Fill Rate, Time to Fill, and the Staffing Funnel

Fill rate and time to fill get reported as recruiting KPIs constantly, and just as constantly get discussed without a dollar figure attached. A five-point fill-rate improvement or a five-day-faster average time to fill sound like good things. Whether they actually matter enough to invest in fixing depends entirely on what they're worth, which is a calculation, not an instinct.

What a fill-rate improvement is actually worth

Fill Rate Financial Impact takes your eligible job orders, current and target fill rate, and average GP per fill, and computes the dollar value of closing that gap. This reframes fill rate from a percentage that sounds better or worse in isolation into a specific number of additional filled reqs and the GP they represent, which is the number that actually justifies (or doesn't) investing in whatever would close the gap, better sourcing, more recruiter capacity, or a change in which reqs you're accepting in the first place.

Is filling faster worth more, or just earlier

Time-to-Fill Financial Impact asks a subtler question than it looks like. Filling a req faster clearly matters to the client, but does it matter financially to the agency, or does it just move the same GP earlier without adding to it? The answer depends on what's constraining your business. If recruiter capacity or req volume is the binding constraint, filling faster frees up capacity to work more reqs, which does add GP. If neither is binding, and reqs would get filled eventually regardless, faster mostly improves client satisfaction and win-rate on future business rather than adding measurable GP to this specific placement. Know which situation you're in before treating "reduce time to fill" as an automatic financial win.

Finding the actual bottleneck

Recruiting Funnel Conversion looks at conversion rates stage by stage, sourced to submitted, submitted to interviewed, interviewed to placed, rather than a single blended fill rate that hides where the real problem is. A low overall fill rate caused by a sourcing problem needs a completely different fix than the same fill rate caused by a submission-to-interview problem, and treating both as "fill rate is low, source harder" wastes effort on the stage that isn't actually broken.

One structural note worth stating directly: funnel stages don't have to narrow monotonically at every point to be healthy. A single job order can legitimately source several candidates, meaning the funnel can widen before it narrows again toward a single placement; what actually signals a problem is a stage that goes from a positive count to zero, not a stage that expands relative to the one before it.

Knowing which reqs are worth working

Not every open req deserves equal recruiter attention. Job Order / Req Expected Value weighs a req's GP against your realistic probability of filling it, producing an expected-value figure that separates "this pays well but we almost never fill this kind of role" from "this pays well and we usually fill roles like it," two reqs that can look identical on GP alone but deserve very different amounts of recruiter time.

Putting the three together

A recruiting operation that's underperforming usually shows up first as a fill-rate or time-to-fill number that looks worse than it should. The funnel conversion breakdown is what turns "something's wrong" into "this specific stage is where candidates are being lost," and req expected value is what tells you whether the fix is process (better conversion at the broken stage) or triage (spend less recruiter time on reqs unlikely to fill regardless). Treat the headline metrics as the smoke alarm, not the diagnosis.

Related

- Staffing Recruiter Economics
- Every recruiter economics engine