

A WORKING MODEL FOR PEOPLE LEADERS

The cost of finding out too late.

What burnout-driven attrition actually costs an Indian company — built from the bottom up, line by line, in rupees, so it survives the meeting where someone asks where the number came from.

What this is, and what it deliberately is not.

This is a costing model, not a manifesto. It exists to turn a feeling every people leader already has — that burnout is expensive long before anyone resigns — into a number a finance team can audit.

It is

A bottom-up model with eight cost lines, each one defined, each one traceable to either a published source or an assumption you set yourself and can be challenged on.

India-specific where that matters: notice periods measured in months rather than weeks, recruitment fees quoted as a share of CTC, and the two-sided notice problem that makes an Indian vacancy longer than a Western one.

Honest about its own weak points. Where the published evidence is thin, this document says so and hands you a method instead of a borrowed statistic.

It is not

A benchmark report. There is no table here telling you what attrition "should" cost in your sector, because the numbers needed to build one honestly are not public.

A multiplier. The widely quoted rule that replacing someone costs half to twice their salary is a starting hypothesis, not a business case. Part 2 explains why, and the rest of the document replaces it.

A case study. The worked example in Part 5 is a constructed company. It is not a customer, and no figure in it is drawn from anyone's real data.

The one rule behind every number in here

No figure appears in this document without one of two labels: a **source**, meaning a named publisher, a date and a sample we could read; or an **assumption**, meaning a value you are expected to replace with your own. Where a number is widely repeated but we could not trace it to a primary document, it is excluded or flagged. Part 10 lists every source with what it does and does not support.

This matters more than it sounds. The reason most wellbeing business cases fail is not that the costs are small. It is that the case arrives built on a statistic nobody can defend, and one unanswerable question from the CFO takes the whole thing down.

CONTENTS

Ten parts, two of which you can skip.

PART 1	The number nobody has	Why the cost is invisible until the exit interview	5
PART 2	Why "half to twice salary" will not survive your CFO	A ten-fold spread, and what to do instead	8
PART 3	The model: eight lines	Definitions, formulae, and the double-counting traps	10
PART 4	The two costs that never reach the attrition report	Presenteeism and absenteeism, at population level	19
PART 5	Worked example: 200 people, in rupees	Two roles costed in full, then the annual aggregate	21
PART 6	What early detection is worth	Break-even, sensitivity, and the value of lead time	25
PART 7	The one-page business case	A template, and the four questions you will be asked	27
PART 8	How you will know if it worked	Making the claim falsifiable before you make it	29
PART 9	Your assumptions log	A blank worksheet, to fill in and attach	30
PART 10	Sources, and what each one supports	Every figure, dated, with its limits stated	31

Two routes through it

IF YOU HAVE	READ	YOU WILL LEAVE WITH
Twenty minutes	Part 2, Part 5, Part 7	A defensible number for your own headcount and a one-page case to take into a budget conversation.
An afternoon	All of it, with Part 9 open	A full model with your own assumptions logged, ready to be challenged line by line.

A note on arithmetic. Rupee figures are written in the Indian convention: one lakh is ₹1,00,000 and one crore is ₹1,00,00,000. All per-day figures assume 240 working days a year unless stated. Rounding is to the nearest ₹1,000 within tables and the nearest ₹0.01 lakh in totals, so columns may differ from their sum by a rupee or two.

Nine inputs do most of the work.

You do not need a data project to run this model. Nine numbers, most of which your HRIS and ATS already hold, will get you to a defensible range. Everything else in the document exists to help you argue about these nine, and the spreadsheet that ships with this PDF has them already wired up.

#	INPUT	WHERE IT COMES FROM	YOURS
1	Annual CTC for the role	Payroll. Use total cost to company, not fixed pay.	₹
2	Notice period served	Contract, then your actual average — the two differ.	days
3	Time from vacancy to offer accepted	ATS. Median, not mean: one freak search skews the mean.	days
4	Joiner's own notice period	Offer records. In India this is usually 60–90 days.	days
5	Agency fee, as % of CTC	Your vendor agreements. Zero if you hire direct.	%
6	Months to 50%, 80% and 100% productivity	The hiring manager. Ask for three numbers, not one.	m / m / m
7	Coverage shortfall while the seat is empty	Judgment. What share of the work simply does not happen.	%
8	Months of pre-exit disengagement	Judgment, informed by exit interviews. See Part 3, line 7.	months
9	Share of voluntary exits attributable to burnout	Your own exit data. There is no credible India benchmark.	%

THE INPUT THAT DECIDES EVERYTHING

Input 9 is the one that moves your final number most, and it is the one with no published answer. Resist the urge to borrow a figure. Go through your last twenty voluntary exits and classify them yourself, with the manager who lost each person. A number you built from twenty real cases, even roughly, will beat any benchmark in the room — because you can describe how you got it.

What to do before the meeting

1. Fill in the nine inputs for your two or three most expensive-to-replace roles, not for an "average employee". Averages hide the roles that actually hurt.
2. Run each role three times — low, base and high — using the ranges in Part 3. Never walk in with a single number.
3. Write your assumptions into Part 9 and attach it. The log is what converts an estimate into an argument.

The number nobody has.

Every company can tell you its attrition rate. Almost none can tell you what a single avoidable exit cost, which is why the conversation about preventing it never gets a budget.

An attrition rate is a rear-view mirror with the resolution turned down. It tells you how many people left. It does not tell you which of them were preventable, what the preventable ones cost, or how long the warning was visible before the resignation arrived. Those three questions are the entire business case, and the standard HR dashboard answers none of them.

The information gap has a specific shape, and it is worth naming precisely, because it determines where the money is. Most organisations learn why someone left **at the exit interview** — that is, at the one moment when the information has lost all of its value. The person has already decided, already negotiated, already told their team. What you receive is an explanation, delivered to someone who can no longer act on it, by someone with no remaining incentive to be precise.

52%

of people who left voluntarily say their manager or organisation could have done something to prevent it.
Gallup, March 2019

51%

say that in the three months before they left, no manager or leader spoke to them about their satisfaction or future.
Gallup, March 2019

Read those two figures together and the structure of the problem becomes clear. The exits were mostly preventable, and the conversation that would have prevented them mostly did not happen. Not because managers do not care — because nothing told them which of their twelve people to have it with, or when.

Why this is a costing problem before it is a wellbeing problem

A people leader who says "our teams are burning out" is making a claim about human experience, and it will be received as one — sincerely, sympathetically, and without a budget line. A people leader who says "we lost seven preventable exits last year at roughly ₹19 lakh each, and here is the line-by-line build" is making a claim about capital allocation. The second claim is harder to make and much harder to dismiss. This document exists to make the second claim possible.

What is actually established, and by whom.

Five sources carry almost all of the defensible weight on this subject in India. Here is each one, what it measured, and the single sentence it entitles you to say.

SOURCE	WHAT IT MEASURED	WHAT YOU MAY CLAIM FROM IT
Aon Annual Salary Increase and Turnover Survey for India 2026–27. 33rd edition. 1,200+ organisations, 45 industries. Published 14 Oct 2026.	Overall attrition and salary movement across Indian employers.	Overall attrition rose to 18% in 2026 , from 16.2% in 2025 and 17.1% in 2024. Aon's own release calls this <i>overall</i> attrition. Press reports describing it as voluntary are wrong — do not repeat that.
Gallup State of the Global Workplace: 2026 Report, India country data. Published April 2026; data collected Jan–Dec 2025.	Engagement, daily stress and life evaluation, by country.	In India, 23% of employees are engaged (down 7 points on the prior three-year average), 18% are actively disengaged, 28% felt stress a lot of the previous day, 17% are thriving, and 53% say it is a good time to find a job where they live.
Deloitte India Mental health and well-being in the workplace, 2022. 3,995 respondents, 12 industries, 60+ expert interviews.	Prevalence, behaviour and employer cost of poor mental health.	Poor mental health costs Indian employers roughly US\$14 billion a year through absenteeism, presenteeism and attrition. 80% reported symptoms in the preceding year; 39% took no step, citing stigma. Always date it 2022.
Gallup This Fixable Problem Costs U.S. Businesses \$1 Trillion. 13 March 2019.	US voluntary turnover cost and preventability.	52% of voluntary leavers say it was preventable; 51% had no conversation about their future in their final three months. The replacement-cost range in the same article is discussed in Part 2 — treat it differently.
WHO Mental health at work, fact sheet. ICD-11 classification of burnout.	Global working days lost; the definition of burnout.	12 billion working days are lost each year to depression and anxiety, costing about US\$1 trillion in lost productivity. Burnout is an <i>occupational phenomenon</i> , not a medical condition, with three dimensions: exhaustion, cynicism, reduced efficacy.

