

Statistical Case Study: Job Promotion Discrimination

Statistical Analysis and Equity Implications

Ryan Beavers

October 28, 2025

Contents

1	Executive Summary	3
2	Case & Data	3
2.1	Data entry	3
3	Methods (Classical)	4
4	Results & Interpretation	5
4.1	1) Contribution vs. Promotion (2×2)	5
4.1.1	Visuals	6
4.2	2) Contribution vs. Job Condition Change (3×3)	6
4.2.1	Heatmap	7
4.3	Candidate Ranks (Context, not dispositive)	7
5	Conclusions & Implications	8
6	Recommendations for Practice	8
7	Appendix A - Transparency	10
8	Appendix B - Reproducibility	11
9	Appendix C - Code	12
9.1	Setup	12
9.2	Data entry	12
9.3	Methods (Classical)	13

9.3.1	1) Contribution vs. Promotion (2×2)	13
9.3.2	2) Contribution vs. Job Condition Change (3×3)	13
9.3.3	Candidate Ranks	14

1 Executive Summary

What's at stake?

Promotion isn't just about titles — it determines who gains access to higher pay, leadership tracks, and meaningful work. When the criteria behind those decisions aren't clearly defined, even the language of “performance” can become a retroactive story we tell to justify outcomes. This case asks a simple but serious question: Do the available data reflect a fair process, or do gaps in transparency leave room for unequal treatment?

What we analyzed?

We reviewed two key comparisons from the case:

Contribution vs. Promotion

Contribution vs. Change in Job Conditions

Because the data are small and include zeros, we used Fisher's Exact Test as the primary method, with Chi-square and effect sizes (Phi and Cramer's V) providing context. Candidate rankings (A–J) were considered as background, not as formal evidence.

What the data suggest?

Contributors were promoted at far higher rates than non-contributors ($p < .05$; large effect).

Contributors were also more likely to receive positive changes in their job conditions, while non-contributors clustered in negative outcomes ($p < .01$; large effect).

Why caution is necessary?

At first glance, this pattern supports the idea that performance drives decisions. But without written standards or a transparent process, “contribution” may reflect who had access to high-visibility work — not who worked the hardest. In many organizations, opportunity itself is unevenly distributed. The real question is not only who contributed, but who was given the chance to contribute.

Bottom line.

The results align with a performance narrative, but the process behind the numbers remains unclear. In systems where criteria are unwritten or inconsistently applied, fair outcomes cannot be assumed — even when the statistics appear clean. True fairness depends on transparency, not just on outcomes.

2 Case & Data

We use the case's two cross-tabulations and a ranking table (A–J; D is the plaintiff). We treat ranks as qualitative context — informative but not sufficient to adjudicate fairness on their own.

Source (assignment instructions): :contentReferenceoaicite:0

2.1 Data entry

Promotion

##	Contribution	Promoted	NotPromoted
##	Contributed	6	0
##	NotContributed	1	3

Table 1: Candidate Ranking (1 = highest; D = plaintiff)

Candidate	Rank
A	1
B	2
C	3
D	4
E	5
F	6
G	7
H	8
I	8
J	8

##	Contribution	JobChange		
##	Contribution	Positive	Negative	Unknown
##	Contributed	4	0	2
##	DidNotContribute	0	7	0
##	Unknown	1	9	2

3 Methods (Classical)

To evaluate whether promotion outcomes were meaningfully associated with candidate contribution, we relied on classical statistical tests suitable for small datasets. Because several cells contained very low or zero counts, we used Fisher’s Exact Test as our primary method. Unlike Chi-square, Fisher’s test does not rely on large-sample assumptions and is better suited for drawing cautious conclusions from limited data.

We also calculated effect sizes — specifically Phi for the 2×2 table and Cramer’s V for the 3×3 table — not to reduce the case to a single number, but to understand the strength of the observed relationships. Instead of treating thresholds (0.1 = small, 0.3 = medium, 0.5+ = large) as rigid cutoffs, we interpreted them in context: a strong statistical association may reflect consistent standards, or it may reflect consistently applied bias. The statistics tell us how strong the pattern is — not whether it is fair.

