9/7/25

The Discrimination Paradox

Lee Jussim, Rutgers

Lee Jussim
Department of Psychology
Rutgers University
53 Ave E
Piscataway, NJ 08854

jussim@psych.rutgers.edu
908-963-9522
Fax: 908-262-2026
Word count: 6365
Keywords: Discrimination

Ethical Approvals
Not applicable. It's a review article with no original data.

Conflict of Interest
I have no conflicting or competing interests.

Funding
None.

Availability of Data and Raw Materials
Not applicable, because there are neither data nor materials to make available. Computations for Table 1 are available in an online supplement at: https://osf.io/rxtjs/

Abstract

Rigorous studies published within the past eight years have found diametrically opposed results regarding racial discrimination. Some have found that racial discrimination is very rare; others that racial discrimination is very common. The paradox is that they are all well-conducted studies. In this paper, I show why there is no paradox, and the two sets of findings are completely compatible.

In the last few years, several studies of racial discrimination have been published, *seeming* to produce wildly different results. The entire set of studies focus on *individual discrimination* rather than "systemic racism." The focus of the present review is four papers including dozens of studies: (1) a meta-analysis of audit studies of employment discrimination (Quillian et al., 2017); and (2) three large sample papers assessing discrimination in either the real world (Campbell & Brauer, 2021; Nødtvedt et al., 2021) or experimental laboratory (Peyton & Huber, 2021). The meta-analysis produced evidence of substantial discrimination, the others of low levels of discrimination.

This is not a comprehensive review of the vast literature on racial discrimination. Instead, my focus is limited to understanding why results of studies finding dramatically different levels of discrimination are compatible with one another.

The Discrimination Paradox

Because recent studies find both very high and very low levels of racial discrimination, it might seem that something is wrong somewhere. Perhaps there are deep flaws in the studies finding little discrimination but not in the ones finding substantial discrimination (or vice versa). Perhaps the situations are too different to justify any comparison.

These are possible, but I don't think so. Instead, I believe that the studies described herein are all strong, credible and generalizable studies. If that is true, we have an apparent paradox of strong studies producing seemingly strikingly contrasting findings. Therefore, I review these studies next.

The Meta-Analysis

Audit studies are a class of experimental studies, conducted in the real world, wherein targets who are otherwise identical (e.g., they have identical or equivalent resumes) differ on some demographic characteristic and apply for something (such as a job). It can be almost any demographic characteristic, but herein I focus primarily on those manipulating whether the job applicants are Black or White. The main outcome is whether Black or White applicants are treated similarly (i.e., in an egalitarian manner) or differently (i.e., discrimination occurs, say, in callbacks or requests for interviews).

A review and meta-analysis (Quillian et al, 2017) found 21 audit studies of racial discrimination in hiring since 1989 and three additional studies going back to 1972. The studies included over 55,000 applications submitted for over 26,000 jobs. There were two headline findings: 1. On average, White applicants received 36% more callbacks than did Black applicants. 2. This difference did not decline between either 1972 or 1989 and 2015. Indeed, there was weak evidence that it had increased over that time.

Audit studies are not without limitations. They typically assess callbacks or interview requests, rather than actual hiring (applicants are fake so there is no one to hire). They focus mostly on entry level jobs. I know of no research examining whether minority-owned business discriminate against White people, which, to evaluate the *net* effect of discrimination, needs to be known. Depending on who one counts as White or not, about 25-40% of Americans are not White (U.S. Census, 2024) and about 20% of business are minority owned (Lee, 2023). If minority-owned businesses discriminate against White applicants, net discrimination will be less than the headline number in audit studies that test for discrimination in White-owned businesses. Because both field and experimental studies have consistently found that minorities, on average, engage in more ingroup favoritism than do majorities (e.g., Leonardelli & Brewer, 2001; Mullen, Brown & Smith, 1992), the extent to which discrimination against White applicants reduces net discrimination against Black applicants nationally may exceed the proportion of minority owned businesses, though I am aware of no research that has directly examined this question.

Despite these limitations, audit studies in general and the Quillian (2017) meta-analysis in particular have major strengths which mean they should not be dismissed. As actual experiments (rather than surveys, studies of "gaps," or other correlational type studies) they can assess whether applicant demographics (in this case, race) *cause* them to receive better or worse employment application treatment. Because they are conducted in the real world, assessing what goes in in real employment situations, they cannot be dismissed as trivial situations concocted in ivory tower laboratories.

36% is high enough to contribute something substantial to racial employment and income gaps. It occurred in more than 1 out of 3 employment applications, over 40 years, and was not any lower in more recent studies. That is a lot of discrimination. And yet, the result seems to conflict with findings from other high quality recent studies of discrimination.

Recent Studies Finding Very Low Levels of Acts of Discrimination

A slew of recent studies has found that racial discrimination occurs very infrequently. I review those next.

Ultimatum game study. The first (Peyton & Huber, 2021) found anti-Black discrimination occurred 1.3 percent of the time. In the study, they had over 700 people play the ultimatum game with either Black or White partners. Variations of the ultimatum game are often used in experimental studies to assess how people make decisions. Typically, the first player proposes to the second how to divide some money. For example, the first player may be given a dollar to divide with a second player but also has the power to make a take it or leave it offer (the ultimatum). For example, the first person may offer 30 cents to the second. If the second player accepts, then the first gets 70 cents and the second gets 30 cents. If the second rejects this division, neither gets anything.

In the Peyton & Huber (2021) study, participants played the ultimatum game 25 times with either Black or White partners, so the total number of offers accepted or refused was over 18,000. They operationalized racial discrimination as occurring when White players rejected offers from Black players that would have been accepted had the person offering been White.

The abstract of the paper emphasizes "racial resentment" and "explicit prejudice." Indeed, the last sentence declares that "explicit prejudice is widespread." However, to be clear about their main result regarding discrimination, I quote from the paper directly (pp. 30-31):

> The first estimate, a 1.3 percentage point decrease (p<.01) in the probability of acceptance, shows that, on average, white responders engaged in anti-Black discrimination by rejecting offers they would otherwise accept if the proposer was white (M1.1).

Discrimination occurred 1.3% of the time. One might wonder why the authors did not highlight this remarkable finding in either the abstract or their discussion.