ONE NUMBER WE COULD NOT VERIFY, AND THEREFORE DO NOT USE

Quarterly attrition disclosures from India's large IT services firms are routinely quoted in sector commentary. We tried to source them directly and failed: the figures live in investor fact-sheet PDFs that block automated access, they are absent from the corresponding regulatory filings, and secondary coverage reported four different numbers for overlapping periods. No benchmark table in this document rests on them. If you need sector comparators, open each company's fact sheet by hand and cite the document.

Is 13% of your workforce burnt out, or 59%?

Both figures are real, recent, and about India. They differ by a factor of four and a half because they measure different things. If you quote one without knowing which, the first person to quote the other will win the argument.

FINDING	FIGURE	WHAT WAS COUNTED
Great Place To Work India, reported September 2026	≈13%	Employees in the <i>highest</i> burnout category. 72% sit in the low-burnout category and about 15% in the middle. Broadly flat since 2022.
McKinsey Health Institute, November 2023 >30,000 employees, 30 countries, fielded Apr–Jun 2023	59%	Employees reporting <i>any</i> burnout symptoms, scored above 3 on a 1–5 scale across four Burnout Assessment Tool dimensions. Reported as the highest of the countries surveyed, against a global average near 20%.

Neither number is wrong. One is a severity threshold and the other is a symptom screen, and they are answering different questions. The distinction has a direct operational consequence, so it is worth holding onto:

Severity measures size the *clinical* response.
Symptom screens size the *organisational* response.

If you are budgeting counselling capacity, the 13% figure is the relevant order of magnitude. If you are deciding whether workload design, manager behaviour and meeting load need to change across the company, the 59% figure is. Using the first to argue against the second is a category error — and a common one in budget meetings.

How to state it without being corrected

The safe construction names the measure inside the claim. "Around one in eight Indian employees reports burnout at the most severe level, and a majority report at least some symptoms" is defensible, specific and survives a challenge. "59% of Indian employees are burnt out" does not, because the word "burnt out" implies the severe category and the figure came from the screen.

Two more distinctions that get people into trouble

Attrition is not turnover. Overall attrition includes involuntary exits — performance exits, restructuring, end of contract. A retention business case rests on *voluntary* attrition, which is the only part you can claim to influence. Aon's 18% is the overall figure; your voluntary share will be materially lower, and it is the one to model.

Edition year is not data year. Gallup's 2026 report carries 2025 data. Deloitte's India figure is from 2022 and is frequently quoted undated as though it were current. Saying "a 2022 Deloitte survey found" costs you four words and removes an entire category of objection.

Intent to leave is not a current Gallup finding for India. The widely circulated claim that nearly half of Indian employees want to quit traces to press coverage of the 2025 edition. The current India country page reports no intent-to-leave figure at all. Do not use it in the present tense.

Why "half to twice salary" will not survive your CFO.

The most quoted figure in retention is that replacing an employee costs between one-half and two times their annual salary. It is a reasonable hypothesis and a terrible argument, for one reason: no published method sits behind it.

The range comes from a Gallup article published on 13 March 2019, which states that "the cost of replacing an individual employee can range from one-half to two times the employee's annual salary," and describes that as a conservative estimate. The same article puts US voluntary turnover at about a trillion dollars a year. It gives no sample, no cost categories and no derivation. We read it directly to confirm this; the absence is in the source, not in the reporting of it.

That alone would be survivable. The problem is that the academic estimates of the same quantity are far lower, and they *do* publish methods:

ESTIMATE	SOURCE	COST OF REPLACING ONE EMPLOYEE
Low	Center for American Progress, 2012	≈21% of annual salary
Low	Washington Center for Equitable Growth	≈23.5% of annual salary
High	Gallup, March 2019 (no method published)	50%–200% of annual salary
Spread	Lowest credible to highest quoted	≈10×

A ten-fold spread is not a measurement. It is the absence of one. And it puts you in an impossible position in a budget meeting: if you quote 200%, a numerate CFO can halve it twice and still be inside the published literature. If you quote 21%, you have argued yourself out of your own case.

Why both ends are wrong for your purposes

The low academic figures are built mostly from recruitment and administrative cost — the things that generate invoices. They are rigorous about what they include and largely silent on what every operator knows is real: the eight months a replacement takes to reach full output, the manager who lost thirty hours to the search, the three colleagues who absorbed the work, and the four months before the resignation when the person had already left in every sense but the paperwork.

The high end includes those things implicitly, but because it never says which or how much, you cannot defend any particular number inside the band. "Somewhere between ₹16 lakh and ₹64 lakh" is not an answer to "what did it cost".

A multiplier is a conclusion. What a budget meeting needs is the arithmetic that produced it.

The rest of this document builds the number from eight defined lines. Each line is either sourced or set by you. The multiplier then falls out as a *result* — and because you can show the build, you can defend the result.

What we do keep from Gallup, and why

The 52% preventability figure and the 51% no-conversation figure are a different kind of claim: they are survey results, reported as percentages of a surveyed population, and they are the two most useful numbers in this entire field. They tell you the *addressable share* of the problem. Keep those. Retire the multiplier.

Five rules that make a cost model hold up.

These are the rules the model in Part 3 follows. They are also, in our experience, the five places where people-team business cases fall apart under questioning.

1 · Cost it per role, never per employee

An "average employee" does not exist, and the average conceals exactly what matters. A senior engineer and a support associate at the same company can differ by a factor of six in replacement cost, because CTC, time-to-fill, ramp length and single-person dependency all move together. Model your two or three most costly-to-replace roles and say so. A case built on three named roles is specific; a case built on an average is a guess with a decimal point.

2 · Count cost to company, not salary

Every formula here uses annual CTC, because CTC is what leaves the bank account. Using fixed pay understates each line by whatever your benefits, employer contributions and variable load come to. If your finance team costs people on a fully loaded basis that is higher than CTC, use theirs — the point is to use the same basis they do, so the conversation is about the model and not about the denominator.

3 · Declare every double count, and kill it

This is where most models quietly inflate. Three traps recur, and all three are addressed explicitly in Part 3: knowledge-transfer hours counted both in the notice period and again as handover; the empty seat counted both as output not produced and as work absorbed by colleagues; and the ramp period counted both as lost productivity and as a training cost. Each pair overlaps. Pick one side of each, say which you picked, and your model becomes conservative by construction.

4 · Present a range, and name the swing factors

A single number invites a single counter-number. A low-base-high range with the two or three assumptions that drive the swing invites a conversation about those assumptions, which is the conversation you want. Our worked example in Part 5 spans 0.32x to 1.52x of CTC for the same role, and the entire spread comes from four inputs.

5 · Publish what you excluded

Counter-intuitively, the exclusion list is the most persuasive page in the document. Employer-brand damage, the chilling effect on internal mobility, lost institutional memory, the opportunity cost of work never attempted — all of these are real, and none of them is in our totals, because we cannot measure them without inventing a number. Saying so converts you from an advocate into an analyst, and it is the cheapest credibility available.

THE SENTENCE TO HAVE READY

"This model excludes employer-brand damage, lost institutional knowledge and second-order effects on morale, because I can't measure them honestly. Everything in the total is either sourced or an assumption I've written down, and you can challenge any of them." Say that early. It ends the methodology argument before it starts, and moves the meeting to the decision.

The model: eight lines.

Seven of these lines are recovered only by preventing the exit. The eighth is recovered by noticing earlier, even when the person leaves anyway — which is why it is the line that justifies detection rather than just retention.

