Philos Stud
https://doi.org/10.1007/s11098-020-01563-3

# What it means to respect individuality

**Xiaofei Liu**[1] · **Ye Liang**[2]

Accepted: 22 September 2020


**Abstract** Using pure statistical evidence about a group to judge a particular member of that group is often found objectionable. One natural explanation of why this is objectionable appeals to the moral notion of respecting individuality: to properly respect individuality, we need individualized evidence, not pure statistical evidence. However, this explanation has been criticized on the ground that there is no fundamental difference between the so-called "individualized evidence" and "pure statistical evidence". This paper defends the respecting-individuality explanation by developing an account of what it means to respect individuality. It combines an idealistic account of respecting individuality and a prioritization account of respecting individuality, and offers a principled way to distinguish individualized evidence from non-individualized evidence.

**Keywords** Individuality · Statistical discrimination · Evidence

## 1 Introduction

Austin, a 40-year-old white businessman, is waiting for taxi outside of a high-end shopping mall in Shanghai. A flower peddler approaches him and asks: "Sir, would you like to buy a flower for this beautiful lady?" pointing to a young Chinese woman standing next to him, whom Austin does not know. The peddler takes them

✉ Xiaofei Liu
liuxiaof@gmail.com

Ye Liang
ye.liang@okstate.edu

[1] Department of Philosophy, Xiamen University, 422 Siming South Road, Xiamen 361005, China

[2] Department of Statistics, Oklahoma State University, 301 MSCS Bldg, Stillwater, OK 74078-1056, USA

Published online: 06 October 2020

to be a couple. Out of curiosity, Austin asks the peddler why he came to this belief. Imagine three different scenarios:

> *Peddler 1*: The peddler responds that he believes that most middle-aged white men are attracted to young Chinese women.
>
> *Peddler 2*: The peddler responds that his company has actually conducted systematic research in the past 10 years, according to which a middle-aged white man and a young Chinese woman appearing next to each other in this shopping district have a high chance to be in a relationship.
>
> *Peddler 3*: The peddler responds that he overheard Austin talking to the woman next to him in Chinese, using phrases like "sweetie", while in fact, Austin was talking to his Chinese adopted daughter on the phone through a Bluetooth headset.[1]

Austin would probably be outraged by the reason provided by Peddler 1, which, like typical racist beliefs, is nothing more than an ill-grounded prejudice. However, he would probably also find the reason provided by Peddler 2 problematic (if not outright offensive), even if he grants the veracity of its statistical basis. By contrast, he would probably find the reason provided by Peddler 3 acceptable, despite feeling awkward.

What would explain the difference between Peddler 2's and Peddler 3's reason, both grounded upon factual evidence? One natural reaction is to say that Peddler 2's reason consists of pure statistical facts about a group (often called *pure statistical evidence*) and therefore disrespects Austin's individuality, whereas Peddler 3's reason consists of evidence concerning Austin's particular characteristics (often called *individualized evidence*) and therefore respects his individuality. Since respecting individuality is an important moral principle, there is a morally significant difference between these two reasons (Miller 1999; Eidelson 2015; Liu 2015). This line of thought seems to be behind several US court rulings too, in which the courts claimed that "mathematic probability" is not the type of "direct" evidence that could be used to justify a claim that an individual actually did something,[2] and that the standard of "reasonable suspicion requires *particularized* suspicion", which cannot be based on such categorical membership as ethnicity or nationality.[3]

Despite its intuitiveness and popularity, trying to distinguish the belief-forming method in *Peddler 2* and that in *Peddler 3* by appealing to respecting individuality faces at least two serious challenges. One challenge is that Peddler 2's evidence, despite being labeled "pure statistical", does include Austin's individual characteristics—his age, his race, and his standing next to a young Chinese woman. Where one stands, for instance, seems a piece of evidence as "individualized" and

[1] For the sake of argument, let's assume that the situation is such that any bystander in the peddler's position would easily perceive Austin as talking to the woman standing next to him and that this is apparent to Austin upon reflection.

[2] *Smith v. Rapid Transit*, 317 Mass. 469, 470, 58 N.E.2d 754, 755 (1945).

[3] *United States v. Montero-Camargo*, 208 F.3d 1122, 1134 (9th Cir. 2000), emphasis in original.

"particularized" as what one says. The only sense in which the latter might be said to be more "direct" than the former is that the degree to which one's saying certain things to another person makes it likely that the two individuals are in a relationship is higher than the degree to which standing next to each other makes it possible that the two individuals are in a relationship. Even this is but an appearance, for it is not hard to imagine cases in which the so-called "pure statistical evidence" confirms a judgment to a greater degree than "individualized evidence". Therefore, whether individuality is respected seems to be the wrong place to look for a distinction—"pure statistical evidence" can be as "individualized" as "individualized evidence".

The second challenge is that "individualized evidence" can be as "statistical" as "pure statistical evidence". For example, the reason why one's saying "sweetie" to another person confirms the belief that the two individuals are in a relationship is presumably the following: *usually*, only people in a relationship use such words to address each other. Thus, both Peddler 2's reason and Peddler 3's reason involve what Frederick Schauer calls "nonspurious but non-universal generalizations", and "what appears to be an individualized analysis is simply an aggregate of stereotypes" (2003: 69).

These two challenges lead some theorists to doubt whether there is indeed any fundamental difference between the so-called individualized evidence and non-individualized evidence (e.g., Levin 1992; Blum 2002; Schauer 2003; Lippert-Rasmussen 2011; Beeghly 2018). If there is no such difference, then the attempt to distinguish cases like *Peddler 3* from cases like *Peddler 2* by appealing to the idea of respecting individuality is undermined. Moreover, if no such difference exists, the very idea of respecting individuality may seem to make no sense. As Michael Levin once puts it: "[t]here is in fact no such principle [to respect individuality]… [p]eople are and must always be judged by the classes to which they belong" (1992: 23).

The purpose of this paper is to investigate what exactly it means to respect individuality in the context of forming a belief about somebody, and the plausibility of distinguishing cases like *Peddler 3* from cases like *Peddler 2* by appealing to this idea. Before starting the investigation, it is helpful to situate this project in the current literature on statistical discrimination and statistically grounded beliefs.

