

## Supporting Information for

### A framing effect in the judgment of discrimination

Christopher K. Hsee and Xilin Li

To whom correspondence may be addressed.

**Email:** [Christopher.Hsee@chicagobooth.edu](mailto:Christopher.Hsee@chicagobooth.edu)

**This PDF file includes:**

- Supplementary text
- Figure S1
- Legend for Dataset S1
- SI references

**Other supporting materials for this manuscript include the following:**

- Dataset S1

## Supporting Information Text

### Supplementary Materials Part 1: Detailed Methods and Additional Results

Some general notes: All studies adopted the same target sample size criterion—150 per cell (e.g., 450 in total if a study had three conditions). To show the generality of our findings, we used a variety of terms (“discrimination,” “prejudice,” and “favoritism”) when asking participants to judge discrimination. We report all measures of judged discrimination such that greater values represent stronger judged discrimination against Min members; if the original ratings were in the opposite direction, we reverse-coded them for consistency. The analyses are based on all responses we collected without any exclusions.

#### Study 1A

Study 1A (pre-registered at [https://aspredicted.org/9QK\\_MGZ](https://aspredicted.org/9QK_MGZ)) consisted of two between-subjects conditions: accept and reject. We used a nationally representative sample of US participants ( $N = 448$ ; 214 females,  $M_{\text{age}} = 42.79$ ) from Prime Panel. (We pre-registered a target sample size of 300 because Study 1A had two conditions, but Prime Panels recruited more participants than requested, probably because the sample had to be large enough to be nationally representative.) All participants read a hiring case:

Kris is the principal of a private school. The school recently recruited new teachers. Many people applied, of whom 20% were female and 80% were male. These applicants were equally qualified, and Kris knew that.

Kris decided which applicants to reject and which applicants to hire. In total, Kris rejected half the applicants and hired the other half.

In addition, participants in the [accept]{reject} condition read the following:

You check the composition of the applicants [hired]{rejected} by Kris and find the following information: 20% of the applicants [hired] {rejected} by Kris were female, and 80% of the applicants [hired] {rejected} by Kris were male.

All participants were then asked whether Kris discriminated against either female or male applicants, and they answered on a 9-point scale anchored by “strongly against female applicants” and “strongly against male applicants.”

### **Study 1B**

Study 1B consisted of an accept condition and a reject condition. All participants ( $N = 302$ ; 133 females,  $M_{\text{age}} = 26.14$ , recruited on Prolific) read a hiring case:

Drew is the HR director of a US company. The company recently recruited new employees. A total of 200 people applied, of whom, 20 were URMs and 180 were non-URMs. (“URMs” stand for underrepresented minorities, including African Americans/Blacks, Hispanics/Latinxs, Native Americans/Alaskan Natives, and Native Hawaiians/Other Pacific Islanders.)

The URM applicants were as qualified as the non-URM applicants, and Drew was aware of this. Because the number of positions was limited, half the applicants had to be rejected.

In addition, participants in the [accept] {reject} condition read:

Drew decided which applicants to [hire] {reject} and which to [reject] {hire}. In total, Drew [hired] {rejected} 100 applicants and [rejected] {hired} the rest. You check whom Drew [hired] {rejected}, and find that Drew [hired] {rejected} 10 URMs and 90 non-URMs.

All participants were asked whether Drew discriminated against either URMs or non-URMs, and they answered on a 9-point scale anchored by “strongly against URMs” and “strongly against non-URMs.”

On the next page, we tested for the equality expectation by asking participants to indicate what they thought would be a fair hiring decision. Specifically, we asked participants in the [accept] {reject} condition: “If you were to decide whom to [hire] {reject} and you want to be fair to both URMs and non-URMs, how many URMs and how many non-URMs would you [hire] {reject}?” Participants chose from 11 options, ranging from “100 URMs and 0 non-URMs” to “0 URMs and 100 non-URMs” in 10% increments.

### **Study 1C**

Study 1C consisted of an accept condition and a reject condition. All participants ( $N = 306$ ; 150 females,  $M_{\text{age}} = 40.11$ , recruited on CloudResearch) were first told:

In this study you will read a case and then evaluate the decision-maker in the case. To make sure you read the case carefully, we will ask you comprehension questions. If you answer all comprehension questions correctly, you will get a \$0.15 bonus. If you answer any of the comprehension questions incorrectly, you will lose the bonus.

