
What Happens After You Hit Submit

*Where your application actually goes — and why half of all
hires come from somewhere else*

The Direct Line

What Actually Happens After You Hit Submit

Chapters 1 and 2 of The Direct Line

These are the first two chapters of the book, exactly as they appear in it. Not a summary, not a teaser — the actual opening.

They're free because they do a job on their own. By the end of Chapter 2 you should be able to identify where your applications are most likely getting stuck, which is something most people never find out. That's useful whether or not you ever read the rest.

How to read the numbers in here. Some of what follows is measured, some is estimated, and some is my reasoning. I mark which is which every time, and there's a short sources section at the end listing every figure with its origin, sample size and limitations. That's unusual in this category, and it's deliberate: two of the most-quoted statistics in career advice turn out to be fabricated, and I'd rather you were able to check me than trust me.

What this won't tell you. It won't tell you the system is always the problem. Sometimes it isn't — sometimes it's a missing requirement, a salary mismatch, or an oversupplied field. Knowing which situation you're in is worth more than any technique, and Chapter 2 covers how to tell.

CHAPTER 1

The Thousand

You already know the number. You saw it on the listing: *over 1,000 applicants*. Or 847. Or "100+" because the platform stops counting.

What you probably haven't done is the arithmetic.

What happens to a thousand applications

One role. A thousand people. Here is one plausible way that plays out.

This table is illustrative, not measured. I've built it from the shape of a hiring process and the two figures at the bottom that *are* sourced. Nobody publishes stage-by-stage data for thousand-applicant roles, and any book that presents such a table as fact is inventing it. I'd rather show you my reasoning than dress up a guess.

Stage	People left	Basis
Applied	1,000	Given
Cleared the automated screening questions	~800	Illustrative
Actually reached by a human being	~80–150	Illustrative
Survived the resume review	~30–50	Illustrative
Got a phone screen	~10–20	Illustrative

Stage	People left	Basis
Put in front of the hiring manager	4–8	Illustrative
Interviewed	3–6	Illustrative
Hired	1	Given

The middle rows are my estimates. What matters isn't their precision — it's the shape, and the shape is not really in dispute: a very large number of applications, a much smaller number of human reviews, and a handful of people who ever reach the person deciding.

Look at the third row. If anything like that is right, most applicants are never *rejected* at all. They're never reached. That's a different thing, and it matters, because "rejected" implies someone looked at you and decided. Mostly, nobody looked.

Two figures give the shape of it. CareerPlug's 2025 Recruiting Metrics Report analysed more than ten million applications from over 60,000 businesses and found an **applicant-to-interview rate of 3%** and an **applicant-to-hire ratio of about 180 to 1** — with wide variation by sector, from around 47 applicants per hire in healthcare to 191 in technology.

Three caveats you should have, because they change how much weight this carries.

That dataset is drawn overwhelmingly from *small businesses*. Your odds at a thousand-applicant role at a large employer are not the same, and are probably worse. The two figures also aren't derived from each other — 3% of applicants interviewing and 27% of interviewees being hired implies roughly one hire per 123 applicants, not 180 — which tells you these are separately measured metrics with different denominators, not a clean funnel. And an average across ten million applications is not your personal probability. Your odds depend on your field, your location, your experience, and which specific

roles you chose.

So use it for scale, not for prediction. What it establishes is that **applicant-to-interview conversion is measured in low single-digit percentages** — not that you personally have a 3% chance each time.

With that said, run the arithmetic. At something like a 3% rate, fifty applications would be expected to produce one or two interviews. Getting none isn't unusual: if applications behaved like independent coin flips at 3%, zero from fifty would happen about a fifth of the time. They don't behave independently — you're not a random draw and neither are the jobs — but the point survives.

Which means: getting very little back from fifty applications is consistent with a normal process in a congested channel. That is not the same as proof that your process is fine. It might not be; Chapter 5 covers the cases where the problem is something real. What it does mean is that silence is weak evidence about you, and most people have been treating it as strong evidence.

This is new, and that's why nobody warned you

Here is the part that matters most, and the part that explains why your parents' advice and the careers blog from 2019 are both useless to you now.

The channel didn't break. It got crowded — and you were in it when that happened.