Chi-square statistics were included only as supportive context and interpreted with caution due to the dataset’s size. Candidate rankings were treated qualitatively, acknowledging their relevance to perceived merit but recognizing that rank alone cannot confirm procedural fairness.

4 Results & Interpretation

4.1 1) Contribution vs. Promotion (2×2)

The data show a striking gap: contributors were promoted at a far higher rate than non-contributors, and the statistical test confirms this difference is unlikely to be due to chance (Fisher’s Exact, $p < .05$, large effect size). On the surface, this appears to validate a performance-based promotion system.

But fairness in promotion is not defined solely by outcomes — it is defined by process. The key question is whether everyone had an equal opportunity to contribute. If “contribution” is informally defined or tied to visibility, access, or managerial favor, then it may reflect opportunity rather than effort. In such systems, those already connected to decision-makers are more likely to be noticed, supported, and ultimately promoted.


When performance standards are unwritten or inconsistently applied, statistics can become a veneer — a way to justify outcomes rather than to explain them. Without transparent criteria, contribution risks becoming a story told after the fact, not a standard fairly applied from the start.

##		Promotion	
## Contribution		Promoted	NotPromoted
##	Contributed	6	0
##	NotContributed	1	3

Table 2: Contribution vs Promotion — Simple Summary

Measure	Value
Promotion rate: Contributed	100.0%
Promotion rate: Not Contributed	25.0%
Fisher’s Exact p-value	0.033
Phi (effect size)	0.802

4.1.1 Visuals



4.2 2) Contribution vs. Job Condition Change (3×3)

The results show a strong association between contribution status and changes in job conditions. Contributors were far more likely to receive positive changes, while non-contributors clustered almost entirely in negative outcomes — a pattern confirmed by Fisher’s Exact Test ($p < .01$) and a large effect size (Cramer’s V).

At face value, this seems to support a performance-driven system. But it raises a deeper structural question: Does contribution lead to better treatment, or does better treatment create the conditions for recognized contribution? If access to visible projects, mentorship, or leadership assignments is unevenly distributed, contribution may not be the cause of advancement — it may be the result of prior opportunity.


When opportunity flows through informal networks rather than explicit policy, the system risks becoming self-reinforcing. Those already recognized continue to gain visibility, which generates more proof of “contribution,” which then justifies further rewards. That isn’t a neutral meritocracy — it’s a cycle. Without written criteria and transparent access to high-impact roles, what appears to be performance can, in practice, be path dependence.

## Contribution	JobChange		
	Positive	Negative	Unknown
## Contributed	4	0	2
## DidNotContribute	0	7	0
## Unknown	1	9	2

Table 3: Contribution vs Job Condition — Simple Summary

Measure	Value
Fisher's Exact p-value	0.001
Cramer's V (effect size)	0.574

4.2.1 Heatmap



4.3 Candidate Ranks (Context, not dispositive)

Candidate D, the plaintiff, ranked 4th overall — a position that signals competitiveness rather than clear exclusion. Rank alone cannot resolve questions of fairness, but it challenges the idea that promotion was an obvious or automatic outcome based purely on merit. Without a documented rubric for promotion, it remains unclear whether D failed to meet transparent standards, or whether those standards were never explicitly communicated in the first place.

Table 4: Candidate Ranks (lower is better)

Candidate	Rank
A	1
B	2
C	3
D	4
E	5
F	6
G	7
H	8
I	8
J	8

##	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
##	1.00	3.25	5.50	5.20	7.75	8.00

5 Conclusions & Implications

The statistical evidence shows a strong association between contribution and promotion, suggesting that performance may play a central role in advancement. If contribution is clearly defined, consistently measured, and equally attainable, this alignment can reflect a fair and principled system.

However, the data alone cannot confirm whether access to contribute was equally available to all candidates. In many organizations, visibility, high-impact assignments, and sponsorship are distributed through informal networks rather than open processes. Under those conditions, “contribution” may reflect access rather than ability — and the language of merit can quietly reproduce structural bias.

This creates the risk of a self-reinforcing cycle: those already recognized are granted opportunities that generate further recognition, while others — equally capable — remain outside the evaluative frame. Without transparency, performance becomes a moving target, and fairness becomes a matter of perception rather than policy.