Participants then read the following case:

A certain library had 60 librarians, of whom 50 were female and 10 were male. These librarians were equally qualified. Due to an economic recession, the library could keep only half the librarians and had to dismiss the other half. Kris—the head of the library—decided whom to keep and whom to dismiss.

Participants were asked to answer the following comprehension questions and were reminded that they would not get the bonus if they answered any of the questions incorrectly:

How many librarians did the library have?

- 60
- 40
- 20

How many of them were female vs. male?

- 50 female / 10 male
- 30 female / 30 male
- 10 female / 50 male

Were these librarians equally qualified?

- yes
- no

Who is Kris?

- one of the librarians
- the head of the library

Regardless of their answers, participants then read the rest of the case, which differed between the [accept] and {reject} conditions:

In total, Kris [kept] {dismissed} 30 librarians and [dismissed] {kept} the rest. Of the 30 librarians Kris [kept] {dismissed}, 25 were female and 5 were male.

Participants were asked whether Kris favored either female or male librarians, and they answered on a 9-point scale anchored by “favored female librarians” and “favored male librarians.”

On a separate page, to check whether participants had the quantitative ability to infer the information that had been provided in the other condition, we asked another

comprehension question and reminded participants that they would not get the bonus if they answered incorrectly:

Given that the library had 50 female librarians and 10 male librarians, and Kris [kept] {dismissed} 25 female librarians and 5 male librarians, then how many female librarians and how many male librarians did Kris [dismiss] {keep}?

- (a) Kris [dismissed] {kept} 25 female librarians and 5 male librarians
- (b) Kris [dismissed] {kept} 20 female librarians and 10 male librarians
- (c) Kris [dismissed] {kept} 15 female librarians and 15 male librarians
- (d) Kris [dismissed] {kept} 10 female librarians and 20 male librarians
- (e) Kris [dismissed] {kept} 5 female librarians and 25 male librarians
- (f) none of the above

Most participants (95.8%) answered all comprehension questions correctly.

## Study 1D

Participants ( $N = 300$ ; 170 females,  $M_{\text{age}} = 30.95$ , recruited on Credamo in China) were assigned to either an accept condition or a reject condition. All participants read the following case (originally in Chinese):

Laoli is the organizer of a nationwide variety show. Laoli recently posted advertisements nationwide to recruit participants. Many people applied, of whom 80% were Northerners and 20% were Southerners. These applicants were similarly qualified, and Laoli knew that. Due to the limited scope of the show, it was impossible to accept all applicants. Laoli decided which applicants to reject and which applicants to accept.

Participants in the [accept] {reject} condition read the following additional information:

Laoli eventually [accepted] {rejected} half the applicants. You check the composition of the [accepted] {rejected} applicants and find the following:

Among the applicants [accepted] {rejected} by Laoli to participate in this nationwide variety show, 80% were Northerners, and 20% were Southerners.

All participants were asked to evaluate Laoli's favoritism on a 9-point bipolar scale anchored by "strongly favored Northerners" and "strongly favored Southerners." We asked about favoritism rather than discrimination to test the generalizability of our findings, assuming that favoritism was the opposite of discrimination. For consistency, we report the results in terms of discrimination.

## **Study 2**

Study 2 adopted a 2 (framing: accept vs. reject) x 2 (objective discrimination: against Min vs. against Maj) between-subjects design. All participants ( $N = 606$ ; 318 females,  $M_{age} = 40.92$ , recruited on CloudResearch in the US) read a layoff case that began:

A certain company in the US had 200 employees, of whom, 70% were Northerners and 30% were Southerners. These employees were equally good and the company knew that.

Due to a recent economic downturn, the company laid off half of its employees and kept only the remaining half.

In the objectively-against-Min condition, participants in the [accept] {reject} condition further read:

You check the composition of the employees [kept] {laid off} by the company, and find the following: [90% of the kept employees were Northerners and 10% of the kept employees were Southerners.] {50% of the laid-off employees were Northerners and 50% of the laid-off employees were Southerners.}

In the objectively-against-Maj condition, participants in the [accept] {reject} condition read the following instead:

You check the composition of the employees [kept] {laid off} by the company, and find the following: [50% of the kept employees were Northerners and 50% of the kept employees were Southerners.] {90% of the laid-off employees were Northerners and 10% of the laid-off employees were Southerners.}

All participants then evaluated the company's discrimination on an 11-point bipolar scale anchored by "Strongly against Northerners" and "Strongly against Southerners."

