

Stop Filtering. Start Finding.

Why evidence-gated hiring beats application-filtering technology – and why the edge it creates belongs to the companies and headhunters who move first, precisely because the rest of the market is converging on the same tools, the same keywords, and the same average.

THE INCUMBENT – FILTERING

Subtractive. Starts with everyone and removes.

Screens for what is legible – exact titles, keyword density, unbroken employment – not for what is good. Optimises the visible cost (a wasted interview) and is structurally blind to the invisible one: the exceptional candidate rejected before any human ever saw them.

THIS METHODOLOGY – EVIDENCE-GATED FINDING

Additive. Starts with the requirement and seeks.

Ranks demonstrated work above claimed credentials, refuses to conclude on thin evidence, isolates its evaluators so consensus means something, red-teams its own recommendation, and measures the false negative – the cost everyone else ignores.

Companion to the hiring-team skill


Grounded in "You Cannot Hire the Future Through a Filter Trained on the Past" ·
Maintained by JJ

github.com/thejustinjames/ClaudeSkills

What filtering technology actually optimises

The horror story of the rogue robot deleting brilliant CVs is fiction. The real story is worse: a chain of reasonable people and reasonable tools, each making a small, sensible compromise, that together produce a system tuned to reject the people you most need – not at random, but *directionally*.

The dilution cascade – five reasonable steps, five losses



The evidence that this is happening at scale



The two asymmetries that keep it invisible. The cost of a bad interview is visible and immediately felt, so the pipeline minimises it ruthlessly. The cost of a false negative – the elevating candidate rejected unseen – is real, often enormous, and completely unmeasured: you cannot grieve a hire you never knew you missed. And hiring managers are punished for visible bad hires, never for invisible missed ones – so the *safest* candidate becomes the most attractive one. A filter trained on the past will always, faithfully, return more of the past.

Filtering technology vs. the evidence-gated method

Every failure mode of the filter has a specific structural countermeasure in the hiring-team methodology. This is not a philosophy difference – it is a mechanism-by-mechanism inversion.

DIMENSION	APPLICATION FILTERING	EVIDENCE-GATED HIRING
Primary signal	The CV – a <i>claim</i> , keyword-matched for legibility. Non-linear paths and unusual combinations parse as noise.	Demonstrated work – code read, things shipped, a paid trial, watching someone reason. Claims recorded only so they can be corroborated.
Requirements	Proxies as knockouts: years-of-experience, exact titles, credentials, "no gaps". A year short = auto-ranked down.	A gates-vs-proxies audit: four kinds of genuine hard gate stated first; everything else demoted to a soft signal, never a knockout.
Candidate universe	Whoever applied. "Best available" quietly means "best of the inbound pile" – and the genuinely excellent are usually employed, not applying.	A finding plan is part of the decision. A 100%-funnel shortlist is a stated finding that invalidates the comparison set until the search runs.
Evaluation	One shared impression, amplified: panel members contaminate each other, and "everyone loved her" feels like converging evidence.	Five isolated roles with disjoint evidence slices – the Team Integrator never sees the CV; the Role Architect never sees a candidate. Agreement means something because it was not shared.
Its own errors	False negatives unmeasured, forever. No one audits the rejected pile; the filter's mistakes are structurally invisible.	A red team attacks the recommendation on resemblance, false negatives, and wrong-problem – and a standing quarterly rejected-pile audit measures what the process rejected.
The conclusion	A ranking, delivered with confidence, resting on whatever the corpus rewarded. Verdict chosen by narrative and interview warmth.	A verdict bound to an assumption ledger: hire / reshape / trial / keep-looking. An untested load-bearing assumption caps the verdict at <i>trial</i> – no exception for a strong interview.
Thin evidence	Proceeds anyway. The ranking engine never says "I don't have enough to rank on."	A hard gate refuses to run: zero demonstrated work → no hire verdict is even possible. The output becomes a work-trial design – facts first, memo later.

The encouraging counter-signal: in one field experiment, candidates screened by an AI-led *capability* interview had a **53%** success rate in subsequent human interviews, against **29%** for those screened the traditional CV way. Same technology, opposite intent – used to *find* capability rather than *filter* credentials, it widens the aperture instead of narrowing it. This methodology is that intent, made systematic.

Why this is an advantage, not just a method

An edge exists where you can systematically access value your competitors cannot see. The filter's blind spots are enormous, shared across the whole market, and getting more uniform as everyone adopts the same vendors – which makes seeing differently the cheapest advantage available in talent.



You recruit from a pool your competitors have agreed not to see

The monoculture rejects the same people everywhere – career changers, returners, the self-taught, the non-linear. Millions of capable people, pre-filtered out of every rival's pipeline by the same handful of vendors. A process that reads **demonstrated work** instead of keyword-matching claims is, functionally, exclusive access to that pool – at below-market competition, because nobody else is bidding.



Fewer regretted hires – because "hire" has to be earned

The costliest hiring outcome is not the wasted interview; it is the confident wrong hire. Here, an untested load-bearing assumption **caps the verdict at trial** – a paid work sample, structured references, pre-committed thresholds – before an offer exists. Interview polish and warm consensus cannot outrank evidence, because the structure never lets them compete in the same currency.