True equity in promotion requires more than consistent outcomes; it requires consistent procedures. Written criteria, clear communication, and documented decision-making are not bureaucratic burdens — they are safeguards against informal favoritism and post-hoc justification. A promotion system that cannot explain itself cannot defend itself.

6 Recommendations for Practice

- Clarify standards. Define what counts as contribution and publish the criteria before decisions are made.
- Track opportunity, not just outcome. Monitor who receives visible assignments, mentorship, and decision-making access.

- Document reasoning. Require brief written rationales for promotion decisions tied directly to stated criteria.
- Review for impact. Periodically examine promotion data by role, department, or demographic to identify patterns of exclusion.
- Rotate pathways to visibility. Ensure that all candidates, not just the well-connected, have access to roles that make contribution possible.

7 Appendix A - Transparency

```
# Expected counts (for transparency; small cells justify Fisher)  
chi_promo$expected
```

```
##               Promotion  
## Contribution    Promoted NotPromoted  
##   Contributed         4.2         1.8  
##   NotContributed       2.8         1.2
```

```
chi_job$expected
```

```
##               JobChange  
## Contribution    Positive Negative Unknown  
##   Contributed         1.2      3.84   0.96  
##   DidNotContribute     1.4      4.48   1.12  
##   Unknown             2.4      7.68   1.92
```

8 Appendix B - Reproducibility

```
sessionInfo()
```

```
## R version 4.5.1 (2025-06-13)
## Platform: aarch64-apple-darwin20
## Running under: macOS Sequoia 15.6
##
## Matrix products: default
## BLAS:   /Library/Frameworks/R.framework/Versions/4.5-arm64/Resources/lib/libRblas.0.dylib
## LAPACK: /Library/Frameworks/R.framework/Versions/4.5-arm64/Resources/lib/libRlapack.dylib;
##
## locale:
## [1] en_US.UTF-8/en_US.UTF-8/en_US.UTF-8/C/en_US.UTF-8/en_US.UTF-8
##
## time zone: America/Chicago
## tzcode source: internal
##
## attached base packages:
## [1] stats      graphics  grDevices  utils      datasets  methods   base
##
## other attached packages:
## [1] broom_1.0.10  knitr_1.50    ggplot2_4.0.0 dplyr_1.1.4
##
## loaded via a namespace (and not attached):
## [1] vctrs_0.6.5      cli_3.6.5      rlang_1.1.6     xfun_0.53
## [5] purrr_1.1.0      generics_0.1.4 S7_0.2.0        labeling_0.4.3
## [9] glue_1.8.0       backports_1.5.0 htmltools_0.5.8.1 scales_1.4.0
## [13] rmarkdown_2.30   grid_4.5.1     evaluate_1.0.5  tibble_3.3.0
## [17] fastmap_1.2.0    yaml_2.3.10    lifecycle_1.0.4 compiler_4.5.1
## [21] RColorBrewer_1.1-3 pkgconfig_2.0.3 tidyr_1.3.1     rstudioapi_0.17.1
## [25] farver_2.1.2     digest_0.6.37  R6_2.6.1        tidyselect_1.2.1
## [29] pillar_1.11.1    magrittr_2.0.4 withr_3.0.2     tools_4.5.1
## [33] gtable_0.3.6
```

9 Appendix C - Code

9.1 Setup

```
knitr::opts_chunk$set(message = FALSE, warning = FALSE, fig.width = 6, fig.height = 4)
options(digits = 3)
library(dplyr)
library(ggplot2)
library(knitr)
library(broom)      # for tidy() summaries
theme_set(theme_minimal(base_size = 12))
```