Unlike Studies 1A–1D, this study involved objective discrimination. The magnitude of the objective discrimination was held constant, but the sign varied by condition: in the objectively-against-Min condition,  $\text{Min}\%_R - \text{Min}\%_A = 50\% - 10\% = +40\%$ ; in the objectively-against-Maj condition,  $\text{Min}\%_R - \text{Min}\%_A = 10\% - 50\% = -40\%$ .

### **Study 3**

Participants ( $N = 304$ ; 137 females,  $M_{\text{age}} = 42.99$ , recruited on CloudResearch in the US) were randomly assigned to an accept condition or a reject condition. First, all participants were asked whether they had supported Biden, Trump, or neither in the 2020 presidential election. Most of our participants (92.1%) were either Trump supporters (32.6%) or Biden supporters (59.5%).

Then, participants received the following instructions:

In this study, we want you to write a short essay (up to 100 words) to explain why you supported the presidential candidate that you did.

We encourage you to write the essay as well as possible. After you finish your essay, one of our research assistants (RAs) will read it, and will either "pass" you

or “fail” you. If the RA (research assistant) passes you, you will get the promised payment without having to do any additional tasks. If the RA fails you, you will have to do another task before you can get the payment, and that task may take you a long time to finish.

Before you proceed, we will give you some information about the RA who is assigned to read your essay. This is to give you an idea of whether the RA is biased against either Biden supporters or Trump supporters. If you think the RA is biased, you have the right to ask us to pick another RA to read your essay. You have only one chance for us to change the RA, and if you want us to change the RA, we will *randomly* pick another RA (from our RA pool) for you.

Participants in the [accept] {reject} condition received additional information about the RA who was assigned to read their essay:

Below is the information about the RA who is currently assigned to read your essay.

So far, this RA has been assigned to read the essays of 48 workers, of whom, 36 were Biden supporters and 12 were Trump supporters.

In total, the RA [passed] {failed} half of the 48 workers and [failed] {passed} the remaining ones. Of the workers the RA [passed] {failed}, 18 were Biden supporters and 6 were Trump supporters.

All participants completed two dependent measures. First, participants were asked whether RA had discriminated against either the Trump supporters or the Biden supporters, and they answered on a 7-point scale anchored by “strongly against Biden supporters” and “strongly against Trump supporters.” Second, participants were asked whether they wanted to keep this RA or wanted us to pick another RA to read their essay, and they made a binary choice. Finally, participants wrote their essays; the essays were read by either the default RA or a different RA, depending on their choice. All participants “passed” and received a bonus.

## Study 4

Study 4 consisted of three framing conditions: accept, reject, and dual. All participants ( $N = 448$ ; 189 females,  $M_{age} = 27.57$ , recruited on Prolific) first read:

In this study, we want you to judge whether a certain rich country favors any ethnic groups in its decisions to deport or keep refugees. You have a chance to receive a bonus if your judgment is accurate; specifically, we will randomly pick 10 participants among those whose judgements are most accurate, and give each of them a \$1 bonus.

Then participants read a refugee case:

A total of 4000 civilians from a poor war-torn nation that consists of two ethnic groups—Ayronians and Byronians—have entered a rich neighboring country as refugees. Of those refugees, 1000 are Ayronians and 3000 are Byronians. The rich country is willing to accept some of the refugees and decides to deport the rest of them.

Participants in the [accept] {reject} <dual > condition learned the following additional information:

[The rich country eventually keeps 50% of the refugees and deports the rest. Of the refugees the rich country keeps, 25% are Ayronians and 75% are Byronians.]