LINE	WHAT IT CAPTURES	TYPICAL SHARE
L1 · Separation and handover	Unproductive notice period, handover received, exit administration	10–20%
L2 · Vacancy and coverage	Output not produced while the seat is empty	10–15%
L3 · Recruitment	Agency fee, sourcing, recruiter and interview-panel time	10–20%
L4 · Onboarding and ramp	The gap between a new joiner's cost and their output	25–35%
L5 · Manager and team drag	Time displaced from other people's actual jobs	10–20%
L6 · Second-order attrition	The exits that follow an exit, in the same team	0–15%
L7 · Pre-exit decay	The months of reduced output before the resignation	10–15%
L8 · Role-specific exposure	Quota, named accounts, key-person risk — stated separately	Role only

Typical shares are from the worked example in Part 5 and are indicative of its assumptions, not a benchmark. Your mix will differ, and the point of running the model is to find out how.

Notation used throughout

CTC = annual cost to company for the role

W = working days a year (assumption: 240)

D = CTC ÷ W = fully loaded cost of one working day

M = CTC ÷ 12 = fully loaded cost of one month

H = D ÷ 8 = fully loaded cost of one working hour

Calendar days convert to working days at 5/7. A 60-day notice period is therefore about 43 working days. Mixing the two is the most common arithmetic error in these models and it inflates every time-based line by about 40%.

L1 · Separation and handover

You pay a full salary for a notice period during which output falls, you pay colleagues to receive a handover, and you pay HR, finance and IT to process an exit. Only the first of these is large.

The notice period is the whole line

Indian notice periods run 30, 60 or 90 days, which is long by international standards and makes this line much bigger here than in equivalent Western models. The cost is not the salary — you would have paid that anyway. It is the *shortfall*: the gap between what they would normally have produced and what they actually produce once they have decided to go.

$$L1a = \text{notice days (working)} \times D \times (1 - \text{productivity during notice})$$

Assumption to set: productivity during notice. We use 50% as a base, 70% low, 35% high. Managers consistently estimate this lower than HR does; ask the manager.

Handover received, not handover given

Here is the first double count. The departing person's handover hours are **already inside L1a** — they are part of what that person did during a notice period you have already costed in full. Counting them again as "knowledge transfer" charges you twice for the same time.

What is *not* yet counted is the time the receiving colleagues spend. That is real, additional and displaced from their own work.

$$L1b = \Sigma (\text{receiving colleagues' hours} \times \text{their } H)$$

Count only the receivers. For a role with documented systems, 12–20 hours across one or two people is typical; where one person held it for four years with nothing written down, 40–60 hours is not unusual — and that gap is itself a finding.

Exit administration

Full-and-final settlement, leave encashment, gratuity calculation, asset and access recovery, the exit interview, and the payroll work that follows. Individually trivial, collectively a day or two of skilled administrative time.

$$L1c = (\text{HR} + \text{finance} + \text{IT hours}) \times \text{their blended } H$$

8–12 hours is a reasonable default. This line is small and you should resist the temptation to pad it; its presence in the model matters more than its size, because leaving it out invites the objection that you have not thought about process cost.

What belongs here and is often misplaced

Statutory payouts — gratuity, leave encashment, notice buy-out — are generally *not* an incremental cost of the exit. Gratuity accrues across the whole tenure and is already provisioned; leave encashment pays for leave that was earned. Including them double-counts against your own balance sheet, and a finance partner will spot it. Count only the cost of processing them.

L2 · Vacancy and coverage

The seat is empty for far longer than your time-to-fill metric suggests, because in India the person you hire has to serve their own notice before they can start.

The two-sided notice problem

Almost every vacancy model uses time-to-fill — vacancy opened to offer accepted — and then stops. In the Indian market that understates the gap badly, because the accepted offer is followed by the joiner's own 60 or 90 day notice period at their current employer. The seat stays empty through both.

Gross gap = time to offer accepted + joiner's notice period

Net uncovered days = gross gap – notice days already served by the leaver

The subtraction matters: while your leaver was working out their notice, the seat was not yet empty. Counting the whole gross gap as vacancy, on top of a full L1, charges you twice for the same calendar weeks.

Worked through: a 65-day search followed by a 60-day joiner notice is a 125 calendar-day gross gap, about 89 working days. Subtract the 43 working days your leaver served, and you have **46 working days of genuinely uncovered seat** — roughly nine weeks in which the work either does not happen or lands on somebody else.

Costing the empty seat

You are not paying the absent person, so this line is not a cash cost. It is forgone output, and it is legitimate to include provided you are honest that the business did not write a cheque for it. Finance teams accept forgone output readily when it is labelled as such and refuse it when it is smuggled in as cash.

L2 = net uncovered working days × D × coverage shortfall

Coverage shortfall is the share of the role's output that simply does not get produced. Base 45%, low 30%, high 55%. If colleagues absorb much of the work, this number falls — and the cost moves to L5 rather than disappearing. See the warning below.

DOUBLE COUNT NUMBER TWO

Work that an empty seat generates goes one of two places: it does not happen, or someone else does it. The first is L2; the second is L5. If you set coverage shortfall at 60% and charge a team of four with three hours a week of absorbed work, you have costed well over 100% of one person's output. Split it deliberately: pick a shortfall, then let L5 carry only the work that was genuinely absorbed, and check that the two together do not exceed the role's full output.

The premium you pay to cover

Where coverage is bought rather than absorbed — a contractor, an agency resource, overtime at a premium, or a backfill at above-market rate because the search is urgent — that premium is a genuine cash cost and belongs here. Count only the premium over what the departed person cost, not the whole contractor invoice: you were always going to pay for that seat.

L3 · Recruitment

The only line that generates invoices, which is why it is the only line most companies can state accurately — and why most business cases stop here and lose.

Agency fees in India

Indian recruitment fees are quoted as a percentage of annual CTC. The conventional anchor is **8.33%**, which is one month's pay, and it is a market convention rather than a regulated rate. Fees rise with seniority:

ROLE LEVEL	TYPICAL CONTINGENCY FEE	NOTES
Entry and mid	8.33% – 12%	8.33% is one month's CTC. Volume commitments are the usual lever for the lower end.
Senior	12.5% – 16.67%	16.67% is two months' CTC.
Leadership and retained search	15% – 25%	Usually with an upfront retainer, payable regardless of outcome.

Source: published rate guidance from Indian staffing and search firms, checked October 2026. These are vendor-published conventions, not independently surveyed benchmarks — use your own executed agreements where you have them. A 60–90 day replacement guarantee is a standard thing to negotiate for and worth more than a point of fee.

The line your CFO has never seen

Interview panel time is almost never costed, and for engineering and senior roles it is frequently larger than the agency fee it sits beside. Work it through for one senior hire: twelve screening calls at forty-five minutes, then six candidates through a panel of three for an hour and a quarter each. That is about **32 hours of expensive technical time** per hire — before coordination, debriefs and the scheduling churn that fills a recruiter's week.

L3 = agency fee ($\% \times \text{CTC}$)
 + internal recruiter hours $\times H_{\text{recruiter}}$
 + Σ (panel hours $\times H_{\text{interviewer}}$)
 + sourcing tools and job boards, amortised per hire
 + assessments and background verification
 + referral bonus, where paid

Amortise annual licence costs — job boards, résumé databases, assessment platforms — across the hires you actually make with them, not across your whole headcount.

A useful aside for the same meeting

Because panel time is the hidden bulk of this line, it is also the line most improved by *not* having to hire. When you present the retention case, the engineering leader in the room is the one who feels L3 most sharply — they are the one whose team loses the 32 hours. It is often the most persuasive single line in the model, and it is almost always missing from the version finance has seen.

L4 · Onboarding and ramp to productivity

The largest line in almost every version of this model, and the one where the published evidence is weakest. So this page gives you a method rather than a benchmark.

We looked for a citable figure for time-to-full-productivity and did not find one we would stand behind. The numbers in circulation — three months, eight months, twelve months — are attributed variously to Gallup, to sales-benchmarking firms and to consultancies, and when traced they either dissolve into blogs citing each other or turn out to measure sales ramp in another market. The reason is structural: "fully productive" has no shared definition, so the studies are not measuring the same thing.