Of course, this study had some important limitations. The participants were Mechanical Turk workers, which is important because they are not a representative sample of Americans. Whether the 1.3 percent figure would generalize to "Americans" is unknowable from this study. Also, whereas the 98.7 percent nondiscrimination is very high, it was not a real-world context. Although this renders its implications for real-world discrimination unclear, the next two papers addressed discrimination in the real world.

University of Wisconsin-Madison college student study. Another study (Campbell & Brauer, 2021) examined college students' behavior as they went about their days on campus at the University of Wisconsin-Madison. It included five surveys, eight experiments and a meta-analysis examining discrimination. Although I focus exclusively on the seven experiments addressing racial/ethnic discrimination and discrimination against Muslims, results were similar for the study addressing discrimination against homosexuals. All studies examined naturally occurring interactions, such as door-holding, asking directions, and sitting next to a target on a bus, as students went about their business on campus.

In each study described here, the researchers enlisted an *actor* — someone to play a part in the study unbeknownst to the students whose discriminatory behaviors they assessed. For example, in the racial discrimination studies, they enlisted Black and White actors, so they could compare student behavior towards a Black person or a White person. Some studies compared behavior towards a White person or a Muslim (operationalized as a female actor wearing a hijab). The actors were trained to behave identically.

The discrimination assessments begin with Study 5, which found that 5% of students held a door for a White person but not for a Black person. Study 6a found that a White actor requesting directions received them 9% more often than an Asian actor and 6% more often than a Muslim actor. In Study 7a, a White actor received help 18% more often than did a Muslim actor, but 20% less often than did an Asian actor. Study 8 found a Muslim actor was treated with more social distance on a bus 6% of the time.

Studies 9a and 9b were job application audit studies (like those of the Quillian et al meta-analysis). Study 9a found that a White applicant received 7% more responses than an Arab applicant. Study 9b found that a White applicant received 8% more responses than did a Black applicant.

They also conducted a meta-analysis of their studies that included only the Black and Muslim targets. There was statistically significant evidence of discrimination, indicating a small overall tendency to favor White targets. Simply averaging the differences for these groups produced an overall discrimination rate of about 8%, which is the same as a 92% non-discrimination rate. Clearly, non-discrimination is the overwhelming pattern, although discrimination occurred systematically, albeit at a low rate.

Of course, these studies also had limitations. They were only conducted among college students at a single university. Their generalizability is unknown. Most of the situations were relatively trivial. Even the audit studies, like audit studies generally, did not examine actual hiring.

Nonetheless, I think the strengths of these studies outweigh their weaknesses. First, there were over 1400 participants across the studies described herein. Another major strength was that discrimination was operationalized in a wide variety of ways. Third, they examined discrimination with respect to several different groups. Fourth, all studies were experiments, so that differences in treatment by race/ethnicity/religion of target reflected causal effects, not mere gaps or correlations.

Discrimination in Airbnb responses. Nødtvedt et al. (2021) examined discrimination in the selection of Airbnb listings among a nationally representative sample of 801 Norwegians. The host was identified either as ethnically Norwegian or ethnically Somali. Overall, there was a 9.3 percent preference for the listing by the Norwegian ethnic.

The authors (in their Public Significance Statement) declared that "When an identical Airbnb apartment was presented with a racial outgroup…host…[people were] 25% less likely to choose the apartment over a standard hotel." So how do I get 9.3%? Here is the relevant text from their results section (p. 521): "Results revealed that participants chose the Airbnb apartment (vs. hotel) significantly more often when the Airbnb host was an in-group member ***(38.4%,*** 95%

CI [33.6, 43.1]) compared to when the host was an out-group member (***28.9%***, 95% CI [24.7, 33.5], $\chi 2(1, 801) = 7.80$, $p = .005$, ***proportion difference = 9.3%***, 95% CI [2.8, 15.8]) [emphasis mine]."

This foreshadows the resolution to the Discrimination Paradox. Nonetheless, the *percent difference* in listings chosen (rather than the *percentage difference of the percentages of listings chosen)* was 9.3%. Actually 38.4 - 28.9 = 9.5, and even with rounding error, I can't get it below 9.4, so either I am missing something, or this is a minor error. Still, 9.3 is what was reported.

The bottom line is that acts of discrimination occurred 9.3% of the time. This meant that outgroup hosts were chosen 25% less often (28.9%) than when the host was an ingroup member (38.4%; 28.9/38.4=75%, i.e., outgroup hosts were selected 75% as often or 25% less often) . One may wonder why the authors chose to emphasize 25%, which appears in the abstract, significance statement, results and discussion, and de-emphasize 9.3% (which only appears in the results even though it is arguably a clearer description of the difference because it does not have a potentially variable denominator).

Of course, this study also had limitations. It was conducted in Norway and whether its results generalize to the U.S. (the focus of this article) or anywhere else is an open empirical question. It only focused on Airbnb, so it provided no information about other contexts in which discrimination might manifest. Nonetheless, it was a well-conducted experiment conducted in the real world with a substantial sample size, so my judgment is that the results deserve to be taken seriously, despite the study's limitations.

The Discrimination Paradox

Taking these findings altogether, we have the Discrimination Paradox. The high quality meta-analysis of audit studies found job discrimination at 36% (Quillian et al., 2017); the recent studies reviewed in detail just above (Campbell & Brauer, 2021; Nødtvedt et al, 2021; Peyton & Huber, 2021), along with many other studies (e.g., see Ferguson & Smith, 2024 for another meta-analysis) find discrimination at very low levels, typically single digits.

When I first discovered the low discrimination studies, I re-examined them very carefully and then did the same with the Quillian et al. (2017) meta-analysis, with an eye towards trying to resolve the seemingly deep discrepancies. Perhaps one or more were weak studies and the others were rigorous; perhaps the different operationalizations of discrimination could explain the different findings. And then I realized that it was not necessary to figure out which studies were not credible because the seemingly contradictory results are entirely compatible.

There is No Discrimination Paradox

There is an apparent conflict, but no actual conflict. To see how and why requires some math at about 7th grade level.

There is no single number for the amount of discrimination a group experiences. Discrimination varies in type (hate crimes, harassment, exclusion, etc.), and there are many different methods for assessing discrimination, which often yield different estimates. Still, to illustrate the Discrimination Paradox, I need to use actual numbers. I use the 36% figure, based on the Quillian et al. (2017) meta-analysis of audit studies of racial discrimination in hiring as the starting point.