{The rich country eventually deports 50% of the refugees and keeps the rest. Of the refugees the rich country deports, 25% are Ayronians and 75% are Byronians.}

<The rich country eventually keeps 50% of the refugees and deports 50% of the refugees. Of the refugees the rich country keeps, 25% are Ayronians and 75% are Byronians. Of the refugees the rich country deports, 25% are Ayronians and 75% are Byronians.>

All participants were asked to evaluate the rich country's discrimination on a 7-point scale anchored by "strongly favors Ayronians" and "strongly favors Byronians." Out of curiosity, we ended with six Cognitive Reflection Test (CRT) questions, which are

designed to measure the participant's tendency to use cognitive reflection instead of intuitive thinking (1). We found a negative correlation between the number of correct answers on the CRT questions and the absolute difference between judged discrimination and the midpoint of the judgment scale (i.e., the correct judgment: no objective discrimination),  $r = -.308, p < 0.001$ . The result indicates that participants with a weaker tendency to use cognitive reflection (i.e., a stronger tendency to rely on intuition) were more susceptible to the framing manipulation.

As promised, after we finished conducting the study, we awarded a \$1 bonus to 10 participants, randomly picked from those whose judgment was most accurate (which we define as rating the rich country as non-discriminatory, since  $\text{Min}\%_A = \text{Min}\%_R$ ).

## Study 5

Study 5 followed a 2 (framing: accept vs. reject) x 2 (control vs. flipside thinking) between-subjects design. Participants ( $N = 604$ ; 267 females,  $M_{\text{age}} = 26.20$ , recruited on Prolific) read one of four versions of a case corresponding to the 2 x 2 design. The case began:

You live in a city that consists of two sides—east side and west side. There is only one college in your city, and it is located at the city center. The college recently recruited students. A total of 400 people applied; among them, 100 were from the east side and 300 were from the west side. On average, the applicants from the two sides were equally good, and the college knew that. But due to its limited capacity, the college had to reject many applicants.

In addition, participants in the [accept] {reject} condition read:

In total, the college [admitted] {rejected} 200 applicants and [rejected] {admitted} the rest. You check the composition of the [admitted] {rejected}

students and find the following: Among the [admitted] {rejected} applicants, 50 were from the east side and 150 were from the west side.

Then, participants in the control condition received the following prompt:

Take a moment to think about the information in the case. After you think, answer the following question: Was the college prejudiced against applicants from either the east side or the west side?

Meanwhile, participants in the flipside-thinking condition received the following prompt:

Take a moment to think about the following: Given that 100 people from the east side and 300 people from the west side applied, and the college [admitted] {rejected} 50 from the east side and 150 from the west side, then how many from the east side and how many from the west side did the college [reject] {admit}? After you figure out these numbers, answer the following question: Was the college prejudiced against applicants from either the east side or the west side?

All participants answered the last question on a 9-point scale anchored by “strongly against applicants from the east side” and “strongly against applicants from the west side.”

## **Study 6**

Study 6 adopted a 2 (framing: accept vs. reject) x 2 (composition: unequal vs. equal) between-subjects design. Participants ( $N = 618$ ; 318 females,  $M_{age} = 39.45$ , recruited on CloudResearch in the US) read one of four versions of a case corresponding to the 2 x 2 design.

In the unequal-composition condition, participants in the [accept] {reject} condition read:

Recently, a US political leader named Kris announced that they wanted to form a bipartisan committee to discuss important policy issues. Many people applied to serve on the committee, of whom 70% were Democrats and 30% were Republicans. These applicants were similarly qualified and Kris knew that.

To keep the size of the bipartisan committee manageable, Kris rejected half the applicants and accepted the other half.

You check the composition of the applicants [accepted] {rejected} by Kris to serve on this bipartisan committee and find the following information: 70% of the [accepted] {rejected} applicants were Democrats, and 30% of the [accepted] {rejected} applicants were Republicans.

The text in the equal-composition conditions was identical except that we replaced 70% and 30% with 50% and 50%. All participants were asked to evaluate Kris's favoritism on a 9-point bipolar scale anchored by "strongly favors Democrats" and "strongly favors Republicans."

## **Study 7**

Study 7 adopted a 2 (framing: accept vs. reject) x 2 (constituent groups: expectedly equal vs. expectedly unequal) between-subjects design. Participants ( $N = 608$ ; 310 females,  $M_{\text{age}} = 40.50$ , recruited on CloudResearch in the US) read one of four versions of a committee formation case.