There have been two main approaches to distinguish a belief-forming practice based on pure statistical evidence (e.g., racial profiling) from a belief-forming practice based on seemingly individualized evidence. One approach tries to locate the difference in some purely epistemic error or flaw in the former: for example, the former is resistant to counter-evidence (Fricker 2007), or it fails to satisfy epistemic principles such as sensitivity (Enoch et al. 2012; Enoch and Fisher 2015) and safety (Blome-Tillmann 2017; Pardo 2018). The other approach is motivated by the "moral encroachment" theory, according to which whether it is epistemically justified to believe something (or to be confident of it to a certain degree) does not just depend on how much evidence one has, but also on what moral features that belief has (e.g., Stroud 2006; Keller 2007; Pace 2011; Moss 2018b). Based on this theory, some theorists recently propose that a belief-forming practice like racial profiling often involves some morally undesirable features, such as stigmatizing

certain social groups or contributing to the disadvantage of socially disadvantaged groups, and thus the threshold of it being epistemically justified should be higher—that is, it requires more than pure statistical evidence for them to be epistemically justified (e.g., Basu 2018; Bolinger 2018; Schroeder 2018; Hellman 2018). This provides another way to explain the difference.

Our project can be viewed, in a broad sense, as falling into the second approach: like the moral encroachment approach, we believe that moral considerations (in this case, respecting individuality) impose certain requirements on our belief-forming process. But unlike the moral encroachment approach, we remain neutral about whether such requirements constitute part of the epistemic justification or something of its own kind. Despite this divergence, our project can be helpful to the moral encroachment accounts in at least two regards. First, it provides a concrete case of how a particular kind of moral consideration impacts (or encroaches on) what we should believe. Second, as Sara Moss noted, most discrimination-related discussions in the moral encroachment literature "focus on moral harms that occur as a result of profiling" (2018b: 194). But such contingent moral features may not be present in all purely statistically grounded beliefs. As we can see from the *Peddler* cases, some objectionable beliefs may signify an "advantage", instead of a "disadvantage" or "harm", enjoyed by the relevant social groups. Appealing to a deontological moral duty such as respecting individuality avoids such limitation.

Section 1 starts with an idealistic account of what respecting individuality requires. We argue that this account should be revised to allow for degrees of respecting individuality. We then develop, based on this account, a general framework for measuring the appropriate degree of respecting individuality. Section 2 introduces a prioritization account of respecting individuality—respecting individuality requires us to prioritize individualized evidence. We then propose a principled way to distinguish between individualized and non-individualized evidence. Section 3 discusses how these two accounts can be combined to handle some real-life cases. Section 4 compares this combined account to some prominent alternative accounts.

## 2 An idealistic account

A recent account of what it is to respect individuality by Kasper Lippert-Rasmussen offers a possible solution to the two challenges mentioned earlier. Lippert-Rasmussen proposes that *X* treats *Y* as an individual if and only if this treatment is "informed by all relevant information… reasonably available to *X*" (2011: 54). On this account, when we form a belief about an individual *Y*—e.g., attributing a certain property *P* to *Y*—based on some identifying properties, whether or not this belief-forming practice respects *Y*'s individuality depends not on whether this or that particular identifying property is considered, but on whether *all* the properties relevant to the attribution of *P* to *Y* are properly considered. Thus, there is no need to call this type of evidence "individualized" and another type "non-individualized": they should all be given proper consideration as long as they are relevant to the attribution of *P* and are reasonably available in the given context.

Even though it avoids the challenges by moving away from a distinction between individualized and non-individualized evidence, this account does not offer any insight on how cases like *Peddler 3* can be distinguished from cases like *Peddler 2*, since both cases can fail to consider all the relevant information. Moreover, respecting individuality, on this account, is a yes-or-no issue.[4] But it is quite implausible to think that there is a clear line such that we respect one's individuality if we reach this line, and fail to respect individuality if we miss even one piece of information, no matter how insignificant it may be. Taking into consideration all the relevant and reasonably available information seems to be an ideal for, rather than a necessary and sufficient condition of, respecting individuality. A more plausible view is that respecting individuality is a matter of degree, something that can be more or less achieved. Maybe what we should say about the *Peddler* cases is that the belief-forming practice in *Peddler 3* respects individuality to an appropriate degree but the belief-forming practice in *Peddler 2* fails to do so. The idealistic account needs to be revised to allow for degrees of respecting individuality.

To determine whether a belief-forming practice respects individuality to an appropriate degree, we first need an account of how to measure the degree to which such a practice respects or disrespects individuality. We believe this measurement can be grounded upon the two factors emphasized in Lippert-Rasmussen's account—i.e. whether a piece of information is "relevant" to the attribution of *P* to *Y* and whether that piece of information is "reasonably available" to *X*. The rest of this section sketches a general framework for measuring the degree to which a belief-forming practice respects individuality and for determining which degree is appropriate in a given context.

Let's start with the measurement issue. Consider first relevance. How relevant some considered or neglected evidence is to the attribution of *P* to *Y* constitutes an important measure of the degree to which the belief-forming practice respects *Y*'s individuality. Let's define *relevance* as the degree to which a piece of evidence, judged by a rational person, confirms or disconfirms the attribution of *P* to *Y* in the given context. The most relevant evidence to the attribution will be evidence that confirms or disconfirms the attribution to the greatest degree.[5] A piece of evidence is irrelevant to the attribution if it neither confirms nor disconfirms the attribution in the given context.

For presentational reasons, we separate two different measures: the degree to which a belief-forming practice *respects Y*'s individuality is understood as how much relevant evidence *has been considered* in that practice, and the degree to which a belief-forming practice *disrespects Y*'s individuality is understood as how much relevant evidence *has been neglected* in that practice (italicized "respect" and "disrespect" are used to signify this limited sense hereafter). These two measures,

[4] Some recent epistemically based accounts for distinguishing between individualized and pure statistical evidence, such as Enoch et al. (2012), seem to hold this "yes-or-no" view too.

[5] Whether a piece of evidence *e* confirms or disconfirms a belief *B* is understood in the Bayesian way: *e* confirms *B* if the probability of *B* being true given *e* is greater than that without *e*; *e* disconfirms *B* if the probability of *B* being true given *e* is smaller than that without *e*; *e* neither confirms nor disconfirms *B* if the probability of *B* being true given *e* is equal to that without *e*.

together, provide a fuller picture of how individuality is respected in a belief-forming practice.