Levels of analysis: Acts of discrimination versus experiences of discrimination. One key to resolving the Discrimination Paradox is understanding that discrimination can be assessed at two different levels of analysis. The 36% figure obtained by Qullian et al. (2017) is based on the differences in callbacks received by Black and White applicants. It is a difference between *the experiences* of Black and White applicants. In contrast, the findings of single digit discrimination (Campbell & Brauer, 2021; Nøtvedt et al., 2021; Peyton & Huber, 2021) addressed *acts* by potential perpetrators of discrimination. This difference is critical in resolving the Discrimination Paradox.

A simple hypothetical example can show the striking differences in discrimination obtained within a single study, depending on whether one focuses on acts or experiences of discrimination. Consider a simple audit study of 50 Black and 50 White job applicants. I'll call this Study A because I return to it later. If 10 Black and 15 White applicants receive callbacks, Black applicants received 33% fewer callbacks. However, there are, at most, 5 acts of discrimination (as shall be explained below, it may be fewer depending on assumptions), out of 100 callback decisions. Acts of discrimination were, at most, 5%. There is no incompatibility between the 33% difference in experiences of discrimination and the 5% rate of acts of discrimination.

Discrimination rates. With respect to discrimination, it is almost always useful to unpack whatever result is presented into its underlying rates. For example, Study B may report something like "Group 1 (e.g., some outgroup) obtained good outcome X (e.g., a job callback, beneficial treatment in the Ultimatum Game, help getting directions, etc.) 25% less often than did Group 2." The problem with this sort of reporting is that it obscures potentially profound differences in frequencies and rates of both acts and experiences of discrimination which can create the appearance of the discrimination paradox.

Any rate is a numerator divided by some denominator. Variations in rates (or proportions), therefore, can derive either from variations in the numerator or the denominator. For example, assume that, like Study A, Study B also found that Black applicants received 33% fewer callbacks. However, in Study B, 30 Black applicants received callbacks whereas 45 White applicants received callbacks. Although in both cases the outgroup received 33% fewer callbacks, there clearly was far more discrimination in Study B than Study A.

Reporting of percentage differences is not "wrong" – it provides useful information about the relative experiences of the two groups. But it is incomplete, at least if one wishes to know *how much* discrimination occurred. However, there are often different plausible choices for both the numerator and denominator of any calculation of discrimination rates, proportions, or the relative experiences of different groups, which are discussed next.

Choice of numerator: Alternative assumptions for what constitutes discriminatory behavior. To resolve the discrimination paradox a necessary condition is to determine *the rate at which acts of discrimination were committed.* Part of this involves choosing the numerator – what value constitutes the amount of discrimination to be divided by a denominator to get some rate (or proportion) of discrimination?

Conceptually, acts of discrimination can be determined by discrepancies from egalitarian behavior. Egalitarian behavior, however, can be determined in two different ways in most studies of discrimination, which produce, respectively, lower and upper bounds on how many acts of discrimination were committed. Under what I term *the zero sum assumption,* the lower bound on discrimination is produced by assuming the upper bound on positive treatment is the actual *total* amount of positive treatment possible (of both groups, e.g., callbacks, choosing an Airbnb listing, etc.). If 30 out of 100 applicants (regardless of group membership) receive callbacks, then the total amount of positive treatment is 30 callbacks, and this is assumed to be the total amount of positive treatment possible.

This is reasonable at least sometimes. For example, employers typically have limited resources and can decide in advance that there will be a fixed number of applicants to be called back (e.g., 30). If this is the case, callbacks are zero sum – for every White person who receives a callback, one less Black person receives a callback and vice versa. Thus, I refer to this as the *zero sum assumption.*

Assuming an equal number of ingroup and outgroup members (e.g., Black and White applicants, as in most audit studies) for simplicity of illustration (although the point holds for unequal numbers), under the *zero sum assumption*, acts of anti-Black discrimination equal half the difference between callbacks to White applicants and Black applicants. For example, under this assumption, if White applicants received 21 callbacks and Black applicants received 15 callbacks, the employer decided in advance to make 36 callbacks. Egalitarian treatment would, therefore, constitute 18 callbacks for each group, meaning that 3 acts of pro-White discrimination occurred. This constitutes a lower bound on discrimination because, given empirical callback frequencies, it is not possible for there to be fewer than three acts of discrimination in this example.

However, in many instances, there may be no fixed number of positive acts. For example, employers may have the flexibility to not set a fixed number of callbacks in advance. Perhaps the applicant pool is particularly good, in which case callbacks may increase. Perhaps the initial callbacks did not produce a viable candidate, so another round of callbacks are initiated. In the extreme, callbacks may be entirely unlimited (up to the total applicant pool), and so I call this *the unlimited assumption.*

Under the *unlimited assumption*, callbacks are not zero sum. Continuing with the same example even though the employer only made 36 callbacks, more could have been made. In this case, discrimination equals the simple difference in callbacks made to White versus Black applicants. Thus, if there were 21 and 15 callbacks to White and Black applicants, respectively, there would be six acts of discrimination. This is an upper bound on acts of discrimination, because, given the callback frequencies, it is not possible for there to have been more than a net of six acts of discrimination.

Discrimination rates: Choice of denominator. Similarly, there are different plausible choices for the denominator, each of which leads to the computation of a discrimination rate that has different interpretations. To determine the rate at which discrimination occurred *overall*, the denominator needs to be the entire set of acts, both discriminatory and nondiscriminatory. If anti-Black discrimination occurred 100 times in an audit study with 500 Black and 500 White job applicants, then, overall, it occurred 10% of the time.

This answers the question, "what proportion of all acts in the study constituted discrimination? One might argue that, in a study finding anti-Black discrimination, the acts towards White applicants are irrelevant. But they are not irrelevant if one wishes to answer the question that led off this paragraph. Its importance is, perhaps, more obvious in a study that had more than two groups. Only one group will receive the "best" overall treatment (not counting a tie). Clearly, if one wants to know "what proportion of all acts in the study constituted discrimination?" one needs to account for both positive and negative treatment of all groups. Indeed, any group *may* be the target of discrimination, and if one group is not subject to any discrimination that is zero discrimination for that group.