In the expectedly-equal conditions, participants in the [accept] {reject} condition read:

Kris is a political leader in the US. Recently, Kris announced that they wanted to form an important nationwide committee to discuss policy issues relevant to every American. One hundred people applied to serve on the committee, of whom, 36 were from the Red States and 64 were from the Blue States.

Kris personally decided which applicants to [accept] {reject} to serve on this important nationwide committee and which applicants to [reject] {accept}. In

total, Kris [accepted] {rejected} 50 applicants and [rejected] {accepted} the remaining ones. You check the composition of the 50 applicants [accepted] {rejected} by Kris to serve on this important nationwide committee and find that 18 of the [accepted] {rejected} applicants were from the Red States and 32 of the [accepted] {rejected} applicants were from the Blue States.

The text in the expectedly-unequal versions was identical except that “the Red States” was replaced with “Rhode Island,” and “the Blue States” was replaced with “the other 49 US states.”

All participants were asked to evaluate Kris’s favoritism on a 9-point bipolar scale, anchored either by “strongly favored people from the Red States” and “strongly favored people from the Blue States” (expectedly-equal condition) or by “strongly favored people from Rhode Island” and “strongly favored people from the other 49 US states” (expectedly-unequal condition).

We conducted a pretest to check the effectiveness of the expectation manipulation. We assigned a separate group of participants ( $N = 202$ ) to the same 2 x 2 design and provided the same scenario from the main study, but we did not provide the actual composition of the accepted or rejected candidates. Instead, we asked those in the [accept] {reject} condition: “Ideally, what should be the composition of the applicants [accepted] {rejected} to serve on this nationwide committee?” The participants chose from 21 options. In the expectedly-equal conditions, the options ranged from “0% from Red States and 100% from Blue States” to “100% from Red States and 0% from Blue States” in 5% increments. In the expectedly-unequal conditions, “the Red States” and “the Blue States” were replaced with “Rhode Island” and “the other 49 US states,” respectively. In the expectedly-equal conditions, the expected proportion of Min members in both the accept condition (45.3%; median = 50%) and the reject condition

(45.0%; median = 50%) were significantly closer to equality than the actual proportion of Min members in the main study scenario (i.e., 36%),  $t(51) = 4.85, p < .001$ , and  $t(51) = 4.57, p < .001$ . By contrast, in the expectedly-unequal conditions, the expected proportion of Min members in both the accept condition (13.6%; median = 5%) and the reject condition (25.2%; median = 5%) were significantly *farther* from equality than the actual proportion of Min members in the main study (i.e., 36%),  $t(48) = 10.93, p < .001$ , and  $t(48) = 2.36, p = .023$ . The pretest results show that our expectation manipulation was effective.

## Supplementary Materials Part 2: Single-Paper Meta-Analysis

We conducted a single-paper meta-analysis (SPM) of the framing effect (2). The analysis included the discrimination judgment data from all conditions of all studies in which our theory predicts a framing effect (all conditions of Studies 1A, 1B, 1C, 1D, 2, and 3; the control conditions of Studies 4, 5, 6, and 7). Since seven of the 10 studies used a 9-point scale, we converted the responses of the studies that used a different scale (an 11-point scale in Study 2; a 7-point scale in Studies 3 and 4 to a 9-point scale prior to the analysis.

The analysis yielded a robust effect: the difference in judged discrimination between the accept and reject conditions (i.e., the framing effect) is 1.14 on the 9-point scale ( $SE = 0.14$ ,  $p < .001$ ). See Figure S1 for a forest plot of the results.

### Supplementary Materials Part 3: A General Theory of Discrimination Judgment

In the main text, we presented our theory as:

$$JD_A = -(\text{Min}\%_A - 50\%) + k \quad \text{Equation 1A}$$

$$JD_R = (\text{Min}\%_R - 50\%) + k \quad \text{Equation 1R}$$

This theory can be generalized to:

$$JD_A = -(\text{Min}\%_A - E\text{Min}\%_A) + k \quad \text{Equation 1A(G)}$$

$$JD_R = (\text{Min}\%_R - E\text{Min}\%_R) + k \quad \text{Equation 1R(G)}$$

where  $E\text{Min}\%_A$  and  $E\text{Min}\%_R$  are the expected proportions of Min members in the accept and reject conditions, respectively.