Given the earlier definition of relevance, these two measures can be specified as follows: the degree to which a belief-forming practice *F respects Y*'s individuality is directly proportional to the total number of relevant evidence that has been considered in *F* and the degree to which each considered evidence confirms the attribution of *P* to *Y*; the degree to which *F disrespects Y*'s individuality is directly proportional to the total number of relevant evidence that is neglected in *F* and the degree to which each neglected evidence would disconfirm the attribution. Thus, we have

(1) The degree to which *F respects Y*'s individuality is directly proportional to $\sum_e \mathrm{O}(e)$, and
(2) The degree to which *F disrespects Y*'s individuality is directly proportional to $\sum_{e'} \mathrm{NO}(e')$.
($e$ = any relevant evidence considered in $F$, $e'$ = any relevant evidence neglected in $F$, $\mathrm{O}(e)$ = the degree to which $e$ confirms the attribution of $P$ to $Y$, $\mathrm{NO}(e)$ = the degree to which $e$ disconfirms it.)

Next, consider reasonable availability. Whether a piece of evidence is reasonably available to *X* depends on two factors: how costly it is to obtain that evidence, and how significant an impact the belief formed in *F* will have on *Y*.

The more costly it is to obtain *e*, the less reasonably available *e* is to *X*. When *e* is simply unavailable, we may think of its cost as being infinitely big. This cost should not be understood as just to include the resources required from *X* to obtain *e* (such as money and time), but also the risks coming with obtaining it that *X* is morally required to weigh in her deliberation. To see this, consider the following case.

> *Vaccine*: A deadly virus is fast spreading in a certain area, against which there is a vaccine. But this vaccine will cause severe discomfort to a small portion of the population who are naturally immune to the virus. It involves only a simple test to find out who is immune, but by the time this test can be done, a great many deaths will have been caused by the virus.

Even if the information of whether any particular individual is immune to the virus is technically easily obtainable, seeking out this information will cause a significant number of deaths, which is a cost that we are morally required to weigh in our deliberation. Thus, there is a good reason to think that the relevant information is not reasonably available.

The impact that the formed belief has on *Y* can also affect whether some neglected evidence is reasonably available to *X*. Consider a revised vaccine case: instead of causing severe discomfort to people who are naturally immune to the virus, the vaccine will kill them. There is now a much stronger reason for obtaining information about individual immunity before deciding whom to vaccinate, even though the cost of obtaining the information remains the same. This case shows that

the degree to which some not-yet-obtained evidence is reasonably available is inversely proportional to the impact the formed belief will have on *Y*.

Thus, we propose the following measure for the degree of reasonable availability:

(3) The degree to which any relevant but neglected evidence *e'* is reasonably available to *X* is directly proportional to I(*y*)—the significance of the impact that the formed belief has on *Y*, and inversely proportional to C(*e'*)—the cost on *X* to obtain *e'*.

Let's now turn to the appropriateness issue—how to determine whether a belief-forming practice respects individuality to an appropriate degree in a given context. Finding a precise way to determine the appropriate degree is difficult and probably controversial. Nevertheless, we find the following two principles to be a good starting point. First, if the degree to which a belief-forming practice *F respects Y*'s individuality is greater than the degree to which it *disrespects Y*'s individuality, then it seems reasonable to believe that *F* has respected *Y*'s individuality, all things considered, to an appropriate degree. In other words, if the evidence that has been considered in *F* confirms the attribution of *P* to *Y* to a greater degree than the evidence that is neglected in *F* disconfirms this belief, there is a good reason to think that *F* properly respects *Y*'s individuality. Second, other things being equal, the less reasonably available the evidence considered in *F* is, the more ready we are to accept the claim that *F* respects individuality; the more reasonably available the evidence neglected in *F* is, the more ready we are to accept the claim that *F* disrespects individuality.

Based on these two principles, we propose the following measure to determine the appropriate degree of respecting individuality:

*Measure of appropriateness*: a belief-forming practice *F* respects *Y*'s individuality to an appropriate degree if $\sum_e(\mathrm{O}(e)\cdot\mathrm{C}(e)) > \sum_{e'}(\mathrm{NO}(e')\cdot\mathrm{I}(y)/\mathrm{C}(e'))$ (where $\sum_e(\mathrm{O}(e)\cdot\mathrm{C}(e))$ means a summation of the adjusted degrees to which each relevant evidence that has been considered confirms the attribution of *P* to *Y*—each adjusted by the cost of obtaining that evidence; $\sum_{e'}(\mathrm{NO}(e')\cdot\mathrm{I}(y)/\mathrm{C}(e'))$ means a summation of the adjusted degrees to which each relevant evidence that is neglected disconfirms the attribution—each adjusted both by the significance of the impact that the formed belief has on *Y* and, inversely, by the cost of obtaining that evidence)[6]; *F* fails to respect *Y*'s individuality to an appropriate degree if otherwise.

This revised idealistic account pictures a general framework for determining whether a belief-forming practice respects individuality to an appropriate degree in a given context. It answers Levin's worry—in the absence of a principled distinction between individualized and non-individualized evidence, how a belief-forming

[6] O(*e*) is not modified by I(*y*) because only when the information is *not yet considered*, the significance of the impact that the formed belief has on *Y* increases the importance to obtain that information and thus the threshold of unavailability.

practice can be said to disrespect individuality. Nonetheless, it is still unable to distinguish *Peddler 2* from *Peddler 3*, as both may fail to reach the appropriate degree of respecting individuality. Thus, the following worry remains: respecting individuality seems to be the wrong place to look for a distinction. Let's now turn to a different account that helps to answer this worry.

## 3 A prioritization account

Respecting individuality often requires us to give priority to a particular type of evidence. Consider the following two cases.

> *Disease*: It is a statistically established fact that, in a certain population, 90% of the Buddhists have a certain disease, while only 45% of the whole population have that disease. John is a Buddhist, and his recent medical examination, which is 90% reliable, says that he does not have the disease.
> *Recruitment*: The junior national basketball team is searching for shooters. Lin applied and shot consistently better than other applicants during the tryout, and tryout performance is 90% reliable for evaluating shooting skills. Meanwhile, Lin is from School A. Records show that 90% of the applicants from A were not qualified shooters.