That distinction matters, and I'll come back to it, because "applications don't work" is both wrong and the wrong lesson.

Ashby, analysing more than a hundred million applications across two hundred thousand jobs, put the current figure at **300+ applications per posting**. Gem's 2026 benchmarking — comparing its own platform against itself, which is the cleaner measure — found recruiters handling **93% more**

applications than in 2021. Greenhouse reported applications per recruiter up **412%.**

Different platforms count differently and serve different employers, so don't lean on any single cross-source comparison. But every dataset published in the last two years points the same direction, and the within-platform figures are the ones that matter: Gem's near-doubling in twelve months isn't an artefact of comparing apples to oranges.

LinkedIn now processes something in the order of eleven to fourteen thousand applications *per minute*.

Meanwhile the capacity to process all this went the other way. Gem found recruiting headcount **down 14% against 2021**, recruiters carrying **13.4 open roles each**, and the average number of interviews per hire **up 33% since 2021**. More applications, fewer people reading them, and a more thorough process for the few who get through.

The top of the funnel widened enormously. The bottom did not keep pace — and the sharpest evidence for that is the recruiter figure: if the number of people doing the reviewing had grown alongside the volume, applications *per recruiter* wouldn't have more than quintupled.

Why it happened

Because applying became free.

Not free in money — free in effort. It used to take an hour to tailor a CV and write a covering letter. Now it takes ninety seconds and a language model. Tools exist, for about twenty dollars, that will apply to two hundred jobs while you sleep.

When the cost of using a channel falls to nearly zero, people use it without limit. That is not a moral failing on anyone's part; it is just what happens. Nearly half of job seekers report applying to *more* positions specifically

because they believe automated filters are screening them out — which is, of course, precisely the behaviour that creates the filtering.

Greenhouse's chief executive has a name for it: the **doom loop**. Candidates use AI to apply to more jobs. Employers respond with more aggressive automated filtering. Candidates escalate again. Round and round.

The consequences show up in the data exactly where you'd expect. Candidates are roughly **half as likely to reach an interview as they were five years ago**, and around **three times less likely to be hired than three years ago** (Ashby 2026; Gem 2025). Fifty-three percent of job seekers reported being ghosted in the past year — a three-year high.

You didn't get worse. The queue got longer.

The number that changes everything

Everything so far has been about how crowded the application channel is. That's the setup. Here is the actual argument, and it's the reason this book exists.

Gem's 2026 report tracked where applications come from and where hires come from — separately. The gap between those two things is the most important fact in this book.

Job boards and company career pages generate roughly 90% of all applications — but account for only about half of hires.

Direct sourcing produces 11% of hires from just 2.6% of applications.

Referrals convert at 11 times the rate of inbound applicants. Internal moves, 32 times.

Overall: a sourced candidate is nearly eight times more likely to be hired than an inbound applicant.

Gem, 2026 Recruiting Benchmarks Report — 165 million applications, 1.2 million hires, June 2021 to May 2025.

Read that again, because it isn't saying what you might assume.

It is **not** saying applications don't work. Half of all hires still come through them. Millions of people get jobs by applying, every year, including probably most people you know.

What it says is that **the application channel is where almost everyone competes, and it is not where hiring is most efficient.** Nine applications in ten arrive through the front door. Only half the hires come out of it. Meanwhile a small minority of candidates who arrive some other way — sourced, referred, already known — account for the rest, at dramatically better odds.

That is the whole thesis of this book, and I want to state it precisely because it is not the usual claim:

You do not need to stop applying. You need to stop making applications your only channel.

The people getting hired disproportionately are not applying harder. They arrived by a different route.

Which raises the obvious question, and it's the one the rest of this book answers: those other routes belong to recruiters, referrals and internal candidates. **If you don't have a referral and nobody is sourcing you, can you get onto that side of the ledger yourself?**

Mostly, yes. That's what Chapters 3 to 10 are about.

Why effort still counts

One supporting idea, and I'll be honest about its limits.

Economists call it **costly signalling**: a signal only carries information if it's expensive to send. Applying used to cost an hour, so applying meant something. It now costs ninety seconds, so it means much less. Research still costs something — which is why it still signals.