9.2 Data entry

```
# Table 1: Contribution by Promotion (2x2)
#           Promoted NotPromoted
# Contributed      6         0
# NotContributed   1         3
tab_promo <- matrix(c(6,0,
                      1,3),
                    nrow = 2, byrow = TRUE,
                    dimnames = list(Contribution = c("Contributed","NotContributed"),
                                    Promotion      = c("Promoted","NotPromoted"))))

tab_promo

# Candidate ranking (context)
ranks <- data.frame(
  Candidate = c("A","B","C","D","E","F","G","H","I","J"),
  Rank = c(1,2,3,4,5,6,7,8,8,8)
)
kable(ranks, caption = "Candidate Ranking (1 = highest; D = plaintiff)")

# Table 2: Contribution by Job Condition Change (3x3)
#           Positive Negative Unknown
# Contributed      4         0         2
# DidNotContribute 0         7         0
# Unknown          1         9         2
tab_job <- matrix(c(4,0,2,
                    0,7,0,
                    1,9,2),
                  nrow = 3, byrow = TRUE,
                  dimnames = list(Contribution = c("Contributed","DidNotContribute","Unknown"),
                                    JobChange    = c("Positive","Negative","Unknown"))))

tab_job
```

9.3 Methods (Classical)

9.3.1 1) Contribution vs. Promotion (2×2)

```
tab_promo
n_promo <- sum(tab_promo)

# Tests

fisher_promo <- fisher.test(tab_promo)
chi_promo <- suppressWarnings(chisq.test(tab_promo, correct = FALSE))
phi_promo <- as.numeric(sqrt(chi_promo$statistic / n_promo))

# Human-friendly summary table

promo_rates <- prop.table(tab_promo, 1)[, "Promoted", drop = FALSE]
promo_summary <- tibble::tibble(
  Measure = c("Promotion rate: Contributed", "Promotion rate: Not Contributed",
    "Fisher's Exact p-value", "Phi (effect size)"),
  Value = c(sprintf("%.1f%%", 100 * promo_rates["Contributed", 1]),
    sprintf("%.1f%%", 100 * promo_rates["NotContributed", 1]),
    formatC(fisher_promo$p.value, format = "f", digits = 3),
    round(phi_promo, 3))
)
kable(promo_summary, caption = "Contribution vs Promotion - Simple Summary")

df_promo <- as.data.frame(as.table(tab_promo))
names(df_promo) <- c("Contribution", "Promotion", "Freq")

df_promo_rates <- df_promo %>%
  group_by(Contribution) %>%
  mutate(Prop = Freq / sum(Freq)) %>%
  filter(Promotion == "Promoted") %>%
  ungroup()

ggplot(df_promo_rates, aes(x = Contribution, y = Prop)) +
  geom_col() +
  geom_text(aes(label = sprintf("%.1f%%", 100*Prop)), vjust = -0.25) +
  scale_y_continuous(labels = function(x) sprintf("%d%", x*100), limits = c(0,1)) +
  labs(title = "Promotion Rate by Contribution", x = NULL, y = "Promoted (percent)")
```

9.3.2 2) Contribution vs. Job Condition Change (3×3)

```
tab_job
n_job <- sum(tab_job)
```

```

# Tests

fisher_job <- fisher.test(tab_job)
chi_job <- suppressWarnings(chisq.test(tab_job, correct = FALSE))
r <- nrow(tab_job); c <- ncol(tab_job)
cramer_v <- as.numeric(sqrt(chi_job$statistic / (n_job * min(r - 1, c - 1))))

job_summary <- tibble::tibble(
  Measure = c("Fisher's Exact p-value", "Cramer's V (effect size)",
  Value = c(formatC(fisher_job$p.value, format = "f", digits = 3),
  round(cramer_v, 3))
)
kable(job_summary, caption = "Contribution vs Job Condition - Simple Summary")

df_job <- as.data.frame(as.table(tab_job))
names(df_job) <- c("Contribution", "JobChange", "Freq")

ggplot(df_job, aes(x = JobChange, y = Contribution, fill = Freq)) +
  geom_tile(color = "white") +
  geom_text(aes(label = Freq), fontface = "bold") +
  scale_fill_gradient(low = "gray90", high = "gray20") +
  labs(title = "Job Condition Change by Contribution", x = NULL, y = NULL, fill = "Count") +
  theme(legend.position = "right")

```

9.3.3 Candidate Ranks

```

kable(ranks, caption = "Candidate Ranks (lower is better)")
summary(ranks$Rank)

```