In *Disease*, we face two contradicting pieces of information: John's being a Buddhist and John's medical examination. One piece confirms the belief that John has the disease, and the other piece disconfirms it, both to the same degree (i.e., 0.9). Despite this equivalence, when it comes to determining whether John should receive medical treatment, we usually rely more on the result of his medical examination, not his being a Buddhist. Likewise, when coaches of the junior national basketball team decide whether to recruit Lin, we believe that they should rely more on his tryout performance, not which school he is from, even if these two pieces of information are equivalent in terms of the degree to which each confirms or disconfirms the belief that Lin is a good shooter.[7]

The above two cases show that when the degree to which each confirms or disconfirms the relevant belief is similar, individualized evidence should nevertheless take precedence over pure statistical evidence. If one judges John or Lin mainly on the basis of the pure statistical evidence, then they have a legitimate reason to complain that their individuality is disrespected. Therefore, respecting individuality also requires that, other things being equal, individualized evidence be prioritized over non-individualized evidence.

Notice that this conclusion is compatible with the Bayesian theory of probability. Take *Disease* as an example. A Bayesian typically takes 0.9 as the *prior probability* for the Buddhist population. John would be considered as a sample point from the population; without any medical examination, the probability that John has the disease is 0.9 (which is typically considered as a "no data" situation). After the

[7] The problem underlying these two cases is often referred to as the "reference class problem".

medical examination, we now have a *data point*, which is that John has no disease, with an accuracy of 0.9. Combining the information from the data and the prior probability leads to the Bayesian *posterior probability*, upon which we make a decision. Nowadays, most Bayesians tend to adopt the objective Bayes theory, which aims to minimize the prior impact and weigh the data point much higher than the population prior. They do so by assigning a so-called *non-informative prior distribution*, which implies a higher degree of uncertainty about the prior information (compared to the data point), and thus reduces the impact of the prior on the posterior probability (e.g., Robert 2007: 127).[8]

But one worry immediately arises: is there a principled way to draw a distinction between the so-called individualized evidence and non-individualized evidence? The two challenges discussed at the beginning suggest *no*. To make the prioritization account plausible, we first have to find a principled way to draw the distinction.

To begin with, let's take a closer look at the epistemic process through which we attribute a property $P$ to an individual $Y$ based on an identifying property $Q$. This process involves an inference of the following form:

*Premise 1*: $Y$ has $Q$.

*Premise 2*: The probability of having $P$ given $Q$ is higher than normal—that is, higher than not having $Q$.

*Conclusion*: Therefore, there is a higher than normal probability that $Y$ also has $P$.

To truly respect $Y$'s individuality, our justification for this inference must be based on something that can be reasonably believed to be true of $Y$.[9] Obviously, we must have good reason for believing that $Y$ indeed has $Q$. But this cannot be all that is required to satisfy the duty to respect individuality: otherwise, all forms of statistically grounded beliefs would satisfy this duty, as they all inevitably require some identifying property to start with. Therefore, the key for satisfying the duty to respect individuality must lie with *Premise 2*—our justification for *Premise 2* must be based on something that can be reasonably believed to be true of $Y$.

But what exactly does this mean? Consider *Disease* again. When we draw the conclusion that John likely has the disease based on the pure statistical fact about the whole population, our justification for *Premise 2*—the probability of having the disease given John is Buddhist is higher than normal—is not based on something that can be reasonably believed to be true of the individual John: there is no

[8] This is, of course, not suggesting that the prior information (or *base rate*, as it is often called) does not count at all. *Informative prior distributions* (based on the base rate) are used to lower the posterior variation if the data are considered inaccurate or highly uncertain; non-informative prior distributions are used when the data are highly certain. In other words, high-quality individualized evidence is prioritized in Bayesian practices. In *Disease*, since the accuracy of the medical test is relatively high, the prior probability is weighed lightly; were the accuracy of the medical test sufficiently low, the prior probability would be weighed more heavily. Di Bello (2019) gives a detailed discussion of the application of Bayesian theory in statistically grounded beliefs..

[9] Respecting individuality is best understood as a subjective moral duty, and thus it does not require our justification to be grounded on something that is objectively true of $Y$.

propensity or mechanism in John that inclines him to have that disease. What grounds the statistical regularity is some external facts about the Buddhist group and the whole population.[10] By contrast, when we draw the conclusion that John likely does not have the disease based on the result of his medical examination, our justification for the statistical regularity in *Premise 2* is based on something that can be reasonably believed to be true of John: we can reasonably believe that there is some mechanism at work in John such that when he has (or has not) certain physiological conditions, he likely does not have that disease.

Therefore, despite that both the evidence based on group statistic and the evidence based on one's medical examination are probabilistic in nature, there is a critical difference between them: what grounds the statistical regularity in the latter can be reasonably believed to be true of the individual, but what grounds the statistical regularity in the former cannot. In this sense, the latter constitutes an instance of individualized evidence, but the former does not. This distinction has nothing to do with whether the identifying property itself is "individualized" or not—that is, it has nothing to do with *Premise 1*; rather, it has to do with why we believe that the identifying property is connected to the attributed property—that is, what grounds the truth of *Premise 2*. We now have an answer to the challenges mentioned at the beginning.

Let us further clarify this distinction by considering three objections. First, one may object that it seems false to say that what grounds the statistical regularity between having a certain disease and being a Buddhist in a certain population cannot be reasonably believed to be true of John. For example, suppose that John indeed has the disease; in this case, something that is true of John (i.e., he has the disease) is part of what grounds the overall prevalence of the disease in the Buddhist population.

In response, we want to emphasize that John's having the disease alone cannot ground the statistical regularity; it is the fact that all the other individuals also have the disease that grounds it—this fact is definitely something external to John. By contrast, when it is John's genetics that grounds our belief that he has the disease, our justification for the statistical regularity is based on the relevant genetic mechanism, which can be reasonably believed to exist in John.

A second objection is this. People often believe that the presence (or absence) of certain physiological conditions indicates the absence of a disease in John not because they have conducted individual examination on John and discovered that the relevant causal mechanism is present in him, but because they have found this causal mechanism in other human beings and assume that it works in John too. The epistemic process usually involves the following two basic steps.

> *Step One*: Through the study of past subjects, we establish a mechanism in which having the identifying property $Q$ is causally connected to having the property $P$.

[10] This echoes Moss's point that "in many situations where you are forming beliefs about a person, you morally should keep in mind the possibility that they might be an exception to statistical generalizations", which Moss calls the "rule of consideration". See Moss (2018a, p. 221).

*Step Two*: We assume that this mechanism is true for new subjects, and draw conclusions about them based on this assumption.

*Step One* is based on information of the subjects being studied; *Step Two* is not based on information of the subjects *being judged*—we do not look into the actual circumstances of these new subjects. Thus, one may object that assuming what is true of other individuals to be true of *Y* without first looking into *Y*'s actual circumstance also disrespects *Y*'s individuality.