But effort can increasingly be faked at volume. AI can generate a plausible personalised message, find a hiring manager, and summarise a company's recent news. The moat isn't that this is impossible to automate. It's that automated versions are currently mediocre in ways practitioners notice, and that most people won't bother.

A real edge, and a shrinking one. Which is an argument for doing this now rather than in three years.

WHAT THE NUMBERS DON'T PROVE

These benchmarks do not mean that every applicant faces the same odds. They do not mean every employer uses AI ranking. They do not mean every unanswered application was rejected by software, or that every long-running posting is a ghost. And they do not mean your CV is fine — it might not be.

They show something simpler: **high-volume application channels create severe competition for attention, and the channels where most people compete are not the channels where most hiring efficiency sits.**

This book is about building a second route to visibility. It is not a claim that the first one is worthless.

Some of what you applied to was never going to hire anyone

One more thing, and then we'll move on to something more useful than bad news.

A meaningful share of the postings you've applied to were not going to result in anyone being hired soon.

The estimates vary a lot, and **they vary because they measure different things** — which is worth understanding rather than glossing over:

- Greenhouse, examining its own platform, found **18–22% of listings showed no hiring activity**. That's an observation about what happened, not about intent.
- An analysis of LinkedIn data put the share of likely non-genuine postings near **27%** — a different platform, a different method.

- A ResumeBuilder survey found roughly **40% of employers said** they post roles without an immediate intention to hire. That's self-reported intent, and "not hiring immediately" is not the same as "no job exists."

So I won't give you a single number, because there isn't an honest one. What's fair to say is that **a substantial minority — somewhere in the region of a fifth to a third, depending on how you count — won't lead to a hire on the timeline you're imagining.**

The reasons are mostly mundane: building a pipeline for a high-turnover role, a placeholder while budget is approved, a compliance posting where an internal candidate is already lined up, or simply nobody closing a filled requisition.

It's real enough to have been legislated. Since 1 January 2026, Ontario employers with 25 or more employees have had to meet several requirements on publicly advertised job postings under the Employment Standards Act, 2000 (Part III.1, ss. 8.1–8.6), introduced by the Working for Workers Four and Five Acts, 2024 and O. Reg. 476/24. Among them:

- **State whether the posting is for an existing vacancy** (s.8.5)
- **Disclose if AI is used to screen, assess or select applicants** (s.8.4)
- Disclose expected compensation or a range, where the range doesn't exceed \$50,000 and the role pays under \$200,000
- Tell anyone they interviewed the outcome within 45 days
- Retain postings and applications for three years

If you're applying in Ontario, read the posting: whether the job is real and whether a machine is reading your application may simply be printed on it. Nowhere else has that yet — but it's a useful illustration that this is a recognised problem, not a job-seeker grievance.

Be careful what you conclude from that about your own history. You can't take a platform-level estimate and apply it to your two hundred applications — the postings you personally encountered aren't a random sample of all postings. What you can reasonably say is that **some meaningful share of what you applied to was never going to hire anyone on the timeline you imagined**, and that you had no way of telling which.

In Chapter 5 I'll show you how to spot most of them in about ninety seconds.

What to do with this chapter

Nothing yet. This chapter has one job, and it isn't tactical.

If you've spent six months applying and hearing nothing, you have almost certainly built an explanation for it, and the explanation is almost certainly about you. Your CV. Your experience. Your age. Something you can't identify but must surely be there, because otherwise why would nobody answer.

Put that down. Not because it's kind to, but because it's inaccurate, and because you cannot act well from inside it. People who believe the problem is them do one of two things: they apply more, which makes it worse, or they stop, which makes it worse.

The problem is a queue. Queues are structural. You don't beat a queue by being better at queueing.

You beat it by not being in it.

THE THOUSAND CALCULATOR

Before the next chapter, work this out. It takes four minutes and it is the last time you'll need to think about it.

· Applications sent in the last 90 days: _____

· Responses of any kind: _____

· Phone screens: _____

· Interviews: _____

Your response rate: $\text{responses} \div \text{applications} = \text{_____}\%$

Your interview rate: $\text{interviews} \div \text{applications} = \text{_____}\%$

For scale: CareerPlug's 2025 benchmark reports ~3% of applicants reaching interview and ~180 applicants per hire — from a predominantly small-business dataset. Not your personal odds; a sense of the order of magnitude.