DO NOT BORROW A RAMP BENCHMARK

If you put "8 months, per industry research" in a business case, the first person who checks will find the citation does not hold, and your whole model goes with it. Your hiring managers know this number better than any published study. Ask them.

The three-point estimate

Ask the hiring manager three questions, not one. A single number invites a shrug, and the shape of the curve matters more than where it ends.

1. In which month does a good hire reach roughly **half** the output of an established team member?
2. In which month do they reach **eighty per cent**?
3. In which month are they genuinely **indistinguishable** from an established peer?

$$L4 = \text{onboarding cost} + \sum_{\text{months}} [(1 - \text{productivity in month } m) \times M]$$

Sketch the curve those three points imply, read off an average productivity per month, and the cost is the area above it. Onboarding cost is the direct, invoiceable part: device, licences, training and administrative hours. The summation is the ramp shortfall, usually five to ten times the onboarding cost.

Worked, for a senior engineer at ₹32 lakh CTC

The manager's three points: half productive in month 2, eighty per cent in month 5, indistinguishable in month 9. Monthly cost M = ₹2.67 lakh.

MONTH	1	2	3	4	5	6	7	8	9	TOTAL
Productivity	25%	50%	60%	70%	80%	85%	90%	95%	100%	—
Shortfall	75%	50%	40%	30%	20%	15%	10%	5%	0%	245%
Cost at ₹2.67L a month	2.00	1.33	1.07	0.80	0.53	0.40	0.27	0.13	—	₹6.53L

Figures in ₹ lakh. A 245% shortfall is about two and a half months of output you paid for and did not get, spread across nine. With ₹0.85 lakh of onboarding, L4 totals ₹7.38 lakh — roughly 23% of annual CTC from this one line.

Double count number three

Do not count formal training hours twice. The new joiner's time in training is already inside the ramp shortfall — they were unproductive and you costed it. What belongs in onboarding cost is the *other* side: the trainer's time, the licence, the course fee.

L5 · Manager and team drag

An exit does not only remove one person's output. It quietly conscripts a manager and three or four colleagues into unpaid recruitment and coverage work, displacing whatever they were hired to do.

The manager's hours

A manager losing someone spends time in three distinct phases, and it adds up to more than most people expect — frequently two full working weeks spread across a quarter, at the most expensive hourly rate in the team.

PHASE	WHAT IT INVOLVES	TYPICAL HOURS
Exit	The resignation conversation, counter-offer deliberation, informing the team, replanning committed work, handover design	6 – 10
Hiring	Role definition, scorecard, résumé review, interviews, debriefs, references, closing the candidate, salary negotiation	25 – 40
Onboarding	Ramp planning, daily availability in the early weeks, review cycles, course correction	20 – 30
Total	Across roughly one quarter	51 – 80

L5a = manager hours × H_{manager}

Use the manager's own CTC, not a generic rate. At ₹45 lakh CTC, H is about ₹2,340, so 63 hours is ₹1.48 lakh — from one exit, in time that was supposed to go to their actual job.

The team's absorbed work

This is the half of the coverage problem that L2 deliberately left out: work that did not vanish, but landed on colleagues. It is worth costing separately because it behaves differently — it is invisible in any output metric, it tends to fall on the strongest performers, and it is the mechanism by which one exit becomes two.

L5b = Σ (colleagues × hours a week × weeks × their H)

Base case: four colleagues, three hours a week, thirteen weeks. At ₹1,600 an hour that is ₹2.50 lakh. Check the constraint from L2 — the absorbed hours plus the uncovered shortfall should not exceed the role's full output.

Why this line predicts the next exit

Absorbed work is allocated by competence, not by capacity. The colleague who picks up the departed person's work is usually the one who could least afford the time, and they receive no additional pay, no title change and frequently no acknowledgement that the load is temporary — because nobody has told them when the backfill starts. That is a textbook description of how workload-driven burnout propagates through a team, and it is the direct mechanism behind L6.

L6 · Second-order attrition

Exits cluster. Everyone who has run a people function has watched one resignation become three in a quarter. We include the line, and we are explicit that we cannot give you a coefficient for it.

The mechanism described on the previous page is well understood and uncontroversial: a departure increases load on the people who remain, removes a colleague they relied on, and signals that leaving is both possible and reasonable. Where the departing person was well regarded, the signal is stronger. Where a team has already lost someone recently, the load effect compounds.

What we do not have is a defensible multiplier for India. We looked. There is no published study giving the probability that an exit in a team of a given size triggers a further exit within a stated window, and an invented coefficient here would undermine the honesty of everything else in this model.

$L6 = (L1 + L2 + L3 + L4 + L5 + L7) \times \text{contagion uplift}$

Treat the uplift as a scenario, never a benchmark, and label it as such on the page you present. We suggest: 0% for an isolated exit in a stable team; 15% where the team has lost someone else within two quarters; 30% where the exit is a manager or a long-tenured anchor in a small team.

How to present a line you cannot source

Show the base-case total *without* it, then show the scenario alongside. The framing that holds up is: "my base case excludes the clustering effect entirely. If you think one exit in a team of six makes a second exit more likely — and the mechanism is obvious — then the number is higher than this, and here is by how much under two scenarios."

That construction is stronger than a sourced coefficient would be, because it hands the judgment to the room instead of asking them to accept yours. In practice, operational leaders tend to argue for a *higher* uplift than the people team proposed, which is a good position to be arguing from.

What you can observe yourself, cheaply

You do not need a study. You need your own last eight quarters of exit data, grouped by team rather than by company. Two questions answer themselves immediately:

- What share of your voluntary exits came from teams that had already lost someone in the previous two quarters?
- How does that share compare with the share of your headcount those teams represent?

If exits are concentrated in teams well beyond their headcount share, you have evidence of clustering in your own organisation — which is worth more in the meeting than any external coefficient, and it took an afternoon.

L7 · Pre-exit decay

Nobody becomes disengaged on the morning they resign. This is the line that was accruing for months while your attrition dashboard showed nothing at all — and the only line that early detection recovers even when the person still leaves.

The evidence that this period exists is not controversial. Deloitte's 2022 India survey found that **33% of respondents continued working despite poor mental health**, against 29% who took time off and 20% who resigned. The largest group did not take leave and did not quit. They stayed, and they worked, at reduced capacity, which is precisely the state this line prices.

Gallup's India data describes the same population from a different angle: **23% engaged, 59% not engaged, 18% actively disengaged**. Those three-quarters who are not engaged are, by definition, producing below their capacity — and they are on the payroll at full cost.

L7 = months of decay × M × average productivity shortfall

Assumptions to set. Decay duration: base 4 months, low 2, high 6. Average shortfall across that period: base 25%, low 15%, high 35%. For a ₹32 lakh role at base case: $4 \times ₹2.67L \times 0.25 = \text{₹2.67 lakh}$.

Why this line changes the argument

Lines 1 to 6 are recovered only if the exit does not happen. That makes them a *retention* argument, and retention arguments have a known weakness: you cannot prove a counterfactual. The CFO asks "how do you know they would have left?" and there is no clean answer.

L7 is different in kind. It accrues whether or not the person eventually resigns, and it shrinks the moment someone intervenes — a workload change, a manager conversation, time with a professional. If you detect the decay in month one instead of discovering it at the exit interview in month five, you recover most of this line **even on the exits you fail to prevent**.

The asymmetry worth stating out loud

A retention programme that prevents no exits at all, but shortens the average decay period from four months to two, still returns half of L7 across every affected employee. That is a floor under the investment that does not depend on winning the counterfactual argument — and it is the single most useful thing in this document for a budget conversation, because it is the one claim nobody can ask you to prove with a control group.

L8 · Role-specific exposure, stated separately

For quota-carrying and client-facing roles the real exposure can dwarf every other line put together. It also attracts the most scepticism, so it goes next to the total rather than inside it.

When a mid-market account executive leaves, lines 1 to 7 come to about ₹17 lakh. The territory they carried might have been ₹2.2 crore of annual bookings. Those two numbers are not comparable, and the mistake people make is adding them.