But there is a critical difference between *Step Two* and a generalization based on pure statistical evidence. *Step Two* involves an assumption underlying all scientific induction. Consider a typical scientific reasoning process: scientists first establish, through observation of a number of past events, a certain causal model—when a set of specific pre-conditions are met, having *Q* leads to having *P*; scientists then assume that this causal model is true for all events such that when the set of pre-conditions are met, having *Q* will always lead to having *P*. This assumption, like the assumption in *Step Two*, is not a belief in any particular connection between specific properties, but rather a belief in the fundamental orderliness of the physical world—for example, "physical forces work in the same way under the same conditions". It works as the foundation of scientific induction. When a counterexample emerges, that is, when scientists observe a new event in which the set of pre-conditions are met but having *Q* does not lead to having *P*, they do not give up the assumption; instead, they go back to revise the proposed model by, for example, changing the set of pre-conditions. In this sense, this assumption is treated like a "nonspurious and *universal* generalization". By contrast, a generalization based on pure statistical evidence, in which any specific individual with *Q* is regarded as likely having *P* simply because a certain portion of the population with *Q* also have *P*, enjoys no such epistemic status.

Here is a third objection, which targets not *Step Two*, but *Step One*. It is common knowledge that scientific theories rely essentially on an inductive jump and thus can be wrong—this is referred to as the problem of "inductive risk" in the philosophy of science literature.[11] More importantly, in such areas as medical science, when a causal model fails to make the correct diagnosis on a single case, what scientists often do is not to go back and adjust the parameters of the model, but simply to accept that there can be exceptions to that model, period. Since even well-established causal models allow for exceptions, it might be argued that relying on such models is not fundamentally different from relying on pure statistical evidence.[12]

In response, we want to distinguish two types of causal models: deterministic models—when a set of specific pre-conditions are met, having *Q* invariably leads to having *P*, and probabilistic models—when a set of specific pre-conditions are met, having *Q* is connected to having *P* with a probability of *z*. The latter can further be

[11] See, for example, Churchman (1948), Rudner (1953), Hempel (1965), and more recently Douglas (2000).

[12] We would like to thank an anonymous reviewer for this journal for raising this objection.

divided into two categories: one whose statistical regularity is based on some internal propensity that can be reasonably believed to be true of every individual in the intended scope, and one whose statistical regularity is based on frequency, which is a property of a population, not of any individual.

Deterministic models do not allow for exceptions, and thus are free from the aforementioned objection. Even though such a model can still be wrong due to inductive risk, this is a flaw in the practice of model-construction, not in the practice of using individualized evidence. A probabilistic model of the first kind allows for exceptions, but as long as it is scientifically well-established, we have good reason to believe that the propensity specified by the model applies to every individual in the intended scope—it's just that this propensity fails to manifest itself in some cases. With both types of models, we are justified to believe that the causal mechanism specified by the model applies to the current individual, and thus forming a belief about this individual based on that mechanism respects his or her individuality. The same thing cannot be said about frequency-based probabilistic models—relying on such models is equivalent to relying on pure statistical evidence. Thus, the third objection is right that there is no fundamental difference between pure statistical evidence and evidence based on probabilistic models of the second kind. But we should not think that all scientific models are of this kind. A key difference between these different models lies precisely in whether the causal mechanism specified by the model can be reasonably believed to be universally applicable.

To sum up: respecting individuality also requires giving priority to individualized evidence over non-individualized evidence. We proposed an account of individualized evidence, according to which whether some evidence counts as individualized evidence does not depend on the nature of the identifying property (for example, where one stands can be individualized evidence in one context but non-individualized evidence in another), but on our justification for believing that the identifying property is statistically connected to the attributed property. If what grounds this statistical regularity can be reasonably believed to be true of the individual at issue then the evidence is individualized evidence, and if not, it is non-individualized evidence. For example, when this statistical regularity is grounded upon a causal mechanism that can be reasonably believed to be working in or on the individual, then the evidence is individualized evidence.

But causal mechanism need not be the only acceptable type of grounding mechanism; logical and conceptual mechanisms can do the same job. For example, seeing a person in a restaurant uniform at a restaurant, one can reasonably believe that he likely works there. This is true even if you are aware that sometimes people dress in a restaurant uniform just to play pranks on costumers.[13] It is the social meaning of a uniform that justifies your belief: unless you have some specific reason regarding this person in uniform for believing otherwise, the mere knowledge of some exceptions will not undermine your justification for applying this social

[13] What is at issue here is whether a belief is subjectively justified, not whether it constitutes knowledge. Thus, those well-known counterexamples, such as the Fake-Barn cases, should not concern us.

meaning of a uniform to the current case. On the other hand, if you indeed have a specific reason to believe that this particular person in uniform belongs to the category of exception, then you probably already have some individualized evidence about that individual. What this shows is that you are *pro tanto* justified in believing that the social meaning of a uniform applies to every uniform, even if you know there are instances in which a person in uniform does not play the corresponding role. Similarly, you are pro tanto justified in believing that a person enjoys equal moral status as you do simply by virtue of her being a human being, even if you know there are circumstances in which this equality is compromised.[14]

For the same reason, it is the social meaning of addressing someone as "sweetie" that the person being addressed is taken to be in a special relationship with the speaker. It is reasonable for Peddler 3 to believe that this social meaning applies to Austin. By contrast, the statistical regularity in *Peddler 2* is not grounded on any such causal, logical, or conceptual mechanism that can be reasonably believed to apply to Austin. It is in this sense that Peddler 3's reasoning respects Austin's individuality, whereas Peddler 2's does not. Thus, we have a reasonable way to distinguish *Peddler 3* from *Peddler 2*.

## 4 Combining the two accounts

Respecting individuality requires us to prioritize individualized evidence. However, this does not mean any belief-forming practice that fails to consider individualized evidence automatically disrespects individuality. Sometimes we have good reason to believe that the degree to which some individualized evidence confirms or disconfirms the attribution of $P$ to $Y$ is much lower than the degree to which some alternative, non-individualized evidence does. In such cases, failing to consider the individualized evidence does not mean we fail to respect individuality to an appropriate degree. A satisfactory account of respecting individuality needs to incorporate both the prioritization account and the revised idealistic account.