If you're near those numbers, your process isn't broken — you're getting the standard outcome of a broken channel. If you're well below them, Chapter 2 will tell you which stage is failing, and it's probably not the one you think.

CHAPTER 2

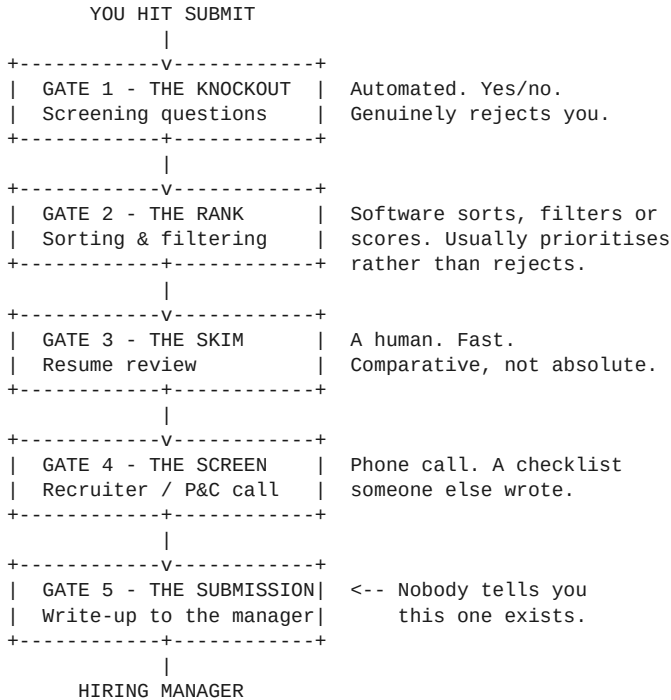
The Five Gates

Most people imagine their application arriving somewhere and being looked at. One door, one decision.

In practice there are usually several, they fail for completely different reasons, and they have completely different fixes. This is why so much job-search advice doesn't work: it's aimed at one stage, and you may well be failing at a different one.

A word on what this model is and isn't. What follows is the five-stage diagnostic model I use in this book. It is not a description of a process every employer runs. Real hiring workflows vary enormously: some employers use all five stages, many use three, a few use two. The order changes. A small company may have the hiring manager reading applications directly. A large one may add assessments, panel reviews or a second recruiter screen that aren't shown here.

The model earns its place by being **diagnostically useful** — it gives you named, distinguishable places where an application can stall, so you can work out which one is likely stalling yours. Treat it as a map, not a photograph.



Notice where the hiring manager sits. Behind all of it. The person who genuinely understands the work you'd be doing is the last to see you — if they see you at all.

Hold that thought.

Gate 1 — The Knockout

This is the only gate where a computer definitely, straightforwardly rejects you. And it isn't mysterious.

When a recruiter sets up a job posting, they configure a handful of questions as pass/fail. Do you have the right to work here. Do you hold the required licence. Do you have the minimum years. Are you within commuting distance. What's your salary expectation. When could you start. Would you

relocate.

Answer one of those in a way that falls outside the configured range and you're out. No algorithm required — it's a filter, in the plumbing sense.

How people fail it: by rushing. You've done this form nine times today. You're clicking through. You answer "no" to the relocation question because you already live there and the question reads oddly. You type a salary figure without knowing the band. You answer a years-of-experience question literally when your actual experience is adjacent and would have counted.

The fix: slow down on the questions and speed up on everything else. Those five minutes are the highest-return five minutes in your entire application, and they're the ones everybody rushes.

And if you think you've fluffed one, you can often fix it. There's an email in Chapter 10 that works more often than people expect.

Gate 2 — The Rank

Now the part you've heard about, mostly wrongly.

You have almost certainly encountered the claim that **75% of resumes are rejected by AI before a human sees them**. It's everywhere. It's in the marketing of nearly every resume-writing service.

I went looking for the source.

It traces back to a 2012 claim made by a company called Preptel, which sold resume-optimisation software. No methodology was ever published. No study was ever produced. Preptel went out of business in 2013.