Deferred is not lost

The honest treatment turns on one question: is the pipeline reassigned or abandoned? In almost every case it is reassigned, which means the bookings are **delayed, not destroyed**. The cost of a delay is a financing cost plus a slipped-quarter problem — real, material, and an order of magnitude smaller than the headline revenue figure.

TREATMENT	WHEN IT APPLIES	WHAT TO COUNT
Deferred bookings Recommended default	Territory is reassigned; pipeline survives the handover.	Value of bookings pushed beyond the period, stated as a cash-flow timing impact. Report the gross amount <i>and</i> make clear it is a shift, not a loss.
Lost bookings	Territory genuinely goes uncovered, or a named relationship leaves with the person.	Gross margin on the specific, nameable opportunities lost. If you cannot name them, you do not have this cost.

$$L8 = (\text{months of territory underperformance} \times \text{monthly quota} \times \text{shortfall})$$

For a ₹2.2 crore territory, five months at 50% coverage is about ₹45.8 lakh of bookings moved out of the period. Present it beneath the total, labelled "deferred revenue exposure, not included above". The discipline of excluding it is what makes the rest of your numbers believable.

Other role-specific exposures worth naming

- **Key-person dependency.** Where one person holds undocumented knowledge of a production system or a regulatory filing, the exposure is operational risk rather than cost. Describe it; do not price it.
- **Named client relationships.** In services businesses, specific accounts may be at risk. Name them or drop the line.
- **Regulated or licensed roles.** Where a vacancy creates a compliance gap, the cost is the cost of the gap, and it is usually knowable precisely.

What this model deliberately excludes


None of the following appears in any total in this document. Each is real. None of them can be measured without inventing a number, and inventing one would cost you the rest of the model.

Employer-brand damage	Review-site ratings, referral rates, offer-acceptance rates. Directionally real, not attributable to a single exit.
Lost institutional memory	The decisions nobody can reconstruct and the context that left with the person.
Morale and discretionary effort	Measurable in surveys, not convertible to rupees without a chain of assumptions long enough to break.
Work never attempted	The projects not started because the team was short. Genuinely large, genuinely unmeasurable.

The two costs that never reach the attrition report.

Everything so far prices an exit. But the Deloitte India estimate rests on three channels, and attrition is only one of them. The other two are larger, and they are running right now across people who have no intention of resigning.

Deloitte's 2022 India survey put the cost of poor employee mental health to Indian employers at roughly **US\$14 billion a year**, through *absenteeism*, *presenteeism* and *attrition*. The behavioural split in the same survey tells you how that cost divides, and the proportions are the opposite of what most companies monitor:



Deloitte India, *Mental health and well-being in the workplace*, 2022. 3,995 respondents across 12 industries. Shares are of all respondents and are not mutually exclusive.

Note which is largest. The group that stays and works through it is bigger than the group that takes leave and bigger still than the group that resigns. Presenteeism is the dominant channel — and it is the only one of the three that appears in no HR report anywhere, because the person is present, on time, and producing something.

A DENOMINATOR WARNING ON THE 20%

That 20% is the share of *all survey respondents* who said they resigned to manage their mental health — not an annual attrition rate, and not the share of resignations that are burnout-driven. It is tempting to use it as "one in five exits is a burnout exit". Do not: the denominators are different and the question did not specify a time window. Input 9 in your own model has to come from your own exit data. This is the single most common misuse of this survey.

Pricing presenteeism without inventing anything.

This is a population calculation rather than a per-exit one, and it needs only four inputs. Two of them you have; two you must set and declare.

Absenteeism = headcount × affected share × extra days a year × D

Presenteeism = headcount × affected share × impaired days a year × D × impairment

Affected shares: Deloitte India's 29% and 33% are the best available India anchors, from 2022. Extra days, impaired days and impairment level are assumptions — set them low, write them down, and let the model be challenged on them.

Worked, for 200 people at a blended ₹25 lakh CTC

Blended D = ₹10,417 a working day. We deliberately choose conservative assumptions: four extra absence days a year for the affected group, and six weeks of twenty per cent impairment for the presenteeism group.

CHANNEL	ARITHMETIC	ANNUAL COST
Absenteeism	200 × 29% × 4 days × ₹10,417	₹24.2 lakh
Presenteeism	200 × 33% × 30 days × ₹10,417 × 20%	₹41.3 lakh
Combined	On a payroll of ₹50 crore	₹65.4 lakh

About 1.3% of payroll, from assumptions chosen to be hard to argue with. Doubling the impairment estimate to 40% — still well below what "working while mentally unwell" plausibly means — takes the combined figure past ₹1 crore.

Presenteeism exceeds absenteeism here by a factor of about 1.7, which matches the pattern in the most detailed work published on the question anywhere. Deloitte's UK analysis — a different market, and quoted here only for the shape of the finding — consistently finds presenteeism to be the single largest component of employer cost, well ahead of absence.

The two global figures, scoped correctly

FIGURE	WHAT IT ACTUALLY SAYS	HOW IT IS USUALLY MISUSED
12 billion working days, US\$1 trillion a year WHO, mental health at work	Global working days lost to depression and anxiety, and the associated lost productivity.	Quoted as the cost of "burnout". It is depression and anxiety specifically, and burnout is classified separately in ICD-11 as an occupational phenomenon rather than a medical condition.
US\$4 returned per US\$1 invested Lancet Global Health commentary on WHO-led work	Return from <i>scaling up treatment</i> for depression and anxiety: about US\$399bn in productivity and US\$310bn in health, against US\$147bn of treatment cost.	Quoted as the ROI on workplace wellbeing spending generally. It is not. It is a treatment-access figure, and applying it to a wellbeing platform is a category error a health economist in the room will catch.

If you want a return figure that genuinely concerns employer-side spending, the UK evidence is the closest useful analogue — Deloitte's recurring *Mental health and employers* series reviews studies of workplace interventions and reports a positive average return. We have deliberately not reproduced the headline multiple here, because we could not open the current edition at source to confirm the figure and its basis, and a return number you cannot cite precisely is worse than none. Treat UK magnitudes as indicative at best; the method transfers, the numbers do not.

Worked example: 200 people, in rupees.

Everything from here on uses one constructed company so the arithmetic is followable end to end. Substitute your own numbers and the structure holds.

THIS COMPANY DOES NOT EXIST

Meridian Systems is a construct, built to be plausible for a Series B Indian B2B software company. It is not a Wanderfly customer, it is not an anonymised customer, and no figure below is drawn from anyone's real data. Every input is an assumption, chosen to sit in a defensible middle of the market and labelled as such. The purpose is to show the method, not to assert a benchmark.

The company

FUNCTION	HEADCOUNT	AVERAGE CTC	PAYROLL	REPLACEMENT DIFFICULTY
Engineering	110	₹28.0 L	₹30.8 Cr	High — long ramp, deep system context
Sales	30	₹22.0 L	₹6.6 Cr	High — quota and territory exposure
Customer success	25	₹12.0 L	₹3.0 Cr	Medium — named account relationships
Leadership	15	₹46.0 L	₹6.9 Cr	Very high — long search, team-wide effects
G&A and People	20	₹13.5 L	₹2.7 Cr	Medium
Total	200	₹25.0 L	₹50.0 Cr	Blended D = ₹10,417 a day

Bengaluru, Series B, 200 people. CTC figures are illustrative market-plausible assumptions for 2026, not survey data. Working year of 240 days throughout.

The two roles we cost in full

A senior backend engineer at ₹32 lakh, because engineering is 55% of headcount and has the longest ramp; and a mid-market account executive at ₹22 lakh, because it is the role where L8 changes the shape of the answer. Both are costed over the next three pages at low, base and high.

Senior backend engineer, ₹32 lakh CTC

D = ₹13,333 a day. M = ₹2.67 lakh a month. One avoidable exit, costed three ways.