Our proposal to unify the two accounts is to give individualized evidence extra weight in the *measure of appropriateness*. More specifically, in a practice $F$ that forms a statistically grounded belief,

> (A) if $e$ (evidence considered in $F$) contains only non-individualized evidence, and $e'$ (evidence neglected in $F$) contains individualized evidence, then to reach the appropriate degree of respecting individuality, the following must be *true*:

[14] One may worry that the "social meaning" of a mechanism can be morally problematic. For example, there might exist a mechanism of racialization that determines how much money people of different race earn, and we can reasonably believe that this mechanism operates on everybody. It follows that, on our account, "race" qualifies as individualized evidence for how much money one earns. We think, if there indeed exists such a mechanism, "race" does count as individualized evidence. In reality, race does not count precisely because we cannot reasonably believe such a universal mechanism exists. Moreover, calling "race" individualized evidence in this case does not suggest that the mechanism itself is morally justified. The latter is a completely different issue.

$\sum_e(\mathrm{O}(e) \cdot \mathrm{C}(e)) >> \sum_{e'}(\mathrm{NO}(e') \cdot \mathrm{I}(y)/\mathrm{C}(e'))$, where >> means "being much greater than";
(B) if *e* contains individualized evidence, and *e'* contains only non-individualized evidence, then to reach the appropriate degree of respecting individuality, the following must be *false*:
$\sum_{e'}(\mathrm{NO}(e') \cdot \mathrm{I}(y)/\mathrm{C}(e')) >> \sum_e(\mathrm{O}(e) \cdot \mathrm{C}(e))$.

Roughly, what this means is that, other things being equal, individualized evidence weighs much more than non-individualized evidence. A belief-forming practice that fails to consider individualized evidence can nevertheless respect individuality to an appropriate degree, but only if the degree to which the individualized evidence neglected in *F* disconfirms the belief is known to be much lower than the degree to which the non-individualized evidence that has been considered by *F* confirms the belief, or the cost to obtain that individualized evidence is too huge, or the impact of the belief on the individual at issue is too insignificant.

Let us now consider some real-life cases that might be thought to cause problems for this combined account. First, consider the following case.

> *Suspect*: A bank robber was reported to be running away in a red Toyota car. You happened to be driving a red Toyota car, and the police pulled you over for a search.

It might be argued that driving a red Toyota car should count as non-individualized evidence on our account, as no causal mechanism can be reasonably believed to be working in or on you that connects driving such a car to being a bank robber. But intuitively driving a red Toyota car is individualized evidence in this context.

We think there is a logical connection between driving a red Toyota car in the neighborhood and being the bank robber: a red-Toyota-car-driving individual in the neighborhood is logically more likely to be the red-Toyota-car-driving bank robber than a non-red-Toyota-car-driving individual. And this logical connection constitutes an individual "propensity"—that is, it can be reasonably believed to be true of you. This is in sharp contrast with the practice of racial profiling—subjecting individuals of certain racial groups to increased scrutiny "based on the expectation… that members of such groups are statistically overrepresented among those involved in certain types of criminal behavior" (Mogensen 2017: 4). There exists no such logical connection between being a member of certain racial groups and being a criminal that can be reasonably believed to constitute an individual propensity: higher criminal rate is a property of a group, not of any individual member of that group. Thus, the justification for connecting membership of certain racial groups to criminal activity is not grounded on something that can be reasonably believed to be true of the individual.

Moreover, as some theorists pointed out, widespread practice of racial profiling tends to disproportionately burden certain racial groups or strengthen existing prejudices against them and thereby exacerbate the disadvantages that they are already suffering (e.g., Lippert-Rasmussen 2014: 169–170; Khaitan 2015: 197; Mogensen 2017). This greater long-term impact increases the weight of the relevant individualized evidence neglected by the belief-forming practice, and thus makes it

harder for practice of racial profiling to reach the appropriate degree of respecting individuality.

Next, let's consider a special type of evidence: testimonial evidence. Testimony is often accepted on the basis of the trustworthiness of the witness. But what grounds the connection between the testimony and our belief that it is true has nothing to do with *the individual that the testimony is about*; it simply depends on facts about the witness, such as how often her testimony about other things turned out to be accurate. Thus, one may worry that our account would mistakenly count testimonial evidence as non-individualized evidence.

When we come to believe that *Y* has *P* based on testimony, our epistemic task lies not in assembling direct evidence about *Y* or *P*, but in assessing the quality of the testimony, which often involves, among other things, inquiring into the witness's grounds for making the relevant claim about *Y*. That is to say, the duty to respect individuality in such cases means not that we have to directly look into *Y*'s particular circumstance, but that we have to take into consideration whether the witness properly respected *Y*'s individuality in forming the testimony. For example, in *Suspect* we could have a testimony that says "*Y* is the bank robber". Suppose that the witness's reason for making that claim is "I saw that *Y* was robbing the bank". Since there is a connection between being seen robbing the bank and actually robbing the bank that can be reasonably believed to apply to *Y* (just as in every scientific experiment, the observation of an event is taken to mean its existence[15]), the testimony is based on individualized evidence about *Y*. If, on the other hand, the witness's reason is "*Y*'s brothers were bank robbers", then the testimony turns out to be based on non-individualized evidence, as there is no causal, logical, or conceptual mechanism that connects one's brothers being bank robbers to one's committing the bank robbery. The duty to respect individuality will then require us not to take the witness's testimony.[16]

## 5 Some alternative accounts

Let's now compare this combined account to some alternative accounts. Thomson (1986) proposes a causal-relation account of what counts as individualized evidence in the legal context. According to Thomson, "individualized evidence for a defendant's guilt is evidence which is in an appropriate way causally connected with the (putative) fact that the defendant is guilty" (1986: 214). A piece of evidence can be appropriately causally connected with the putative fact either because it causes

[15] As we discussed in the previous section, the possibility that the observation is incorrect is a flaw internal to the observational model, not one pertaining to respecting individuality.

[16] This explains, in the famous Blue Bus Case, why a witness's testimony of seeing a blue bus passing through the area should be treated differently from the pure statistical evidence that a bus from the blue bus company has a high probability of passing through that area during that specific time. Even though a witness's testimony can be mistaken, it is formed through an epistemic process that respects the defendant's individuality. Consequently, insofar as the testimony's reliability is not much lower than that of non-individualized counterevidence, using that testimony properly respects the defendant's individuality.

that fact, or it is caused by that fact, or they share a common cause. For example, Lin's good tryout performance and the putative fact that he is a good shooter share a common cause—his good shooting skills, and this is why, on Thomson's account, good tryout performance counts as individualized evidence for Lin's being a good shooter. Thomson believes that an appropriate causal connection generates a "(putative) guarantee" that the relevant belief is true (1986: 214).