So the number that has frightened millions of people into buying resume services came from a resume service, thirteen years ago, and was never substantiated. The company that made the claim didn't survive the following summer.

But — and this matters — something real has changed since then, and I'm not going to pretend otherwise.

AI screening is genuinely growing, and fast. In 2025, 35% of hiring managers reported using AI to screen resumes. In 2026 it was **58%** — the largest single-year jump of any recruitment use case (Resume Genius, 2026 Hiring Trends Report). Eighty-seven percent of hiring managers say their company uses AI somewhere in recruitment. Adoption estimates range from 27% to 87% across different surveys, mostly because they define "AI" differently, but the direction is not in question.

Here are the two numbers almost nobody quotes, and they're the ones that should actually concern you:

Only about 29% of companies maintain full human oversight over all AI rejection decisions. Around 21% permit AI to reject candidates at any stage without human review. And in the same Resume Genius survey, 19% of hiring managers said they use AI to deliberately screen out applications before any human review. That's employers describing their own practice.

So what's actually true? It depends on the employer, and there isn't one universal mechanism. "AI screening" covers several quite different things: keyword and Boolean filters, search and sorting tools, match scores, assessment platforms, workflow automations that move or tag candidates, and newer generative-AI features that summarise or rank applicants. Different employers use different combinations, configured differently.

What most of them have in common is this: **they mostly prioritise rather than eliminate.** They help a reviewer decide who to look at first.

Which sounds better than rejection until you do the arithmetic from Chapter 1. If a thousand people apply and the reviewer gets through the first hundred, being a hundred places further down is, in practice, the same outcome as

being turned down. You weren't rejected. You were never reached.

A caveat on that 21% figure, because it's the kind of statistic that gets stretched. "Permits AI to reject at some stage without human review" is not the same claim as "AI reads and bins your CV before a human sees it." It could mean an automated knockout on work authorisation. It could mean an assessment score cut-off later in the process. The survey doesn't tell us which, and neither will I. What it does establish is that fully automated rejection exists at a meaningful minority of employers — which is more than was true three years ago, and less than the headline suggests.

The fix, and it's smaller than you'd like: make sure a matching system can recognise that the thing you did is the thing they asked for. If the posting says "FP&A" and your CV says "financial planning and analysis," write both. If they name a specific tool, name that tool. Use ordinary section headings. Keep your contact details in the body of the document, not in a header or footer. Export a real PDF, never a photo or a scan.

Do not stuff keywords in white text. It's detectable, people do check, and being caught doing it ends your candidacy in a way that no missing keyword ever would.

That's the whole fix. About ten minutes. Anyone selling you more than that for this gate is selling you comfort.

Gate 3 — The Skim

A human being now looks at your application briefly — and the honest answer is that nobody knows exactly how briefly.

You'll see "six seconds" or "seven point four seconds" quoted everywhere. Those come from small eye-tracking studies of recruiters reviewing resumes in controlled conditions, and they've been repeated so often that they've hardened into fact. Treat them as an indication that initial review is *fast*, not

as a measured parameter of your particular application. A reviewer with four roles to fill and eight hundred applications behaves differently from one with a single vacancy and forty.

Here's the thing nobody says about this gate, and it changes how you should write:

You are not being measured against the job description. You are being compared to whoever else applied this week.

It's a sort, not an assessment. The reviewer isn't asking "is this person good enough?" They're asking "is this person obviously more relevant than the next forty in this pile?" That's why identical applications get different outcomes at different companies, and why the same CV can fail in March and succeed in June.

Remember also that this person is under real pressure. Gem's 2026 benchmark found recruiters handling **93% more applications than in 2021** while managing **13.4 open roles at once**, on teams that are 14% smaller than they were.

The fix: the top third of your first page has to make the match obvious without scrolling. Your title in their language. Your closest comparable piece of work. The hard requirements visible immediately. That's it. The bottom of page two is decoration.

Gate 4 — The Screen

You get a call. Twenty to thirty minutes, usually with someone from People & Culture or an internal recruiter.

What's actually happening on this call is not what most candidates think.

The person is running a checklist that somebody else wrote. Screening guides are typically built for a role family and signed off by the hiring manager — so the criteria being applied are the manager's stated

requirements, translated into questions a non-specialist can ask and score.