LINE	BASE-CASE ASSUMPTIONS	LOW	BASE	HIGH
L1 Separation	60-day notice (43 working days) at 50% productivity; handover received by two colleagues, 32 hours; 10 hours of exit administration	2.32	3.47	4.58
L2 Vacancy	65 days to offer accepted + 60-day joiner notice = 89 working days, less 43 already served = 46 uncovered, at 45% coverage shortfall	0.44	2.76	6.31
L3 Recruitment	Agency at 8.33% of CTC; 25 recruiter hours; 32 panel hours; assessments and verification	1.24	3.54	5.11
L4 Ramp	50% productive at month 2, 80% at month 5, 100% at month 9 — a 245% cumulative shortfall; plus ₹0.85L onboarding	4.18	7.38	13.00
L5 Manager and team	63 manager hours at ₹2,340; four colleagues at three hours a week for 13 weeks	1.45	3.97	8.39
L6 Second-order	Excluded from base case. High case applies a 15% uplift to the whole cost of the exit, decay included	—	—	6.45
L7 Pre-exit decay	Four months at a 25% average productivity shortfall	0.80	2.67	5.60
L8 Role-specific	None attributable for this role	—	—	—
Total	₹ lakh	10.43	23.80	49.44
As a multiple of CTC		0.33x	0.74x	1.55x

All figures ₹ lakh. Low, base and high differ only in the assumptions listed in Part 3 — notice-period productivity, search and notice durations, coverage shortfall, ramp shape, absorbed hours and decay. No line is sourced to a benchmark; every one is either arithmetic or a declared assumption.

What the spread tells you

The same role, same company, same method, ranges from a third of annual CTC to more than one and a half times it. Four inputs account for almost all of the movement: the ramp shape, the length of the uncovered gap, whether an agency was used, and the duration of the pre-exit decay. Those four are where the conversation should go, and they are all knowable from your own systems.

A result worth noticing

The band this bottom-up model produces — roughly **0.3x to 1.6x of annual cost to company** — lands inside the half-to-twice range that Part 2 took apart. That is not a vindication of the multiplier; it is the reason the multiplier survived so long. The difference is that you can now say where the number came from, which line dominates it, and which assumption to argue about. That is the whole point.

Mid-market account executive, ₹22 lakh CTC

D = ₹9,167 a day. M = ₹1.83 lakh a month. Carrying a ₹2.2 crore annual territory — which is where this role stops resembling the last one.

LINE	BASE-CASE ASSUMPTIONS	BASE
L1 Separation	60-day notice at 40% productivity — lower than for the engineer, because a departing seller's pipeline activity falls away faster; account handover received, 24 hours	2.80
L2 Vacancy	55 days to accept + 45-day joiner notice; 28 working days uncovered at 50% shortfall	1.28
L3 Recruitment	Agency at 10% of CTC; recruiter and panel time	2.87
L4 Ramp	50% at month 3, 80% at month 6, 100% at month 10 — slower than engineering, because territory knowledge and pipeline both have to be rebuilt	5.92
L5 Manager and team	55 manager hours; sales-engineering and CS coverage of live deals	2.30
L7 Pre-exit decay	Four months at 25%	1.83
Total, lines 1–7	₹ lakh · 0.77× CTC	17.00

L8, stated separately and not added above:

₹2.2 Cr ÷ 12 × 5 months × 50% shortfall = ₹45.8 lakh of bookings deferred

This is a cash-flow timing impact on the assumption the territory is reassigned and the pipeline survives. It is not ₹45.8 lakh of lost revenue, and presenting it as such is the fastest way to lose the room. If specific named opportunities are genuinely abandoned, count gross margin on those, name them, and move them into the total.

The comparison that matters

On lines 1 to 7 the two roles are strikingly similar as a multiple of CTC — 0.74× and 0.77×. That consistency is useful: it means a blended multiple is a reasonable tool for aggregating across a workforce, provided you derived it from a bottom-up build on your own roles first.

What differs entirely is the exposure sitting beside the total. For the engineer it is zero. For the seller it is twice the direct cost, in deferred bookings, concentrated in whichever quarter the exit lands in. A CFO who does not care about ₹17 lakh will care about a quarter that slips, and that is the version of this page to put in front of them.

Where burnout risk and cost concentration meet

It is worth checking whether your highest-cost-to-replace roles are also your highest-workload roles. In most software companies they are the same people: the engineers holding system context nobody else has, and the sellers carrying the largest territories. Cost exposure and burnout exposure are correlated, which means a detection programme aimed at the heaviest-loaded teams is also aimed at the most expensive exits — and that is a better targeting argument than a company-wide average will ever give you.

What it costs Meridian in a year.

Three steps: how many people leave, how many of those were burnout-driven, and what each one costs. Only the first comes from a published source.

STEP	BASIS	LOW	BASE	HIGH
Total exits	200 × 18% overall attrition (Aon, India, 2026)	36	36	36
Voluntary exits	70% voluntary share (assumption — use your own split)	25	25	25
Burnout-attributable	20% / 28% / 35% of voluntary (judgment — input 9)	5.0	7.1	8.8
Cost per exit	0.33× / 0.74× / 1.55× of ₹25L blended CTC	₹8.2L	₹18.6L	₹38.6L
Attrition cost	Exits × cost per exit	₹0.41 Cr	₹1.31 Cr	₹3.41 Cr
Presenteeism and absence	Population model, Part 4	₹0.22 Cr	₹0.65 Cr	₹1.41 Cr
Total annual exposure	On a payroll of ₹50 crore	₹0.63 Cr	₹1.97 Cr	₹4.82 Cr
As a share of payroll		1.3%	3.9%	9.6%

Three things are worth saying about that base case of **₹1.97 crore**, or 3.9% of payroll.

- It excludes everything on the exclusion list in Part 3 — brand, memory, morale, work never attempted — so it is a floor, not an estimate of the total harm.
- Only one input in the whole chain comes from a published source. The rest are assumptions you can see and challenge, which is the model working as intended.
- The spread between low and high is a factor of 7.6. Anyone presenting a single number for this is either hiding the range or has not looked for it.

The sentence that travels

"Seven avoidable exits a year, at about ₹19 lakh each, plus roughly ₹65 lakh of reduced output from people who stayed. Call it ₹2 crore, and that is before anything I can't measure." It fits on one line of a slide, it survives being repeated second-hand, and every component of it is defined somewhere in this document.

What early detection is worth.

Now invert the model. If one avoidable exit costs ₹17.8 lakh, the question is no longer whether a programme is affordable. It is how few exits it has to prevent.

Exits that must be prevented to break even = annual programme cost ÷ cost per exit

This is the single most useful reframing available to a people leader, because it converts an open-ended spending question into a small, concrete, falsifiable target.

ANNUAL PROGRAMME COST	AT ₹8.2L PER EXIT (LOW CASE)	AT ₹18.6L PER EXIT (BASE CASE)	AT ₹38.6L PER EXIT (HIGH CASE)
₹6 lakh	0.7 exits	0.3 exits	0.2 exits
₹12 lakh	1.5 exits	0.6 exits	0.3 exits
₹24 lakh	2.9 exits	1.3 exits	0.6 exits
₹50 lakh	6.1 exits	2.7 exits	1.3 exits

Exits a year that must be prevented for the programme to pay for itself. Against a base case of seven burnout-attributable exits and an addressable pool of roughly 13 preventable voluntary exits, the thresholds in the middle column are small enough to be testable within a year.

The honest ceiling on what you can claim

Do not model prevention at 100%. Gallup's finding is that **52%** of voluntary leavers believed their exit was preventable — which is simultaneously the strongest argument for doing something and a hard limit on what any programme can claim. Of Meridian's 25 voluntary exits, about 13 sit in the addressable pool. Of the seven burnout-attributable ones, a programme that reached half of them would be doing well.

**Realistic case: 7.1 burnout exits × 52% addressable × 50% successfully intervened
= 1.8 exits prevented = ₹34 lakh recovered a year**

That is a deliberately pessimistic chain — two independent 50% haircuts on top of a conservative per-exit cost. It still clears a ₹12 lakh programme by nearly three times over. Presenting the pessimistic chain first, and then noting that the upside is four times larger, is far more persuasive than leading with the upside.

The value of knowing sooner.

Prevention is the headline benefit and the hardest to prove. Lead time is the quieter one, and it accrues on the exits you fail to prevent — which makes it the floor under the whole investment.