However, it is also useful to answer the question, "At what rate was treatment of the outgroup per se discriminatory?" To answer this question in a Quillian et al. (2017) type of audit study, one must determine the rate at which Black applicants experienced discrimination. The denominator for answering this question is no longer the total number of acts. It is the total number of acts (both positive and negative) towards the Black applicants. Therefore, the number of Black applicants is the relevant denominator. If there were 100 acts of discrimination and 500 Black applicants, then 20% of the acts towards Black applicants were discriminatory.

One might think that only positive events (such as callbacks) matter for determining discrimination because negative events (such as noncallbacks) are nonevents. This, however, is not justified if one wishes to determine any rate of discrimination (excluding negative events might be useful for answering other questions besides "what was the rate of discrimination?"). Deciding not to call an applicant back is a decision in the same way that is deciding not to admit a college applicant or not grant a mortgage to a homebuyer. This is why, in a classic case, comparable *proportions* of men versus women admitted to different graduate programs was the basis for exonerating Berkeley from a suit alleging sex discrimination (Bickel, Hammel & O'Connell, 1975). A proportion in such cases is the ratio of those admitted to the total of those who applied, including those who were rejected. It is not enough to know whether more men than women (or vice versa) were admitted to some program, even assuming a comparable distribution of qualifications. If more men applied to a program, more could be admitted, in the complete absence of sex discrimination. Decisions to reject applicants are just as crucial for determining discrimination as are decisions to admit applicants.

Again, it is not that the choice of one denominator or another is right or wrong. Rather, different information is provided by choosing as the denominator the total set of acts versus the total experiences of the group against which discrimination occurred. Whether researchers "should" always be explicit about reporting rates and, when they do, about what denominator they used and why, and what the subsequent results mean, is a function of many considerations beyond the scope of this paper. Regardless, doing so would clearly help resolve the discrimination paradox because being explicit about these choices and their interpretations could help explain why widely varying yet entirely compatible "discrimination" results can be reported from the same data.

Next, I show that, although these differing assumptions lead to differences in estimates of acts and rates of discrimination, all can resolve the discrimination paradox. Because they focus on acts of discrimination rather than experiences of discrimination, they consistently produce far lower estimates of discrimination than are produced by comparing the simple proportion of time some group receives less positive treatment than is received by some other group.

Example 1. 500 Black and 500 White applicants. Consider a simple hypothetical audit study: 1000 people apply for a job. 500 are Black and 500 are White, with equivalent records. Further assume that, as per Quillian et al.'s (2017) results, White applicants receive 36% more callbacks than do Black applicants. In this hypothetical, I assume that a total of 236 callbacks are received, so White applicants received 136 callbacks and Black applicants received 100; i.e., White applicants received 36% more callbacks. Under the zero sum assumption, the total number of potential callbacks equals the actual number of callbacks, i.e., 236. If there were no discrimination, Black and White applicants would receive identical numbers of callbacks, 118 each.

How much discrimination must be *enacted* to get to 36%? To get to 36%, i.e., 136 callbacks for White applicants, discriminatory *acts* need to have occurred 18 times. (118+18=136). There are *1000 applicants* in this hypothetical. This means that, overall, out of the 1000 decisions to either call back an applicant or not, *acts of discrimination* occurred 18/1000 times, or ***1.8% of the time.***

If one wishes to answer the question, "What proportion of the time do Black applicants experience discrimination?" the numerator stays the same (18) but the denominator changes.

Rather than using all 1000 decisions to call back an applicant as the denominator, it is the 500 Black applicants. Therefore, the answer to this question is 3.6% of the time (18/500).

Under the unlimited (callbacks) assumption, callbacks to White applicants are the benchmark for fair treatment, so Black applicants should have received 136 callbacks, but only received 100. The difference, 36, is the number of acts of discrimination against Black applicants. 36 acts of discrimination out of a total of 1000 acts is 3.6%.

If one wishes to answer the question, “What proportion of the time do Black applicants experience discrimination?” under the unlimited assumption, the numerator (36) stays the same, but the denominator becomes 500, the number of Black applicants. Therefore, the answer to this question is 7.2%.

Despite their differences, these calculations resolve the discrimination because they all show that differences in experiences of discrimination occur are much higher than are acts of discrimination. They show how 36% discrimination from the targets’ standpoints results from acts of discrimination occurring only 1.8% to 7.2% of the time, depending on the question one asks and the assumptions one makes about what constitutes discrimination. ***There is no substantial conflict between the results of Quillian et al’s (2017) meta-analysis, and those of the recent studies finding single digit levels of discrimination.***

More examples. Table 1 displays several alternative scenarios, all showing the same type of phenomenon, whereby 36% of discrimination from the target’s standpoint (or slightly more) results from acts of discrimination occurring, depending on assumptions and questions, between 0.24% and 14.4% of the time. The first three rows purposely used equal numbers of Black and White applicants to make it easy to see the math underlying the resolution to the Discrimination Paradox. The second line captures **Example 1.**

The last three rows use numbers that more closely approximate the Black and White population proportions in the U.S. It would get tedious if I walked through it all because the computations get more complicated when the number of Black and White applicants are not equal. However, the online supplement is an Excel spreadsheet that includes the computations of every entry in Table 1. Therefore, I next go through the computations for a single example, showing how 37.5% experiences of discrimination result from 0.24% acts of discrimination in the last row (overall rate of acts of discrimination under the zero sum assumption).

First, I show where 37.5% Black experiences of discrimination comes from. In this example, Black applicants make up 20% of the applicant pool (200/1000); White applicants make up 80% (800/1000). If there was no discrimination, Black applicants would get 20% of the callbacks and White applicants would get 80%. In the example, I assumed that there are 52 callbacks in total, i.e., 5.2% of the applicants receive callbacks. 20% of 52 is 10.4. 80% of 52 is 41.6. In the absence of discrimination, the expected number of callbacks for Black applicants is 10.4; for White applicants, 41.6. Obviously, there are not fractional people, but those are the expected numbers of callbacks for Black and White applicants, respectively, in the absence of discrimination.

To conform to Quillian et al.'s (2017) findings of 36% discrimination, however, Black applicants would get fewer than 10.4 callbacks and White applicants would get more than 41.6. As indicated in the note to Table 1, in my example, I assume only a whole person can receive a callback. Therefore, the proportion of White applicants receiving callbacks has to exceed the proportion of Black applicants receiving callbacks by at least 36%.