Which leads to the most useful thing to understand about this stage: **for specialist roles, the person screening you often can't fully assess the substance of your work** — and is therefore checking your answers against the requisition rather than evaluating them on their merits. If you're an engineer, they're generally not evaluating your system design.

Be careful with this, though. Plenty of internal recruiters have deep domain knowledge, particularly ones who've recruited in a field for years, and a good screener is genuinely assessing communication, motivation, career logic and how you handle a question — all of which are substantive. Don't walk into the call assuming the person is a box-ticker. Assume instead that **they need to be able to tick the box**, and make that easy for them while treating them as a professional.

They're checking four things:

1. **The must-haves, verified.** Years, licences, authorisation, non-negotiable skills. Binary.
2. **Whether you actually want this job.** Managers' time is rationed, and screeners get burned by candidates who were only browsing. They're testing for interest in *this* company, specifically.
3. **Salary alignment.** Often the only genuinely decisive item on the call, and the fastest way to be removed.
4. **Whether you'll be easy to put in front of a manager.**

And here's the reassuring part, which nobody tells nervous candidates: the bar for passing is low. You generally don't fail a competent screener unless you actually lack a stated must-have, you're evasive about a deal-breaker, or you come across as someone difficult to deal with. Sounding flat or bored is a real risk — it lands worse on a phone call than it does in person — but that's about it.

Of the people who get screened, roughly **30–50% advance to a first interview**; in technical roles it's more like 20–35%.

The fix, in one line: give them the tick-box phrase first, then the evidence. Not "well, at my last job I was involved in a project where..." but "Yes — five years of multi-entity consolidation. Most recently six entities across three ledgers, where we took the close from eleven days to four." Now they can tick and quote you. Tell the story first and they may never extract the qualification at all.

Gate 5 — The Submission

This is the one nobody tells you about, and it's the one that explains your worst experience of job hunting: the good phone screen, the warm conversation, the "we'll be in touch" — and then nothing, ever.

Here is what happens after a good screen.

The person generally doesn't simply pass you along. In a structured process there is usually some point at which the screener **summarises you** for the hiring manager — sometimes a written narrative, sometimes a scorecard, sometimes a conversation — and you go forward as part of a shortlist rather than as an individual file.

The mechanism here is better documented than you might expect — not in career advice, but in the manuals of the software employers actually use.

Workable's documentation for interview scorecards is explicit. Scorecards can be attached to **phone screen** stages. They're added automatically to the candidate's profile and are "accessible to the entire team." The interviewer rates the candidate against each skill or question — and at the end must "include an overall rating and comment (**mandatory field**) providing additional details as to why this candidate is or isn't a good fit for the job."

Greenhouse works comparably: scorecard feedback is stored on candidate profiles and visible to recruiters and hiring managers, and it offers a dedicated Hiring Manager Review stage where managers examine applications alongside prior feedback. Indeed's own guidance for employers tells screeners to prepare a recording sheet, take detailed notes, and leave time after each call to write them up.

So the underlying claim isn't speculative: **in structured hiring systems, what you say in a screen may be converted into notes, ratings, comments or a summary that becomes visible to the people deciding whether you advance.** In Workable, writing that justification isn't optional.

How widely this applies, I can't tell you. Two vendors' documentation and one employer guide establish that the mechanism exists and is built into widely-used tools. They don't establish what proportion of employers work this way, or what happens in any particular screen.

What isn't universal is the form. It might be a prose paragraph. It might be a numeric scorecard, a thumbs rating with comments, a conversation in a pipeline review, or an AI-generated summary a human signs off. Different employers, different systems, different habits. I don't know which yours uses, and I'm not going to pretend to.

Three things follow, and all three are invisible from where you're standing.

One: someone may be choosing. A screener can end up with several viable candidates and a smaller number of interview slots. Where that happens, the question stops being "is this person qualified?" and becomes "which of these go forward?" — a comparative judgement you're not in the room for.

Two: the summary is a compression. However it's recorded, a twenty-minute conversation becomes something much shorter. Whatever isn't specific and memorable is likely to be lost, and what survives is what the manager sees.