Recall the structure from Part 3. Lines 1 to 6 are binary: either the person stays and you avoid all of it, or they leave and you pay all of it. Line 7, the pre-exit decay, is different — it is proportional to *how long the decay ran before anyone acted*. Shorten that window and you recover part of the line regardless of the eventual outcome.

DETECTION LEAD TIME	DECAY WINDOW	L7 PER EXIT	RECOVERED	ACROSS 7 EXITS
None — found out at the exit interview	4 months at 25%	₹2.08 L	—	—
14 days	3.5 months at 25%	₹1.84 L	₹0.24 L	₹1.7 L
30 days	3 months at 25%	₹1.56 L	₹0.52 L	₹3.7 L
60 days	2 months at 25%	₹1.04 L	₹1.04 L	₹7.4 L

Using the blended ₹25 lakh CTC, so M = ₹2.08 lakh. Recovery assumes intervention arrests the decay rather than reversing it, which is the conservative reading. Nothing here assumes a single exit was prevented.

Seven and a half lakh a year from lead time alone is not a dramatic number, and it should not be presented as one. Its value is argumentative rather than arithmetic: it is the part of the case that **does not require you to win the counterfactual**. When the CFO asks how you can possibly know those people would have left, the answer is that for this component you do not need to know — the decay was happening, you measured it, and you shortened it.

Why lead time has to be built in, not added later

A lead time of 14, 30 or 60 days is only meaningful if something is measuring continuously. An annual engagement survey cannot produce it — by the time the results are collated and read, the window has closed, and the person whose scores moved has either recovered or resigned. This is the practical reason detection and intervention get discussed together: the measurement cadence *is* the lead time, and a yearly cadence has none.

What to do with the two numbers together

Present prevention as the case and lead time as the floor. "If this works as intended we prevent roughly two exits a year, worth ₹34 lakh. If it prevents none at all but shortens the average decay by two months, it still returns ₹7.4 lakh. The programme costs ₹12 lakh." That is a range with a defensible downside, and ranges with defensible downsides are what get funded.

The one-page business case.

Copy this structure. It is deliberately short, it leads with the ask, and it puts the methodology where it belongs — available on request, not in the first paragraph.

1 · The ask, in one sentence

₹_____ a year for _____, to reduce avoidable, burnout-driven attrition in _____.

2 · The problem, in two

We lose about _____ people a year voluntarily. Roughly _____ of those exits show burnout or workload as a primary driver in our own exit data, and we currently learn about it at the exit interview, when nothing can be done.

3 · The three numbers

Cost of one avoidable exit, bottom-up: ₹_____ lakh (range ₹_____ to ₹_____).

Annual exposure from burnout-attributable exits: ₹_____ lakh.

Exits we must prevent to break even: _____.

4 · How the number was built

Eight defined cost lines, per role, not a salary multiplier. Assumptions logged and attached. Published sources used for attrition rate and preventability only; everything else is our own data or a stated assumption.

5 · What I have deliberately excluded

Employer-brand damage, lost institutional knowledge, morale effects and work never attempted. All real, none measurable without inventing a figure.

6 · How we will know whether it worked

Baseline: _____. Measured by: _____. Compared against: _____. First review: _____. Kill criteria: _____.

7 · The decision I need

Approve / decline ₹_____ for _____ months, with a review at _____.

PUT SECTION 6 IN BEFORE YOU ARE ASKED

A business case that specifies in advance how it could be shown to have failed is read completely differently from one that does not. It is also the section most people leave out, because it feels like handing over a weapon. It is the opposite: it is the difference between a request for belief and a proposal for an experiment, and only one of those gets renewed.

The four questions, and what to say.

In our experience a CFO asks these in roughly this order. All four are fair, and all four have good answers if you have done the work in Parts 3 to 5.

1 · "Where did the cost-per-exit number come from?"

The answer is never a citation; it is a build. "Eight lines, per role. The biggest is the ramp period, which my engineering lead estimated at nine months to full productivity — that alone is ₹6.5 lakh. The second biggest is the uncovered seat, which runs 46 working days because the person we hire has to serve their own notice. Here is the table." Hand over the table. The question does not come back.

2 · "How do you know burnout caused it rather than pay?"

Be straight: you do not, with certainty, and attribution is the weakest link in the chain. What you have is your own classification of your own last twenty exits, done with the managers who lost those people. Give the share, give the method, and give the sensitivity — "if I'm wrong by a third in either direction, the case still clears break-even." A stated uncertainty you have already quantified is not a vulnerability.

3 · "Would they have left anyway?"

This is the counterfactual question and it cannot be won head-on. Do not try. Split the claim instead: "For prevention, I can't prove the counterfactual — which is why I've haircut it twice and the case still works. For the pre-exit decay, I don't need to: that output was being lost while they were still employed, and detecting it earlier recovers part of it whether or not they stay." Part 6 exists for this moment.

4 · "Why not just pay people more?"

Treat it as a genuine alternative, because it is one, and compare them honestly. A 5% across-the-board increase on Meridian's ₹51.5 crore payroll is ₹2.6 crore a year, permanent and non-reversible. The detection programme is ₹12 lakh, time-bound and cancellable. They are not substitutes — pay fixes a pay problem — but if the exits in question cite workload, manager behaviour or exhaustion rather than compensation, then a pay rise is a very expensive instrument aimed at the wrong target. Your exit data tells you which problem you have.

The one thing not to do

Do not answer any of these four with an external statistic. The moment the conversation becomes a contest of citations, you have lost the advantage you actually hold — which is that you know your own attrition, your own ramp times and your own exit interviews, and nobody else in the room does. The published sources in this document exist to establish that the problem is real and widespread. Your own data is what establishes that it is yours.

How you will know whether it worked.

Decide this before you spend anything. A retention programme with no pre-agreed measurement plan will be judged twelve months later on whatever the attrition number happens to say — which is almost entirely noise at the scale most companies operate at.

Why the headline attrition rate cannot be your metric

At 200 people with 25 voluntary exits a year, preventing two exits moves voluntary attrition from 12.5% to 11.5%. That one-point move is well inside the normal year-to-year variation driven by hiring market conditions, a single team reorganisation or one manager leaving and taking people with them. You would be unable to distinguish a successful programme from a quiet year, in either direction.

THE TRAP

If the company-wide attrition rate is the agreed measure, you have accepted a test your programme cannot pass and cannot fail on its merits. Negotiate the measure at the point of approval, not at the point of review.

What to measure instead

MEASURE	WHY IT WORKS	CADENCE
Lead time to intervention	Days between a risk signal appearing and a documented action. This is the programme's actual mechanism, it is fully within your control, and it moves within weeks.	Monthly
Coverage	Share of teams with current data. A detection programme that nobody uses fails for a reason worth knowing early.	Monthly
Risk-score trajectory in treated teams	Did the measured risk actually fall after intervention? Directional, team-level, and the most direct evidence available.	Monthly
Regretted voluntary attrition in treated vs untreated teams	The outcome measure. Restricted to regretted exits, compared between groups rather than against last year.	Two quarters
Burnout-attributable share of exits	Your input 9, re-measured. If the programme works, this share should fall even if the total does not.	Two quarters

A design that is actually achievable

A randomised trial is not available to you — you cannot withhold support from half your workforce, and you should not want to. What is available is a **staggered rollout**: begin with the three or four teams where cost exposure and workload are highest, keep the rest as a comparison group for one or two quarters, then extend. That gives you a difference-in-differences comparison, it is ethically unobjectionable because everyone is reached eventually, and the sequencing is defensible on cost grounds alone.

Be honest about its limits when you present it. The teams you choose first are the ones with the worst problem, so they have the most room to improve — regression to the mean will flatter you. Say so before someone else does, and the finding that survives is more credible for it.

Pre-register four things: the baseline, the comparison group, the review date, and what result would make you stop.

Write them into the approval document. The fourth is the one that earns you the benefit of the doubt on everything else.

Your assumptions log.

Fill this in and attach it to the business case. It is the document that converts your estimate into an argument, and it takes about twenty minutes.