Table 1: Resolving the Discrimination Paradox

Black Applicants	White Applicants	Callbacks to Black Applicants	Callbacks to White Applicants	Experiences of Discrimination Rate	Zero Sum Callbacks Acts of Discrimination Rate, All Acts	Unlimited Callbacks Acts of Discrimination Rate, All Acts	Zero Sum Callbacks Acts of Discrimination Rate, Only Black Applicants	Unlimited Callbacks Acts of Discrimination Rate, Only Black Applicants
500	500	200	272	36.00%	3.60%	7.20%	7.20%	14.40%
500	500	100	136	36.00%	1.80%	3.60%	3.60%	7.20%
500	500	50	68	36.00%	0.90%	1.80%	1.80%	3.60%
200	800	75	408	36.00%	2.16%	2.70%	10.80%	13.50%
200	800	38	207	36.18%	1.10%	1.38%	5.50%	6.88%
200	800	8	44	37.50%	0.24%	0.30%	1.20%	1.50%

Note: The Experiences of Discrimination column always equals at least 36%, to correspond to the findings of Quillian et al. (2017). Because only whole people can receive callbacks, when 36% produced a fraction, I rounded up. This is why, in the last two hypothetical examples, callback discrimination rates were slightly above 36%. “All Acts” includes callbacks to both Black and White applicants. “Only Black Applicants” only includes callbacks to Black applicants. See text for explanation of the differing assumptions under zero sum and unlimited callbacks.

As I show here, this results in Black applicants receiving only 8 (rather than 10.4) callbacks and White applicants receiving 44 (rather than 41.6) callbacks. Eight is 4% of the 200 Black applicants. 44 is 5.5% of the 800 White applicants. 5.5 is 37.5% higher than 4.0 (5.5-4.0=1.5; 1.5/4.0=.375). White applicants received 37.5% more callbacks (relative to their proportion of the applicant pool) than did Black applicants (relative to their proportion of the applicant pool).

Next I show this is produced by an overall acts of discrimination rate of ***0.24%*** under the zero sum assumption:

1. There are 1000 applicants.
2. To get from 10.4 callbacks for Black applicants to 8 (in the table) and from 41.6 to 44 callbacks for the White applicants, is 2.4. This means 2.4 acts of discrimination took place. Although there are no fractional acts of discrimination in this case, it is still the best estimate of discrimination. Had I used 2 acts of discrimination, the Experiences of Discrimination would have been below 36%; had I used 3, it would have been well above it.
3. 2.4 acts of discrimination out of 1000 applicants is ***0.24%.***

A Formal Simulation Yielding Similar Patterns

One contribution of the present paper comes from its simplicity. The research reviewed was straightforward, with the key results simply how often acts of discrimination or experiences of discrimination occurred. The examples showing that very low frequences of acts of discrimination can produce substantial experiences of discrimination required only some relatively straightforward math.

However, one limitation of this approach is that, to maintain simplicity, only a few examples were used. This limitation is, however, addressed by prior research using formal simulations of sex discrimination. A series of simulations determined the effects of acts of discrimination that varied from 0% to 4% on experiences of discrimination (Hardy III et al., 2022). Of course, 0% means no bias. However, even acts of discrimination as low as 2.2%, under some simulated conditions produced women receiving job offers at about half the rate of men.

With respect to understanding bias outside the lab, however, this simulation study has its own significant limitations. A meta-analysis of 85 audit studies of occupational gender discrimination found overall evidence of bias *in favor of women* (driven primarily by strong biases against men in female-typical jobs), weak evidence of bias against women before 2009 and no evidence of bias against women after 2009 (Schaerer et al., 2023). Thus, the Hardy III (2022) study does not generalize to biases against women in hiring in real world audit studies.

Nonetheless, Hardy III (2022) should not be dismissed entirely. Its results are consistent with those of the present paper and can be described this way: *If* even low frequencies of acts of discrimination occur, then experiences of discrimination can be quite large. The key is not to forget the "if" because if there is no bias, there will be no experiences of discrimination. Furthermore, if the bias is in the opposite direction than many might have expected, the analysis herein and in their simulations still apply, but work in the opposite direction.

A Case Study from Research

I now return to the Nødtvedt et al. (2021) Airbnb study described previously in the section on studies finding very low levels of discrimination as a concrete, real example of the Discrimination Paradox. To briefly recap, participants rated only one Airbnb apartment and were asked if they would choose the apartment they were shown. They found Norwegians were 25%

less likely to choose an Airbnb with a Somali host than with a Norwegian host. This occurred because Norwegians chose a Norwegian host 38.4% of the time and a Somali host 28.9% of the time (which equals 9.5%, but they reported it as 9.3%).

Choosing an Airbnb apartment fits the *zero sum* assumptions I described previously. When one chooses an Airbnb apartment, one *does not* choose any other Airbnb apartment. When one chooses an Airbnb apartment with a Norwegian host, one does not choose an Airbnb apartment with a Somali host (and vice versa). 400 people were shown a Norwegian hosted Airbnb and 401 were shown one hosted by an Somali (801 total). 38.4% of 400 is 153.6, so I round this up to 154 because one cannot choose a fraction of an Airbnb. 28.9% of 401 is 115.9, which I round up to 116.

How many *acts of discrimination* were there? To determine this under the zero sum assumption, one must first determine *what would nondiscrimination choices look like?* A total of 270 apartments were chosen (154 + 116). Therefore, if there was no discrimination, their participants would have chosen 135 apartments hosted by Norwegians and 135 by Somalis. Acts of discrimination, therefore, are the difference between 154 and 135, or 19. 19/801=.024. Acts of discrimination occurred 2.4% of the time.

Even though actually choosing an Airbnb apartment is zero sum, one could argue that, *in this study*, the decisions were not zero sum. More people could have chosen more apartments than actually chose apartments. Even if one believes the *unlimited assumption* here is more appropriate, acts of discrimination only occurred 4.7% of the time (the difference between 154 and 116 divided by 801, i.e., 38/801=4.7%).

In the Airbnb case, one could argue that the denominator for discrimination among postings by Somalis should be the total number of Somali listings, not the total number of listings. Thus, this value would 38/401 or 9.5%, which is the proportion of times Somali listings experienced discrimination out of all Somali listings.