Thomson's account faces at least two major worries (e.g., Blome-Tillmann 2015; Eidelson 2015; Gardiner 2019). One worry is that this account may mistakenly turn some purely statistical evidence into individualized evidence. For example, it is the statistics about middle-aged white men appearing in the shopping district that caused the peddler to form the relevant belief in *Peddler 2*. Another worry is that this account denies even the possibility of individualized evidence when "the putative fact does not obtain" or "the evidence misleads" (Gardiner 2019: 184). For example, in *Peddler 3*, since Austin is in fact not in a relationship with the woman standing next to him, there cannot be any causal connection between the peddler's evidence (i.e., Austin's uttering the word "sweetie") and the putative fact that Austin is in a relationship with that woman—this fact simply does not obtain. Yet it does not follow that there cannot be any individualized evidence for the peddler's belief, false as it is.

Despite some appearance of resemblance, our account is very different from Thomson's. Our account does not require there to be an actual causal connection between the evidence and the putative fact (i.e., the attribution of $P$ to $Y$); instead, what our account requires is that the statistical regularity between the evidence and the putative fact be grounded on some mechanism that can be reasonably believed to apply to the individual in question. This mechanism can be causal, logical, or conceptual in nature. The emphasis on a more inclusive notion of *grounds for the statistical regularity*, rather than on an actual causal relation, enables our account to avoid the problems that Thomson's account faces.

Eidelson (2015) proposes an alternative account of respecting individuality. According to Eidelson's account, to treat people as individuals is to respect the fact that they are autonomous agents. Respecting people's autonomy, in part, requires us to give "reasonable weight" to "relevant" and "reasonably available" evidence of how one "has exercised her autonomy in giving shape to her life" (2015: 144). Using this autonomy-focused account, we can distinguish pure statistical evidence from individualized evidence in the following way: a belief-forming practice that relies merely on pure statistical evidence fails to give reasonable weight to relevant and reasonably available evidence of the ways that people have exercised their autonomy to shape their lives; by contrast, a belief-forming practice that relies mainly on individualized evidence gives proper weight to such evidence.

This account also faces serious problems. On the one hand, it remains unclear precisely what counts as "relevant" evidence that reflects one's autonomous choices. From the autonomy point of view, as long as a piece of evidence reflects autonomous choices, there should in principle be no difference between pure statistical relevance and the kind of relevance that individualized evidence bears. For example, one's religion or nationality can be as much reflective of one's autonomous choices, and thus as relevant, as one's action or choice of words.

Therefore, in cases such as *Disease*, Eidelson's account, as it stands, offers no answer to how to distinguish pure statistical evidence from individualized evidence.

On the other hand, by focusing exclusively on autonomy, Eidelson's account is unable to explain discrimination against non-autonomous beings, such as infants and animals, since such objects presumably are not (yet) able to "exercise autonomy to shape their lives".[17] It is also unable to explain discrimination that picks out individual trait that reflects no exercise of autonomy. For example, Asian males may be regarded as unattractive because of the alleged characteristics of "Asian appearance", such as epicanthic fold and low nasal bridge (Liu 2015). But not all Asian males have these characteristics. So, there is good reason to think that viewing Asian males as unattractive merely by virtue of the alleged "Asian appearance" disrespects individuality. But this cannot be accounted for by the autonomy-focused account, since one's physical appearance usually does not reflect any exercise of autonomy. Concerns about individuality in interpersonal relationships do not completely overlap with concerns about autonomy.

Some other theorists recently propose to distinguish individualized evidence from pure statistical evidence by appealing to such epistemic principles as sensitivity and safety. Your belief that $p$ based on the evidence $e$ is *sensitive* just in case, in the closest possible world where $p$ is false and you have $e$, you do not believe that $p$ (Greco 2012: 196; Enoch et al. 2012: 204). Your belief that $p$ based on the evidence $e$ is *safe* just in case, in close possible worlds where you believe that $p$ on the basis of $e$, $p$ is true (Greco 2012: 196; Pardo 2018: 68). Based on such principles, these theorists argue that beliefs based on pure statistical evidence are either not sensitive or not safe, but beliefs based on individualized evidence are sensitive and safe.

It is unclear that sensitivity really helps to establish such a distinction. For example, Peddler 3's belief based on the evidence that Austin appears to be saying "sweetie" to the woman standing next to him fails to be sensitive: even in the actual world, it is not the case that the belief is false and the peddler has the evidence, and he does not hold that belief. Yet, the evidence is indeed individualized evidence.[18] Safety faces difficulty too. For example, in *Peddler 2*, if the group statistic is high enough (e.g., 90% of middle-aged white men appearing next to a young Chinese woman in this area are actually in a relationship with that woman), the peddler's belief will be safe—in close possible worlds where another middle-aged white man appearing next to a young Chinese woman, they are actually in a relationship. However, the relevant evidence is pure statistical evidence. Safety, in its essence, is meant to capture the reliability of the belief-forming method (Greco 2012: 193); thus as long as a piece of evidence has sufficiently high likelihood to produce true belief, it will satisfy the safety principle, even if it is purely statistical.

[17] Lippert-Rasmussen raised a similar worry by using cases involving "non-autonomous minors" (2011: 53). It is debatable whether minors are truly "non-autonomous". But the general point behind his worry is valid.

[18] For more detailed discussion on the difficulties facing the sensitivity account, see Blome-Tillmann (2015) and Pardo (2018).

Moreover, even if appealing to these epistemic principles somehow can avoid the difficulties mentioned above, they seem unable to capture the moral requirement of respecting individuality. Respecting individuality is essentially a subjective duty, but sensitivity and safety, construed in terms of possible worlds, are objective in nature. Therefore, respecting individuality constitutes a unique dimension in evaluating statistically grounded beliefs that a pure epistemic approach fails to capture.

## 6 Conclusion

Generalization plays an indispensable role in our everyday belief-forming practice. Finding the right person for help by looking at what one wears, understanding what people want by listening to what they say, and diagnosing diseases by reading one's symptoms all employ some statistical evidence and sometimes pure statistical evidence. But the fact that pure statistical evidence is justified in some contexts does not mean it is justified in all contexts. In this paper, we argued that statistical generalization is subject to one specific moral constraint—respecting individuality, and we developed an account of what it means to respect individuality.