Three: it may be recorded during or after the conversation. Indeed advises employers to schedule time after each call for note-taking, which tells you the practice exists — not that every screener works that way. Either way, the less specific your answers were, the more gets flattened.

Put those together and you get the sentence I think matters most in this book:

At this stage, you are not represented by your CV. You are represented by however someone else has summarised you.

You have never seen that summary. You probably never will. It may be the most consequential thing written about you in the whole process, and you had almost no influence over it.

In Chapter 10 I'll show you how to write it for them — properly, helpfully, in a way that's genuinely useful rather than looking like you're managing them. Almost nobody does it. Whether it helps is a question Chapter 10 is honest about.

Which gate are you failing?

Here's why this chapter matters more than any CV advice you'll ever read.

Applicants overwhelmingly reach for the CV, because it's the only part of the process they can see and control. But a CV rewrite only affects one stage. If your applications are stalling at the screening questions, or after a phone screen, three months of redrafting changes nothing about the thing that's actually stopping you.

You can narrow this down yourself, right now, from what's already happened to you.

Read the middle column as "where to look first," not "what went wrong." Every one of these patterns has several possible causes, and you can't tell them apart from outside. What the table does is stop you working on

a stage that isn't your problem.

What you're seeing	Likely places to investigate	What you could act on
Automated rejection within hours	A screening-question knockout — but also requisition closure, location or work-authorisation rules, duplicate handling, or a prior-application rule	Re-read your screening answers first; it's the cheapest thing to check
Total silence, weeks, no acknowledgement	Never reached in the queue — or a posting that was never going to hire	Requisition language; ghost checks (Ch 5)
Rejection after 1–2 weeks	A human looked and moved on — or a stronger field	Top third of page one
You get screened, then silence	The stage after the screen — but equally a stronger candidate, an internal hire, a pause, a budget change, or simply no follow-up	Chapter 10
Warm screen, positive signals, nothing	Same range of causes. Warmth isn't a signal about the outcome	Chapter 10
Repeated first interviews, no progression	Probably not the stages at all	Something real — see Chapter 5

Do this now, honestly, before you read on. Ten minutes.

Treat it as a hypothesis, not a diagnosis. You're reasoning backwards from a pattern to a likely cause, with incomplete information — the same thing a doctor does with symptoms, and with the same fallibility. Sometimes silence after a good screen means a better candidate turned up. Sometimes an automated rejection is exactly what it says. The value here isn't certainty. It's that you now have something specific to test, instead of a vague sense that

something is wrong with you.

WHAT'S DOCUMENTED HERE, AND WHAT ISN'T

Documented: that structured hiring systems capture screen feedback as notes, ratings and written justifications; that this record sits on your candidate profile; and that it's visible to the hiring team. This comes from the vendors' own product documentation — Workable, Greenhouse — and from Indeed's guidance to employers.

Not documented, and therefore not claimed: how many candidates a given employer puts forward, how much discretion an individual screener has, and what motivates their choices. You'll see those absent from this chapter. I've seen them asserted confidently elsewhere and I can't source them.

If you work in recruitment or talent acquisition and any of this is wrong, I want to know: support@thedirectline.com.au. Corrections go into the next edition with attribution.

Where that leaves you

Let's put the two chapters together, because they say something specific.

Applications are a channel, not the whole of hiring. A crowded one — around 90% of all applications arrive through job boards and career pages, and about half of all hires come out of them. Applying works. Millions of people get jobs that way every year. But you are competing there with almost everyone else.

High volume creates a competition-for-attention problem, not a quality problem. When a role attracts hundreds of applicants, most are never carefully assessed — not because anyone decided against them, but because a triage system is optimised for throughput. Whatever number of stages a particular employer runs, that's what they're for: reducing a very large

number of people to a handful without anyone reading a very large number of things closely.

Meanwhile, the routes fewer people use produce disproportionate outcomes. Direct sourcing: 11% of hires from 2.6% of applications. Referrals converting at eleven times the rate of inbound applicants. Sourced candidates nearly eight times more likely to be hired. That's not a trick or a loophole — it's what the hiring data says about where hires actually come from.