In this study, Somali hosts were selected 25% less often than were Norwegian hosts ("experiences" of discrimination, in quotes because they created fake Airbnb listings so no one actually experienced discrimination), even though acts of discrimination only occurred 2.4% of the time. If one prefers the unlimited assumption, discrimination occurred 4.7% of the time. Even if one believes that only Somali listings should be the denominator, discrimination occurred 9.5% of the time.

Alternative Phenomena for Which One Could Compute Useful Rates That Do Not Address the Discrimination Paradox

Other rates one could compute might be important and interesting, but they are beyond the scope of the present paper, because they do not address the discrimination paradox.
The perspective presented here neither invalidates nor contests the informativeness of other computations. Instead, it shows how the much higher figures for experiences of discrimination than for acts of discrimination are not just mathematically compatible, but mathematically inevitable, entirely a function of the different phenomena being computed. It also goes beyond prior work by providing a mathematically clear and coherent method for determining overall rates of acts of discrimination – they are is simply the frequency of such acts divided by the total opportunities to discriminate.

Implications

The Discrimination Paradox becomes more extreme when outcomes are more competitive. Table 1 shows that, as callbacks (or, worse, jobs) become harder to obtain because fewer acts of discrimination are needed to produce larger experiences of discrimination. The proportion of acts of discrimination necessary to equal or exceed the Quillian et al. (2017) 36% rate of experiencing discrimination is very low when there are few callbacks.

Consider a simple example in which there are 500 Black and 500 White identically qualified applicants, and there were only four job openings. With no discrimination, 2 jobs would go to Black applicants and 2 would go to White applicants. Because there are only four jobs, this situation fits the zero sum assumption. In this case, a single act of discrimination would mean that three White applicants but only one Black applicant would receive a job offer. White applicants are 200% more likely to receive a job offer (three versus one), even though there are only 1/1000 or 0.01% acts of discrimination (or 1/500, or 0.02%, if one only considers the Black applicants). Thus, for more competitive selection situations (i.e., where fewer applicants make it into, say, a job or college), minimal levels of acts of discrimination can produce large disparities.

Small biases can be larger than they seem. Another implication is that minimal levels of acts of discrimination can have a substantial impact on the targets of discrimination. There are long-running debates about whether even small biases are important. For example, some have argued that even small biases can have societally important effects, in part, because they accumulate over time (Greenwald, Banaji & Nosek, 2015; Martell, Lane & Emrich, 1996). Others have argued, however, that *presuming* biases accumulate is a far cry from *providing evidence that they actually accumulate* (e.g. Oswald et al., 2015) or that studies that have tested for accumulation have generally not found it (Jussim, Careem, Goldberg, Honeycutt, N. & Stevens, 2025).

Nonetheless, the resolution to the Discrimination Paradox shows that some small biases do indeed likely produce larger disparities than one might assume, e.g., if one merely knew that acts of discrimination only occur in the single digits. Even if small biases do not accumulate, they still can inherently produce large amounts of discrimination, especially for highly competitive positions that are awarded to only a small proportion of applicants.

Everyone is right, even those who find the other side's views repugnant. The Discrimination Paradox can give some insight into *unnecessary* socio-political tensions. One side, let's call them ***anti-racists,*** often argue that Black people experience substantial levels of discrimination. Another side, let's call them ***racism skeptics*****,** deny this, arguing that discrimination is far lower than it once was and that the science does not find consistent evidence of strong discrimination. These two sides are often quite hostile to and dismissive of one another. Some anti-racists take the racism skeptics' views as proof that the racism skeptics are actually racists. For example, Williams, Faber & Duniya (2022) referred to the highly accomplished late psychologist Scott Lilienfeld as a racist because he had critiqued microaggression research. Some interpret some arguments of anti-racists as evidence that they are hopelessly anti-scientific (Krylov & Tanzman, 2023). But the Discrimination Paradox shows that, at least with respect to their understandings of discrimination, they both *may well be right*, at least some of the time and maybe much of the time.

"Systemic racism"? The resolution to the discrimination paradox can also give insight into some claims about "systemic racism." The term is in quotes because it is often used so casually among both some academics and some in the wider society that it has become almost impossible to pin down what the term means. The discrimination paradox is relevant here because, as explained below, almost any discussion of systemic racism implies that discrimination is widespread. If so, then there is an onus on those making the claim to produce quantitative evidence about just how widespread discrimination is.

Wikipedia (Institutional Racism, 2025) defines systemic racism as synonymous with institutional racism — policies and practices that produce discrimination, advantages, or harms to particular groups. I like this definition, because it is reasonably clear and coherent. Systemic racism involves a system, not just individual acts of bigotry. As such, it creates a clear onus on those invoking systemic racism to point to the *specific systems, political or economic policies or practices* producing discrimination. Both inside and outside of peer reviewed articles, systemic racism is sometimes invoked as if it "explains" inequality — but without identifying the specific policies or practices that cause inequality, the explanation is as scientifically vacuous as answering "Why did the chicken cross the road?" with "It was God's will."

For example, Banaji, Fiske & Massey (2021) referred to systemic racism as inequalities intrinsic to societal structure. However, they give no specific examples of a "societal structure" that is intrinsically racist, and then simply highlight the existence of inequalities as evidence for "systemic racism." Thus, the claim of systemic racism is tautological:

- Why is there a racial gap? Because of systemic racism.
- How do you know there is systemic racism? Because there is a racial gap.

Because systems supporting racism (slavery, Jim Crow, redlining) have existed, it is clearly possible to identify such systems. However, when no such specific systems are identified, the attribution of inequality to systemic racism is not empirically justified. However, if systemic racism is defined in this manner, empirical studies of *individual manifestations of prejudice or discrimination* cannot be invoked as "systemic racism" because an individual's beliefs, attitudes, or behavior is not a system.