#	ASSUMPTION	LOW	BASE	HIGH	WHERE YOUR VALUE CAME FROM
1	Annual CTC for the role	₹	₹	₹	
2	Working days a year	240	240	240	
3	Notice days served (working)				
4	Productivity during notice	70%	50%	35%	
5	Handover hours received				
6	Days to offer accepted				
7	Joiner's notice period (days)				
8	Coverage shortfall	30%	45%	55%	
9	Agency fee, % of CTC				
10	Internal recruiter hours				
11	Interview panel hours				
12	Month reaching 50% productivity				
13	Month reaching 80%				
14	Month reaching 100%				
15	Onboarding direct cost	₹	₹	₹	
16	Manager hours across the cycle				
17	Colleagues absorbing work × hours × weeks				
18	Second-order uplift	0%	0%	15%	
19	Months of pre-exit decay	2	4	6	
20	Average decay shortfall	15%	25%	35%	
21	Voluntary share of total attrition				
22	Burnout-attributable share of voluntary exits				
	Cost per exit	₹	₹	₹	÷ CTC = ___ ×
	Annual exposure	₹	₹	₹	___ % of payroll

Defaults shown are the base-case values used in Part 5, offered as starting points only. Every cell in the final column matters more than the number beside it: "my engineering lead's estimate" and "industry research" are read very differently.

Sources, and what each one supports.

Every externally sourced figure in this document, with the date, the sample, and — more usefully — the claim it does *not* support.

SOURCE	SUPPORTS	DOES NOT SUPPORT
Aon, Annual Salary Increase and Turnover Survey for India 2026-27 33rd edition · 1,200+ organisations, 45 industries · published 14 Oct 2026 · aon.com	Overall attrition of 18% in 2026, 16.2% in 2025, 17.1% in 2024. Salary growth of 9% actual in 2026 and 9.1% projected for 2027.	Voluntary attrition — the release reports overall attrition only. Sector-level attrition splits, which are not in the release we read.
Gallup, State of the Global Workplace: 2026 Report , India country data Published April 2026 · data collected Jan–Dec 2025 · gallup.com	India: 23% engaged, 18% actively disengaged, 28% daily stress, 17% thriving, 53% say it is a good time to find a job.	Burnout — Gallup does not report a burnout measure. Intent to leave, which the current India page does not report at all.
Gallup, This Fixable Problem Costs U.S. Businesses \$1 Trillion 13 March 2019 · gallup.com	52% of voluntary leavers say their manager or organisation could have prevented it; 51% had no conversation about their future in their final three months.	Replacement cost. The half-to-twice-salary range in this article is published without method, sample or cost categories — see Part 2. It is US data throughout.
Deloitte India, Mental health and well-being in the workplace 2022 · 3,995 respondents, 12 industries, 60+ expert interviews	~US\$14 billion annual cost to Indian employers via absenteeism, presenteeism and attrition. 80% reported symptoms in the preceding year; 39% took no step, citing stigma; 33% worked through it, 29% took leave, 20% resigned.	A current figure — it is 2022 data. An annual burnout-attrition rate: the 20% is a share of all respondents with no stated time window, not a share of resignations. See the warning in Part 4.
McKinsey Health Institute, Reframing employee health November 2023 · >30,000 employees, 30 countries · fielded Apr–Jun 2023	India at 59% reporting burnout symptoms, the highest of the countries surveyed, against a global average near 20%. Measured as a score above 3 on a 1–5 scale across four Burnout Assessment Tool dimensions.	Severe burnout prevalence — this is a symptom screen, not a severity threshold. Corroborated across outlets but we could not read the country-level figure in the primary PDF; treated here as corroborated rather than verified.
Great Place To Work India Reported September 2026	~13% of employees in the highest burnout category, 72% in the low category, ~15% in the middle; broadly flat since 2022.	Any comparison with symptom-screen figures without stating the difference in measurement. Sample size and method were not available in the coverage we could read.

Continued overleaf.

Sources, continued.

SOURCE	SUPPORTS	DOES NOT SUPPORT
WHO, <i>Mental health at work</i> fact sheet, and ICD-11 who.int	12 billion working days lost annually to depression and anxiety, costing ~US\$1 trillion in lost productivity. Burnout as an occupational phenomenon with three dimensions, not a medical condition.	A cost of burnout specifically — the figure covers depression and anxiety. Any employer-level or India-level cost.
Lancet Global Health commentary on WHO-led work	US\$4 returned per US\$1 invested in <i>scaled-up treatment</i> for depression and anxiety: ~US\$399bn productivity and ~US\$310bn health against ~US\$147bn cost.	Return on workplace wellbeing programmes. This is a treatment-access figure and does not transfer to employer-side spending — see Part 4.
Deloitte UK, <i>Mental health and employers</i> 2024 edition	The qualitative finding, consistent across editions and widely corroborated, that presenteeism is the largest single component of employer cost — larger than absence.	Any specific cost total or return multiple. We could not open the current edition directly, so no figure from it is reproduced here. Indian magnitudes in any case: the method transfers, the numbers do not.
Center for American Progress (2012); Washington Center for Equitable Growth	Published, method-bearing replacement-cost estimates of ~21% and ~23.5% of annual salary.	Indian cost structures, and the lines this model adds — ramp, manager drag, pre-exit decay — which these estimates largely exclude.
Indian recruitment fee conventions Published rate guidance from Indian staffing and search firms, checked Oct 2026	8.33% of CTC as the conventional contingency anchor, rising to 12.5–16.67% for senior roles and 15–25% for retained search.	An independently surveyed benchmark — these are vendor-published conventions. Use your own executed agreements where available.

Deliberately absent. Quarterly attrition disclosures from Indian IT services firms (unverifiable without opening each investor fact sheet by hand; secondary coverage reported four conflicting figures). Any published time-to-full-productivity benchmark (no traceable primary — Part 3, line 4 gives a method instead). Any coefficient for second-order attrition (no published study — Part 3, line 6 treats it as a scenario). The claim that nearly half of Indian employees want to quit (traces to 2025–edition press coverage; not a current finding).

If you are evaluating a detection platform.

Ten questions worth asking any vendor, including us. They are ordered so that the ones that disqualify come first.

1. **Does it watch people or measure patterns?** Screenshots, keystrokes and activity tracking are surveillance. They damage the trust that any honest signal depends on, and they tend to make the problem worse.
2. **What is the smallest group it will report on?** If a team of three gets an individual-level readout, the tool is identifying people, whatever the label says. Ask for the floor and ask what happens below it.
3. **What does the employee see about themselves?** Anything measured about someone that they cannot see is being done to them rather than with them.
4. **Where does the data sit, who can read it, and for how long?** Ask for retention periods and the access list in writing, with your Legal team in the room.
5. **What is the measurement cadence?** This is your lead time, and Part 6 shows what it is worth. An annual survey cannot produce one.
6. **What happens after a signal?** A risk score with no intervention path attached is a dashboard, and dashboards do not change outcomes.
7. **Is there a route to a human professional?** Self-serve content is not a response to someone in genuine difficulty.
8. **How will we know it worked?** If the vendor cannot answer this better than you can, you will be the one explaining the renewal.
9. **What does it cost at our headcount, in rupees, for a year?** Then divide by your cost per exit from Part 5, and you have your break-even in exits.
10. **What does the vendor refuse to claim?** The most useful question of the ten. Anyone who claims everything has thought about none of it.

About Wanderfly

Wanderfly is an employee wellbeing platform built for Indian companies, around one idea: that people leaders should find out about burnout risk before the exit interview, not during it. We call the method **Detect → Intervene → Retain**. Detection is team-level and pattern-based, never individual surveillance. Intervention runs two ways — initiatives for the organisation, and vetted professionals for the person who needs one.

What we will not claim: that we can predict who will resign, that a risk score replaces a manager's judgment, or that any of the arithmetic in this document becomes easier because you bought software. The model in here works regardless of what you choose, and we would rather you ran it.

Edition one, October 2026. All figures checked 26 October 2026. Written by the Wanderfly team in Bengaluru. The worked example throughout is a construct and not a customer. Corrections are genuinely welcome: info@wanderfly.in

Every calculation in this document also ships as a live spreadsheet — change the inputs, watch the total move — downloaded alongside this PDF at wanderfly.in/resources. Companion guides, including the attrition rate calculator, are at wanderfly.in/guides