And you can now locate your own bottleneck. That's what the diagnostic was for. If you're being knocked out on screening questions, that's a twenty-minute fix. If you're never reaching a human, that's about targeting and readability. If you keep getting screened and going quiet afterwards, the problem is happening at a stage you've never seen.

Which leaves one thing outstanding.

Knowing where you're stuck doesn't tell you what to do instead. The routes that produce disproportionate hires — sourcing, referrals, being already known — belong to recruiters, to people with connections, and to internal candidates. If nobody is sourcing you and you don't have a referral, none of that is available to you by default.

But it is available. It's just not automatic.

Recruiters get hiring managers on the phone every working day, at companies with nothing advertised, about people those managers have never heard of. The mechanism is well established, it's ordinary commercial practice rather than a secret, and there is nothing in it that requires you to deceive anyone or to be someone you're not.

You now know what's happening to your applications. The rest of this book is about how to change the outcome: how to work out which opportunities are worth real effort, how to find the person who owns the problem, how to

understand that problem well enough to say something useful about it, and how to open a conversation that a busy stranger is glad to have.

It starts with a phone call that built an industry.

Sources

Every quantitative claim in these two chapters, with what it measures and how much weight it deserves.

Gem, 2026 Recruiting Benchmarks Report. 165M+ applications, 15M candidates, 1.2M hires, June 2021 – May 2025, from Gem's customer base (predominantly US, tech-weighted). Source for: job boards and career pages generating ~90% of applications but ~half of hires; direct sourcing 11% of hires from 2.6% of applications; referrals converting at 11× inbound and internal mobility at 32×; sourced candidates ~8× more likely to be hired; recruiters handling 93% more applications than 2021; 13.4 open roles per recruiter; recruiting headcount down 14%; interviews per hire up 33%. *Limitation: one platform's customers, not the whole labour market.*

CareerPlug, 2025 Recruiting Metrics Report. 10M+ applications from 60,000+ businesses, 2024 data, **predominantly small businesses**. Source for: 3% applicant-to-interview; 180 applicants per hire; ~191 in technology, ~47 in healthcare. *Limitation: small-business weighted, and note that 3% × 27% interview-to-hire implies roughly 123 applicants per hire rather than 180 — these are separately measured metrics, not a single clean funnel.*

Ashby, 2026. 100M+ applications across 200,000+ jobs. Source for: 300+ applications per posting, roughly triple 2021.

Resume Genius, 2026 Hiring Insights Report. Survey of 1,000 US hiring managers. Source for: 87% report AI somewhere in recruitment; 58% use AI to screen resumes, up from 35% in 2025; 19% use AI to deliberately screen out applications before human review; 49% personally review resumes while

using the ATS to organise them; 37% allow ATS screening on employer-set criteria.

Greenhouse. 18–22% of listings on its platform showed no hiring activity. *Measures observed activity, not intent.*

ResumeBuilder. ~40% of surveyed employers said they post without immediate hiring intent. *Measures stated intent. Does not establish that a posting was fictitious.*

Ontario, Employment Standards Act 2000, Part III.1 ss. 8.1–8.6. Introduced by the Working for Workers Four and Five Acts 2024 with O. Reg. 476/24; in force 1 January 2026 for employers with 25+ employees.

The "75% of resumes are rejected by ATS" claim. Traced to a 2012 marketing claim by Preptel, a resume-optimisation company with no published methodology, which ceased trading in 2013. Referenced here only to debunk it.

The "six to eight second resume review." Derives from small eye-tracking studies in controlled conditions. Used here only as an indication that initial review is fast, not as a measured figure.

The illustrative 1,000-applicant funnel and the five-stage model. Mine. Built from the evidence above and from how hiring processes are structured. Neither is a validated instrument; both are labelled as models where they appear.

If you find an error in any of this, I'd like to know. Email support@thedirectline.com.au.

*These are Chapters 1 and 2 of **The Direct Line** — a 121-page method for creating a legitimate conversation with the person who owns the role, instead of competing only in the application queue. The remaining eight chapters*

cover which opportunities are worth real effort, how to find the person who owns the problem, how to say something useful about it, fifteen worked messages across email, LinkedIn and the phone, the fifteen-minute conversation, and nine common brush-offs with responses to each.

<https://thedirectline.com.au>