For example, if systemic racism refers to *prejudice* (as in the Banaji et al. (2021) article), this is an attitude, and, therefore, an individual level phenomenon; it is not any type of a system. Although one could argue that some system causes individual prejudice, without identifying the specific system that has done so, the claim is, again, scientifically vacuous. A similar analysis appears in an article titled "The Structure of Racism in Color-Blind, "Post-Racial" America by one of the foremost proponents of critical race theory (Bonilla-Silva, 2025). While purporting to identify new systems of racism, Bonilla-Silva (2025) simply makes claims about inequality and individual discrimination and calls it a "system." This article is not the place to provide a full critique of Bonilla-Silva's analysis, so one example will suffice. He claims (p. 1363) "…a new dominant racial ideology emerged that I have labeled as 'color-blind racism…'" But individuals, not systems, hold ideologies. If one wishes to argue that some "system" causes ideology, racist or otherwise, one is still left having to identify the specific system doing so. Even if one wishes to argue that systems and individual ideologies exist in an endless feedback loop, one is still left having to identify a particular system; one cannot simply invoke "the system" in some quasi-supernatural manner to explain a phenomenon without identifying a particular system and showing empirically that it has some causal effects on ideology.

There are exceptions to this sort of specificity failure. Some articles have identified specific policies or practices alleged to constitute systemic racism. However, these are often easily contestable and, as such, might be persuasive primarily to those already sympathetic to the general perspective yet be unpersuasive to skeptics. Contested evidence, or contested conclusions based on agreed-upon evidence, often do not constitute scientific facts – which requires establishing something so definitively that it would be perverse to believe otherwise (Gould, 1981). For example, Braveman et al. (2022) highlight the water crisis of Flint Michigan (a predominantly Black small city) as an example of modern systemic racism without attempting to evaluate whether water crises are *generally* more common in predominantly Black areas (a necessary condition for invoking *systemic* racism, as opposed to, for example, an idiosyncratic case of poor government). Thus, Braveman et al (2022) failed to establish systemic racism as the scientific explanation for the Flint water crisis because they failed to even address, let alone rule out, alternative explanations.

Furthermore, discrimination can and does occur at *the individual level, without any institutional or organizational practices involved* (as shown by the studies reviewed above). Thus, to invoke "systemic racism" as an explanation, one cannot point merely to evidence that "discrimination exists" or that "inequality exists." If A explains B, then A and B must be different variables. If systemic racism *explains* inequality, one cannot refer to evidence of inequality as if it is evidence of racism. Instead, at minimum, two criteria must be met to justify the conclusion that systemic racism causes inequality: 1. The *particular racist system has to be identified independent of the inequality;* and 2. *One must provide evidence that the particular system so identified caused inequality.*

Last, sometimes the term *systemic racism* seems to mean *there is a lot of racism out there* (e.g., Bonilla-Silva, 2015; Syed, 2021). Racism is *"systemic,"* in this view, in its pervasiveness. This is a different perspective than one which requires identifying a particular system. It is plausibly interpretable as a form of *concept creep* (Haslam, 2016) – the definition of systemic racism is expanded to include not just systems of racism but any context that the user of the term deems as having "too much" racism or racism above some usually unarticulated threshold.

Regardless, if "systemic racism" is used to mean "*there is a lot of racism out there"* then studies of individual prejudice and discrimination (such as audit studies and the other studies reviewed herein) are relevant. Perspectives arguing for large amounts of racism are not consistent with the rigorous studies finding minimal levels of acts of discrimination reviewed herein. But, given that even such minimal acts of discrimination can produce substantial experiences of discrimination, it raises the question: What can and should be done to mitigate discrimination?

What to Do?

Wrong answers first. The empirical research on the effectiveness of diversity and implicit bias trainings for reducing discrimination is generally weak and equivocal (e.g., Forscher et al., 2019; see reviews by (Mogiliski, Jussim, Wilson & Love, 2025; Paluck, Porat, Clark & Green, 2021). Paluck et al., (2021, p. 533), for example, concluded that "… much research effort is theoretically and empirically ill-suited to provide actionable, evidence-based recommendations for reducing prejudice." Mogiliski et al., (2025, p. 5) likewise concluded, "The real-world success of DEI programming is poorly documented…" The Discrimination Paradox may help explain why success of such programs is difficult to document – because of a floor effect. If acts of discrimination are as rare as indicated in the studies reviewed above, it is likely to be exceedingly difficult to *reduce behaviors with frequencies already near 0 to even less frequent levels* via such trainings. It also seems like a colossal waste of resources – money, time and effort -- if nearly all of those subjected to such trainings are already almost never (or never) engaging in discrimination.

Then, of course, there is the preferential selection form of affirmative action. If we know that some group is victimized by discrimination, why not compensate for it with anti-racist discrimination? In the U.S., discriminating on the basis of race in employment is illegal except under rare circumstances. As I am not a global legal scholar, I have purposely refrained from commenting on its legality outside the U.S.

Right answers? First, no one really knows how to eliminate discrimination. Second, there are ample reasons to believe humans have an evolved tendency to be suspicious of (at best) and hostile to (at worst) outgroups (e.g., Fox, 1992). This does not mean such suspicion or hostility cannot be overcome or unlearned, but it is probably an uphill battle, one that probably will never be completely won — or, put differently, *completely eradicating* prejudice and discrimination does not seem likely anytime soon.

Third, decades of research on stereotypes and prejudice might be usable to produce better ideas. For example, there is abundant evidence that, when perceivers have and attend to a great deal of relevant individuating information (i.e., detailed information about the qualifications and experience of a particular job or college applicant), they overwhelmingly use that information, rather than make judgments based on stereotypes of demographic categories (Kunda, 1996). Indeed, this was one of the other findings in the Airbnb study (Nødtvedt et al., 2021). In one condition, they provided participants with reviews of the apartment, held constant for those listed by Norwegians and Somalis. In this condition, there was no discrimination.

Thus, another strong contender for eliminating discrimination on the job or in admissions is to adopt practices that emphasize focusing on and evaluating merit (Abbot et al., 2023). Meta-analysis has also found that emphasis on merit more strongly predicts nondiscrimination than does emphasizing identity-blind or multicultural approaches (Leslie, Bono, Kim, & Beaver, 2020).

This answer – focusing on individuating information and merit -- is probably imperfect and will probably reduce but not fully eliminate all discrimination. Some people are outright bigots and would be unwilling to focus exclusively on merit. But it is probably the least bad answer we have so far, because it is likely to accomplish something without being illegal as are racial preferences in the U.S.

References

Abbot, D., Bikfalvi, A., Bleske-Rechek, A. L., Bodmer, W., Boghossian, P., Carvalho, C. M., ... & West, J. D. (2023). In defense of merit in science. *Journal of Controversial Ideas*, *3*(1).