You hire difference on purpose – and it compounds

The filter selects for resemblance, so incumbent teams drift toward monoculture – the whole org slowly fills with people who can *defend the decision to hire them*. This method's framing question – **"what can this person do that the last good hire could not?"** – plus a red-team lens dedicated to resemblance, selects for capability the team lacks. Each hire widens what the team can do; the compounding runs in opposite directions.



You measure the only cost nobody else measures

The quarterly rejected-pile audit turns the false negative from an invisible loss into a managed metric. Two conversations-worth of talent found in a 20-application sample means the screening criteria get re-audited before the next req opens. Competitors cannot improve a number they have never once looked at.



Decisions you can defend – to boards, regulators, and candidates

Every claim traced to evidence, protected characteristics structurally excluded, a named human decision-maker, an auditable memo with its evidence-quality footer intact. As automated screening comes under regulatory pressure (hiring is a high-risk category under the EU AI Act), an evidence-anchored human process is not just better hiring – it is a smaller compliance surface.

The honest cost: this is slower per candidate than a keyword filter – deliberately. It scales worse, which is exactly why it works: it forces looking at people as individuals rather than rows. The economics only make sense because the alternative's cheapness is an illusion – the filter's speed is paid for in unmeasured false negatives and regretted hires, the two most expensive line items in talent.

The craft, made visible – and billable

A genuinely good headhunter has always been the antidote to the funnel: they do not filter, they hold a map – they know which signals are real, can tell a proxy from a true requirement, and go and *find* the person. This methodology is that craft written down. For the headhunter, that does two commercially important things: it sharpens the work, and it makes the work **demonstrable to clients**.



Differentiation you can show, not assert

Every recruiter claims to "really understand the role." A headhunter running this method hands the client artifacts no CV-forwarder can produce: the **gates-vs-proxies audit** of the client's own JD (usually worth the fee by itself), the irreducible requirement in one sentence, and a shortlist where every candidate carries demonstrated-work evidence rather than keyword matches. Search over subtraction, on paper.



The economic argument for retained work

Contingency fees reward speed-to-fill, which trains everyone toward the template that converts fastest – the average. This method is the professional instrument that justifies **paying for understanding, not throughput**: the evidence file, the isolated evaluations, and the decision memo are visible work product a retainer buys. It converts "trust me" into "here is the analysis".



Fewer failed placements – the rebate killer

A placement that fails in month four usually fails on an assumption nobody tested: "can operate at this scale", "wants to build, not advise". The ledger forces those assumptions into the open and attaches a **trial card** to each – a paid work sample, a structured reference protocol with thresholds set before the call. Every assumption tested before offer is a guarantee-period rebate that never happens.



Your candidate relationships survive the process

The method's dignity rules – hard gates stated first, paid trials, evidence-grounded reasons for every rejected finalist – are exactly how a headhunter protects the asset a funnel burns: the long-term relationship with excellent people who said no this time. The rejected finalist treated well is next year's placement, and the referrer after that.



An honest broker's answer to "just send more CVs"

When a client pushes for volume, the evidence gate is the professional's refusal made legible: **"a shortlist now would be fiction – here are the three cheapest things to learn this week."** The headhunter who can say that, with a structured plan attached, is operating as an advisor rather than a supplier – and advisor is the seat the fees follow.

The same tool, pointed the right way. AI used as a rejection engine reproduces the past at scale. In a headhunter's hands, pointed at *discovery* – reading public work, surfacing the non-obvious background that maps onto the need – it does what the funnel never could. The winners of the next decade of search are the practitioners who pair deep role understanding with machines aimed at finding. This skill is the operating manual for exactly that pairing.

Adopting the method

The methodology is packaged as a Claude skill – `hiring-team` – that runs the whole harness: workspace, evidence gate, isolated role analysis, red team, ledger, verdict, trial cards. Everything is open and inspectable in the repository.

For a hiring company

1

Start with one consequential req

A senior or first-of-its-kind role, where the cost of a regretted hire is highest and the funnel's weaknesses bite hardest.

2

Run the gates-vs-proxies audit on the live JD

Twenty minutes. What it demotes from "requirement" to "signal" is usually the first visible win – and the first widening of the candidate pool.

3

Let the gate fail honestly once

The first run usually stops at the evidence gate. The Evidence Acquisition Plan it produces – a work trial, two structured reference calls, a week of measurement – is the method teaching the org what evidence looks like.

4

Schedule the rejected-pile audit

Quarterly, twenty applications, read by the role-owner. The only measurement the false negative will ever get.

For a headhunter

1

Lead with the audit as the pitch

Run gates-vs-proxies on the prospect's live JD in the first meeting. It demonstrates the craft – proxy from requirement, in front of them – before any fee discussion.

2

Make the evidence file the deliverable

Shortlists arrive as evidence dossiers: demonstrated work read, references conducted with the structured protocol, hesitations captured verbatim. Priced accordingly.

3

Attach trial cards to every offer recommendation

Pre-committed thresholds, written before the trial runs. Every tested assumption is a rebate that never happens and a reference that compounds.

4

Keep the verdict vocabulary

Hire / reshape / trial / keep-looking. "Keep looking, and here is the finding plan" is the sentence that turns a supplier into an advisor.

The one-sentence case. Every part of the standard pipeline is built from the past – the JD corpus, the "successful candidate" profile, the screening rules. But the talent you now need has changed shape, and the people who change companies rarely look like the people who maintain them. You cannot hire the future through a filter trained on the past – so stop filtering for the person you already hired, and go find the one you cannot yet imagine.