

FairPost.

An AI bias checker that helps companies write job descriptions for the best candidate — not the most credentialed one.

Fix the **problem** — not help people *survive* it.

Most job descriptions are written to impress internally, not to help real people understand the role. Inflated requirements and exclusionary language push qualified candidates to self-select out — and companies shrink their talent pool before the first application arrives.

What's broken

- Gender-coded language and inflated requirements screen people out before they apply
- Recruiters lack the time, tools, or training to audit postings for bias
- No step forces teams to separate must-haves from nice-to-haves
- Pay and schedule details are routinely omitted, eroding candidate trust
- Contradictions go undetected — *"entry-level role, 5+ years required"*

WHO'S HURT

Early-career candidates

Career changers with transferable skills

Women & underrepresented groups

Immigrants & non-native speakers

The company — a shrunken talent pool

A role posted as "Junior Project Coordinator."

THE POSTING ASKED FOR

3-5 years of experience

PMP certification

Advanced Excel

WHAT THE JOB ACTUALLY WAS

Scheduling meetings

Tracking tasks

Updating status reports

A strong candidate read it, assumed they weren't qualified, and didn't apply. The company never knew why.

PRIMARY

The Recruiter

Internal recruiters, HR coordinators & hiring managers

GOAL

Post a role fast that attracts qualified applicants without legal or reputational risk

PAIN

No feedback loop tells them when their language is exclusionary

NEEDS

Fast suggestions to accept or reject — not a lecture

SECONDARY

The DEI / HR Leader

HR directors, DEI program managers & People Ops leads

GOAL

Standardize fair posting practices and demonstrate measurable progress

PAIN

No data on which postings deter diverse applicants, or why

NEEDS

Dashboards, team analytics & exportable reports for leadership

"Every hiring manager wants to add more requirements. By the time it's posted, it sounds far more senior than the actual role — and strong candidates don't apply."

— Recruiter, former colleague

-
- 01 AI Rewriter** Rewrites the posting in clear, plain language — less jargon, tighter responsibilities, legible to a wider audience

 - 02 Requirement Splitter** Forces every requirement into must-have vs. nice-to-have, so teams see when they're over-filtering talent

 - 03 Bias Detector** Flags gender-coded, elitist & culture-fit language — with plain-language reasons and suggested alternatives

 - 04 Transparency Prompt** Catches missing pay, schedule type, or logistics and prompts the recruiter to add them before publishing

 - 05 Fairness Score** A single 0-100 signal across language, requirements, transparency & readability

87 /100

OVERALL FAIRNESS

A fast, scannable signal of posting quality — consistent across teams and over time.



A high score reflects the quality of the job description only — never a guarantee of a fair hiring process.

Identified, but deliberately excluded to keep the MVP tight and shippable.

× **Candidate "Apply Fit" tool**

Different user surface and business model

× **Multilingual translation**

Complex; core bias detection comes first

× **Anonymous candidate feedback**

Needs a two-sided network to be useful

× **ATS integrations**

High technical lift — better suited to v2

× **A/B testing framework**

Requires a sufficient user base first

× **Team approval workflows**

Enterprise feature, not MVP

Did it work?

PRIMARY

+20 pts

Average Fairness Score lift, edited vs. original

>60%

Bias-flag suggestions accepted by recruiters

2+

Requirements removed per posting on average

Are they using it?

ENGAGEMENT

- Weekly active recruiters per account
- Postings run through FairPost per recruiter, per month
- Time-to-complete a full rewrite session

Is it working commercially?

BUSINESS

- Trial-to-paid conversion rate
- Net Promoter Score among recruiter users
- Account expansion — teams adding more seats

A tool built to reduce bias must be held to a higher standard than most.

The AI can be biased

Audit the model regularly for disparate impact across demographic groups

Transparent flags

Every flag explains why — recruiters should feel informed, never judged

Human oversight

AI proposes, the recruiter approves — it suggests, it never decides

No false confidence

A high score is not a DEI certification — and we say so plainly

Data privacy

Be explicit about how posting data is stored, used, and whether it trains models

The biggest risk earns its own slide →

The risk I'd **lose sleep** over.

FairPost could become a **compliance shield** — a high score that lets companies look fair on paper while biased hiring continues behind it. Accountability quietly shifts from real decisions into the tool itself.

THE GUARDRAIL

Track downstream outcomes — not just posting quality — and make that data visible to DEI leaders rather than hiding it behind a score.

- | | | |
|-----------|---------------------------------------|---|
| 01 | Candidate Apply Fit View | Shows candidates how they match — including transferable skills |
| 02 | ATS Integration | Embeds into Greenhouse, Lever & Workday — kills friction |
| 03 | Team Dashboard & Analytics | Fairness trends across postings, by department and over time |
| 04 | A/B Testing for Postings | Test two versions; measure which attracts stronger, more diverse applicants |
| 05 | Compliance Module | Flags pay-transparency law violations; generates EEOC-compliant language |
| 06 | Browser Extension | Run FairPost on any job board or ATS without leaving the page |

Q1 MOST CRITICAL

Are we improving hiring **outcomes** — or just improving job descriptions?

Q2 What model powers bias detection — a fine-tuned open-source model, or a third-party API?

Q3 How do we price it — per seat, per posting, or by company size?

Q4 How do we handle industries with legitimately specific requirements (healthcare, law)?

Q5 Who owns the Fairness Score methodology, and on what review cadence?

01

The hardest decision

Recruiter-first vs. candidate-first. The candidate side ships easier — but only the employer side fixes how descriptions get written. I chose system-level impact over short-term usability.

02

What I'd do differently

Make the browser extension the MVP entry point — intervening the moment a job is written, inside the tools recruiters already use. Starting broad made the MVP feel large.

03

On AI products

Success has less to do with model quality than with whether the system changes human behavior. The real challenge isn't intelligence — it's placement in workflow, incentives, and enforcement.