Banaji, M. R., Fiske, S. T., & Massey, D. S. (2021). Systemic racism: individuals and interactions, institutions and society. *Cognitive research: principles and implications*, *6*(1), 82.

Bickel, P. J., Hammel, E. A., & O'Connell, J. W. (1975). Sex Bias in Graduate Admissions: Data from Berkeley: Measuring bias is harder than is usually assumed, and the evidence is sometimes contrary to expectation. *Science*, *187*(4175), 398-404.

Bonilla-Silva, E. (2015). The structure of racism in color-blind, "post-racial" America. *American Behavioral Scientist*, *59*(11), 1358-1376.

Braveman, P. A., Arkin, E., Proctor, D., Kauh, T., & Holm, N. (2022). Systemic and structural racism: Definitions, examples, health damages, and approaches to dismantling. *Health Affairs*, *41*(2), 171-178.

Campbell, M. R., & Brauer, M. (2021). Is discrimination widespread? Testing assumptions about bias on a university campus. *Journal of Experimental Psychology: General, 150*(4), 756–777. https://doi.org/10.1037/xge0000983

Ferguson, C. J., & Smith, S. (2024). Race, class, and criminal adjudication: Is the US criminal justice system as biased as is often assumed? A meta-analytic review. *Aggression and Violent Behavior*, *75*, 101905.

Forscher, P. S., Lai, C. K., Axt, J. R., Ebersole, C. R., Herman, M., Devine, P. G., & Nosek, B. A. (2019). A meta-analysis of procedures to change implicit measures. *Journal of Personality and Social Psychology*, *117*(3), 522.

Fox, R. (1992). Prejudice and the unfinished mind: A new look at an old failing. *Psychological Inquiry*, *3*(2), 137-152.

Gould, S. J. (1981). *Evolution as fact and theory.* The Unofficial Stephen Jay Gould Archive. https://wise.fau.edu/~tunick/courses/knowing/gould_fact-and-theory.html

Greenwald, A. G., Banaji, M. R., & Nosek, B. A. 2015. Statistically small effects of the Implicit Association Test can have societally large effects. *Journal of Personality and Social Psychology*, 108: 553-561.

Hardy III, J. H., Tey, K. S., Cyrus-Lai, W., Martell, R. F., Olstad, A., & Uhlmann, E. L. (2022). Bias in context: Small biases in hiring evaluations have big consequences. *Journal of Management*, *48*(3), 657-692.

Haslam, N. (2016). Concept creep: Psychology's expanding concepts of harm and pathology. *Psychological inquiry*, *27*(1), 1-17.

Jussim, L., Careem, A., Goldberg, Z., Honeycutt, N. & Stevens, S. T. (2025). IAT scores, racial gaps, and scientific gaps. In *The Cambridge handbook of implicit bias and racism,* pp. 374-411. (J. A. Krosnick, T. H. Stark and A. L. Scott, Eds). New York: Cambridge University Press.

Krylov, A.I. and Tanzman, J. (2023). Critical Social Justice subverts scientific publishing. *European Review* ***31***, 527 – 546.

Kunda, Z., & Thagard, P. (1996). Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological review*, *103*(2), 284.

Lee, L. (January 4, 2023). Who owns America's businesses: Minority business ownership by sector. United States Census Bureau. https://www.census.gov/library/stories/2023/01/who-owns-americas-businesses.html

Leonardelli, G. J., & Brewer, M. B. (2001). Minority and majority discrimination: When and why. *Journal of Experimental Social Psychology*, *37*(6), 468-485.

Leslie, L. M., Bono, J. E., Kim, Y. S., & Beaver, G. R. (2020). On melting pots and salad bowls: A meta-analysis of the effects of identity-blind and identity-conscious diversity ideologies. *Journal of Applied Psychology*, *105*(5), 453.

Martell, R. F., Lane, D. M., & Emrich, C. (1996). Male-female differences: A computer simulation. *American Psychologist, 51(2),* 157-158.

Mogiliski, J., Jussim, L., Wilson, A. & Love, B. (2025). Defining diversity, equity and inclusion (DEI) by the scientific (de)merits of its programming. *Theory and Society,* 1-14.

Mullen, B., Brown, R., & Smith, C. (1992). Ingroup bias as a function of salience, relevance, and status: An integration. *European Journal of Social Psychology*, *22*(2), 103-122.

Nødtvedt, K. B., Sjåstad, H., Skard, S. R., Thorbjørnsen, H., & Bavel, J. J. V. (2021). Racial bias in the sharing economy and the role of trust and self-congruence. *Journal of Experimental Psychology: Applied*, *27*(3), 508-528. https://doi.org/10.1037/xap0000355

Oswald, F. L., Mitchell, G., Blanton, H., Jaccard, J., & Tetlock, P. E. (2015). Using the IAT to predict ethnic and racial discrimination: small effect sizes of unknown societal significance. *Journal of Personality and Social Psychology, 108*(4), 562–571. https://doi.org/10.1037/pspa0000023

Paluck, E. L., Porat, R., Clark, C. S., & Green, D. P. (2021). Prejudice reduction: Progress and challenges. *Annual Review of Psychology*, *72*(1), 533-560.

Peyton, K., & Huber, G. A. (2021). Racial resentment, prejudice, and discrimination. *The Journal of Politics*, *83*(4), 1829-1836.

Quillian, L., Pager, D., Hexel, O., & Midtbøen, A. H. (2017). Meta-analysis of field experiments shows no change in racial discrimination in hiring over time. *Proceedings of the National Academy of Sciences*, *114*(41), 10870-10875.

Schaerer, M., Du Plessis, C., Nguyen, M. H. B., Van Aert, R. C., Tiokhin, L., Lakens, D., ... & Gender Audits Forecasting Collaboration. (2023). On the trajectory of discrimination: A meta-analysis and forecasting survey capturing 44 years of field experiments on gender and hiring decisions. *Organizational Behavior and Human Decision Processes*, *179*, 104280.

Syed, M. (2021). The logic of microaggressions assumes a racist society. *Perspectives on Psychological Science*, *16*(5), 926-931.

U.S. Census (July 1, 2024). Quick facts, United States. https://www.census.gov/quickfacts/fact/table/US/PST045224

Wikipedia (March 18, 2025). Institutional racism. https://en.wikipedia.org/wiki/Institutional_racism