Equations 1A(G) and 1R(G) are identical to Equations 1A and 1R except that we replaced 50% with  $E\text{Min}\%_A$  and  $E\text{Min}\%_R$ . In other words, the original theory is a special case of the general model in which observers expect equality in the compositions of the accepted and rejected groups (i.e.,  $E\text{Min}\%_A$  and  $E\text{Min}\%_R = 50\%$ ).

Combining Equations 1A(G) and 1R(G) yields:

$$JD_A - JD_R = (E\text{Min}\%_A + E\text{Min}\%_R) - (\text{Min}\%_A + \text{Min}\%_R) \quad \text{Equation 2(G)}$$

As explained in the main text, as long as the decision-maker accepts half the candidates and rejects half the candidates (or the decision-maker shows no objective discrimination), the following must be true:

$$(\text{Min}\%_A + \text{Min}\%_R)/2 = \text{Min}\%_C \quad \text{Equation 3}$$

Combining Equations 2(G) and 3 yields:

$$J_{DA} - J_{DR} = (E\text{Min}\%_A + E\text{Min}\%_R) - 2 \times \text{Min}\%_C \quad \text{Equation 4(G)}$$


which means:

$$J_{DA} > J_{DR} \text{ if } \text{Min}\%_C < (E\text{Min}\%_A + E\text{Min}\%_R)/2 \quad \text{Inequality 1(G)}$$

This inequality is a general version of Inequality 1, and it predicts a general framing effect: a decision-maker will be judged as more discriminatory against Min members in the accept condition than in the reject condition as long as the proportion of Min members in the candidate pool is smaller than the *average* of the expected proportions of Min members in the accept and reject conditions. This general framing effect occurs regardless of whether there is objective discrimination and regardless of whether the expected compositions are the same or different in the accept and reject conditions.

The general model not only explains the typical framing effect but also predicts and explains other effects. For example, the model predicts that if the proportion of Min members in the candidate pool is *greater than the average* of the expected proportions of Min members in the accept and reject conditions, the framing effect will reverse. This explains why we observed the opposite framing effect in the expectedly-unequal conditions of Study 7: In those conditions,  $\text{Min}\%_C = 36\%$ ,  $\text{EMin}\%_A = 13.6\%$ , and  $\text{EMin}\%_R = 25.2\%$ , so  $\text{Min}\%_C$  exceeds the average of  $\text{EMin}\%_A + \text{EMin}\%_R$ .

The general theory also can predict what will happen if one group is less qualified than the other. Suppose that 20 rich students and 80 poor students applied for financial aid, and the school granted half the applicants financial aid and denied the other half. Here, most observers likely would believe that the Min members (the rich applicants) are less qualified to receive financial aid than the Maj members (the poor applicants). But, according to our general theory, observers will still exhibit a framing effect: the school will be judged as more discriminatory against the rich applicants in the accept condition than in the reject condition, regardless of the exact composition of the accepted and rejected applicants. The reason is that, in this case,  $\text{EMin}\%_A$  is likely very small (close to 0%), and  $\text{EMin}\%_R$  is likely very large (close to 100%); that is, people likely expect 0% of the rich applicants to receive financial aid and 100% of the rich applicants to be denied financial aid. Thus, the average of  $\text{EMin}\%_A$  and  $\text{EMin}\%_R$  is around 50%, which is larger than  $\text{Min}\%_C$  (20%). According to Inequality 1(G),  $\text{JD}_A$  will be greater than  $\text{JD}_R$ , namely, there will be a framing effect.



**Figure S1. Forest Plot of Single-Paper Meta-analysis** Effect estimates are given by the squares and vertical bars for SPM estimates; 50% and 95% intervals are given by the thick and thin lines, respectively. The average sample size per condition in each study is given by the size of the square.

## **SI References**

1. S. Frederick, Cognitive reflection and decision making, *J. Econ. Perspect.* 19, 25-42.
2. B. McShane, U. Böckenholt, Single-paper meta-analysis: Benefits for study summary, theory testing, and replicability. *J. Consum. Res.* 43, 1048-1063.

**Dataset S1.** Data can be found at  
[https://osf.io/k8xsb/?view\\_only=2ef058bb88d84f0b9141c5145161488f](https://osf.io/k8xsb/?view_only=2ef058bb88d84f0b9141c5145161488f).