On our account, respecting individuality is a matter of degree. The appropriate degree depends on the comparative weight between the relevant evidence that has been considered and the relevant evidence that has been neglected. Moreover, unlike other accounts in the literature, our account holds that whether individuality is appropriately respected does not depend on the nature of the identifying property, nor does it just depend on the identifying property's degree of relevance to the attributed property (i.e., the degree to which having the identifying property confirms or disconfirms having the attributed property). Although the degree of relevance matters, whether individuality is appropriately respected depends also on our justification for it. When the degree of relevance is grounded upon something that can be reasonably believed to be true of the relevant individual (such as some physical propensity), we call the evidence "individualized evidence", which should be prioritized over non-individualized evidence, other things being equal. By calling it "individualized evidence", we do not pretend that our use of this phrase perfectly matches how people normally use it; our purpose is simply to highlight this often neglected step in forming statistically grounded beliefs and its significance to respecting individuality.

Of course, respecting individuality is not the only relevant moral factor in evaluating statistically grounded beliefs. While other moral factors have been extensively discussed in the literature, respecting individuality as a legitimate reason for thinking why statistical discrimination is wrong has often been quickly dismissed, on the ground that there is no real difference between the so-called individualized evidence and non-individualized evidence. We hope this paper provides some reason for reconsideration.

**Acknowledgements** We want to thank Kasper Lippert-Rasmussen, Peter Vallentyne, Eddy Keming Chen, Wilfried Hinsch, Ryan Hornbeck, and one anonymous reviewer for this journal and one anonymous

reviewer for another journal for very helpful criticisms and suggestions on earlier drafts. This work was presented at the Philosophy Departments of Xiamen University and Nanjing University, and we want to thank the participants, especially Andrea Strollo, for helpful comments. We are also indebted to the "Principles of Cultural Dynamics" program at Free University of Berlin, which funded the research leave that enabled the writing of this paper.

## References

Basu, R. (2018). The wrongs of racist beliefs. *Philosophical Studies*. https://doi.org/10.1007/s11098-018-1137-0.

Beeghly, E. (2018). Failing to treat persons as individuals. *Ergo, 5*(26), 687–711.

Blome-Tillmann, M. (2015). Sensitivity, causality, and statistical evidence in courts of law. *Thought, 4,* 102–112.

Blome-Tillmann, M. (2017). More likely than not'—Knowledge first and the role of bare statistical evidence in courts of law. In J. A. Carter, E. C. Gordon, & B. Jarvis (Eds.), *Knowledge first—Approaches to epistemology and mind* (pp. 278–292). Oxford: Oxford University Press.

Blum, L. (2002). *"I'm not a racist, but…": The moral quandary of race*. Ithaca: Cornell University Press.

Bolinger, R. (2018). The rational impermissibility of accepting (some) racial generalizations. *Synthese*. https://doi.org/10.1007/s11229-018-1809-5.

Churchman, C. W. (1948). *Theory of experimental inference*. New York: Macmillan.

Di Bello, M. (2019). Trial by statistics: Is a high probability of guilt enough to convict? *Mind, 128,* 1045–1084.

Douglas, H. E. (2000). Inductive risk and values in science. *Philosophy of Science, 67*(4), 559–579.

Eidelson, B. (2015). *Discrimination and disrespect*. Oxford: Oxford University Press.

Enoch, D., & Fisher, T. (2015). Sense and 'sensitivity': Epistemic and instrumental approaches to statistical evidence. *Stanford Law Review, 67,* 557–611.

Enoch, D., Spectre, L., & Fisher, T. (2012). Statistical evidence, sensitivity, and the legal value of knowledge. *Philosophy & Public Affairs, 40*(3), 197–224.

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford: Oxford University Press.

Gardiner, G. (2019). Legal burdens of proof and statistical evidence. In D. Coady & J. Chase (Eds.), *The Routledge handbook of applied epistemology* (pp. 179–195). New York: Routledge.

Greco, J. (2012). Better safe than sensitive. In K. Becker & T. Black (Eds.), *The sensitivity principle in epistemology* (pp. 193–206). Cambridge: Cambridge University Press.

Hellman, D. (2018). The epistemic commitments of nondiscrimination. *Virginia Public Law and Legal Theory Research Paper No. 2018-60*.

Hempel, C. G. (1965). Science and human values. In C. G. Hempel (Ed.), *Aspects of scientific explanation and other essays in the philosophy of science* (pp. 81–96). New York: Free Press.

Keller, S. (2007). *The limits of loyalty*. Cambridge: Cambridge University Press.

Khaitan, T. (2015). *A theory of discrimination law*. Oxford: Oxford University Press.

Levin, M. (1992). Responses to race differences in crime. *Journal of Social Philosophy, 23*(1), 5–29.

Lippert-Rasmussen, K. (2011). 'We are all different': Statistical discrimination and the right to be treated as an individual. *Journal of Ethics, 15,* 47–59.

Lippert-Rasmussen, K. (2014). *Born free and equal?: A philosophical inquiry into the nature of discrimination*. Oxford: Oxford University Press.

Liu, X. (2015). No fats, femmes, or Asians. *Moral Philosophy and Politics, 2*(2), 255–276.

Miller, D. (1999). *Principles of social justice*. Cambridge: Harvard University Press.

Mogensen, A. (2017). Racial profiling and cumulative injustice. *Philosophy and Phenomenological Research*. https://doi.org/10.1111/phpr.12451.

Moss, S. (2018a). *Probabilistic knowledge*. Oxford: Oxford University Press.

Moss, S. (2018b). Moral encroachment. *Proceedings of the Aristotelian Society, 118*(2), 177–205.

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

Pardo, M. (2018). Safety vs. sensitivity: Possible worlds and the law of evidence. *Legal Theory, 24,* 50–75.

Robert, C. (2007). *The Bayesian choice: From decision theoretic foundations to computational implementation*. New York: Springer.
Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science, 20*(1), 1–6.
Schauer, F. (2003). *Profiles, probabilities, and stereotypes*. Cambridge: Harvard University Press.
Schroeder, M. (2018). Rational stability under pragmatic encroachment. *Episteme, 15*(3), 297–312.
Stroud, S. (2006). Epistemic partiality in friendship. *Ethics, 116*(3), 498–524.
Thomson, J. J. (1986). Liability and individualized evidence. *Law and Contemporary Problems, 49*(3), 199–219.