of Evidence", *Stanford Law Review* 51: 1477.
- Pundik, Amit (2008), "Statistical Evidence and Individual Litigants: A Reconsideration of Wasserman's Argument from Autonomy", *International Journal of Evidence and Proof* 12: 303–24.
 - (2020), "Predictive Evidence and Unpredictable Freedom", Oxford Journal of Legal Studies 40, 238-264.
- Reed, Baron (2010), "A defense of stable invariantism", Noûs 44 (2): 224-244.
- Reiss, Julian and Ankeny, Rachel A. (2016), "Philosophy of Medicine", *The Stanford Encyclopedia of Philosophy*, https://plato.stanford.edu/entries/medicine/#RandContTriaEvidBase-Medi
- Ross, Jacob & Schroeder, Mark (2014), "Belief, Credence, and Pragmatic Encroachment", *Philosophy and Phenomenological Research* 88 (2): 259-288.
- Rothschild, Daniel (2019), "What it takes to believe", Philosophical Studies 177, 1345-1362.
- Schroeder, Mark (2019), "Persons as Things", in Oxford Studies in Normative Ethics 9, 95-115.
 - (2018a), "When Beliefs Wrong", Philosophical Topics 41 (6): 115-127.
 - (2018b), "Rational stability under pragmatic encroachment", *Episteme* 15 (3): 297-312.
- Smilansky, Saul (1994), "The Time to Punish", *Analysis* 54 (1): 50 53.
- Smith, Martin (2017), "When Does Evidence Suffice for Conviction?", *Mind* 127 (508): 1193-1218.
 - (2016), Between Probability and Certainty: What Justifies Belief, (Oxford University Press).
- Stanley, Jason (2005), Knowledge and Practical Interests, (Oxford University Press).
- Stroud, Sarah (2006) "Epistemic partiality in friendship", Ethics 116 (3): 498-524.
- Thomson, J. J. (1986), "Liability and Individualized Evidence", in *Rights, Restitution, and Risk*, William Parent (ed.) (Cambridge: Harvard University Press: 225–250).

- Vargas, Manuel (2013), Building Better Beings: A Theory of Moral Responsibility (Oxford: Oxford University Press).
- Weatherson, Brian (2005), "Can We Do Without Pragmatic Encroachment?", Philosophical *Perspectives* 19: 417–443.
- Williamson, Timothy (forthcoming-a), "Knowledge, credence, and the strength of belief", in Towards an Expansive Epistemology: Norms, Action, and the Social Sphere, A. Flowerree and B. Reed. (eds.) (London: Routledge). Available at http://media.philosophy.ox.ac.uk/docs/people/williamson/Strengthofbelief.pdf
 - (forthcoming-b), "Justifications, excuses, and sceptical scenarios", in The New Evil Demon, Dutant, J., & Dohrn, D. (eds.), (Oxford: Oxford University Press). Available at https://www.academia.edu/10023313/Justifications Excuses and Sceptical Scenarios
 - (2017), "Acting on knowledge", in Knowledge First Approaches to Epistemology and Mind, Carter, A.; Gordon, E. & Jarvis, B. (eds.), (Oxford: Oxford University Press: 163-183).
 - (2013), "Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier cases in epistemic logic, Inquiry 56(1): 15–29.
 - (2011), "Knowledge First Epistemology," in The Routledge Companion to Epistemology, Sven Bernecker & Duncan Pritchard (eds.), (Routledge: 208-218).
 - (2000), Knowledge and its Limits, (Oxford: Oxford University Press).