Topic 08

University Business

What this course is, and how to use it

This is the eighth course in University Business: **organisational behaviour and HRM** – the study of how people actually behave at work, and how to build people processes that are fair, measurable and honest about their evidence.

The running case is **Nimbus Kitchen's front-of-house crew**: eight people, one Saturday rush, and the company's first hires as it grows from kitchen to brand. Every unit works the same crew forward – perception, motivation, teams, leadership, selection, performance, inclusion – with the survey, funnel and meeting numbers recomputable from the crew sheet at the end.

- Units 8.1–8.3 take the individual into the team: perception, motivation, and the design of working groups.
- Units 8.4–8.6 build the system: leadership, culture, selection, development, performance and change.
- Unit 8.7 audits it all: bias, inclusion and ethics with evidence.

How to read each unit: the question first, then the theory with its evidence honestly graded, then two escalating worked examples on the crew, the mistakes people actually make, and a three-line summary.

Individual behaviour and perception

Why it matters. The Saturday rush looks different from the pass and from the oven – and the manager judges what she sees. Perception is not a camera; attribution is not a court. Both can be audited.

The idea, with a picture. A customer films a chef snapping at a waiter; the clip goes round the office. Asked "why?", 25 of 40 staff (62%) say **personality** – "he's abrasive"; 15 (38%) say **circumstances** – "the ticket printer jammed during a double-booked hour". Same clip, two attributions.



Three questions decide an attribution - and the fundamental attribution error is skipping two of them

The definitions. **Selective attention** and **perceptual bias** (halo, recency, contrast) shape what is noticed. **Attribution theory** : behaviour is judged on **consistency** (does the person act this way over time?), **distinctiveness** (only here, or everywhere?), **consensus** (does everyone act this way here?). Low distinctiveness + low consensus → person; the reverse → situation. The **fundamental attribution error** is over-weighting personality and skipping the situation.

The derivation. The chef's ledger: consistency high (short on three of four Saturdays), distinctiveness LOW (also short at the farmers' market stall), consensus LOW (the calm colleague on the same tickets) – the cube points at the person ✓. But add the withheld fact – the printer jammed on all three occasions – and the situation returns. Attribution arithmetic is only as good as its evidence list; the audit question is always "what would I need to see to flip the verdict?"

> **Worked example 1 – the halo audit.** The crew rates the new waiter "brilliant" after one flawless Saturday. The manager checks the record: great with tables, poor at stock (three stock-outs). The halo effect – one vivid strength grading the whole person – had already priced him. Repair: rate the criteria separately (unit 8.6's appraisal is built for this), and the halo loses its seat at the table.

> **Worked example 2 – escalate to stereotype risk.** The oldest crew member is "slow with the tablet" – one instance, one busy day. Calling it an age fact commits the ultimate attribution error (a stereotype acting as a person-verdict). The evidence: two orders mis-keyed in month one, none in month two after a five-minute training. Track *trends within the person* before invoking *categories about the group* – and note the cost of the alternative: stereotyped staff disengage precisely when you need them learning.

Common mistakes. Judging on the vivid instance (recency). Grading the whole

person from one trait (halo). Explaining others by personality and yourself by circumstances (the actor–observer gap).

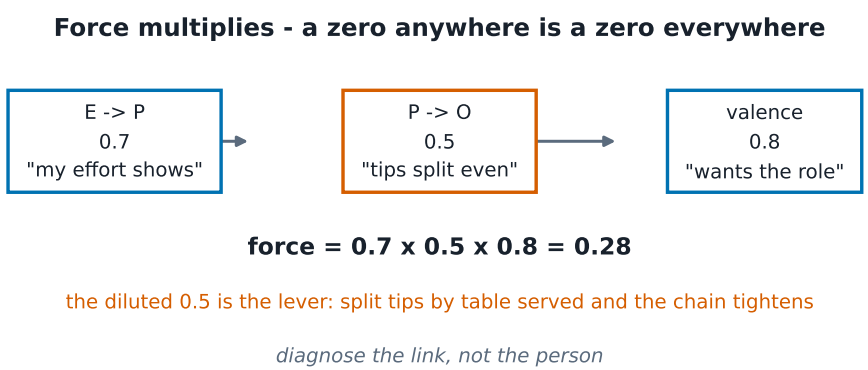
Summary. Perception selects; attribution explains. Consistency, distinctiveness, consensus – ask all three or the verdict is a mood.

Motivation

This unit is about **motivation** .

Why it matters. The Saturday rush is run on discretionary effort – the ten percent nobody can check. Motivation theory is the design language for that effort: what starts it, what directs it, and what sustains it.

The idea, with a picture. Two families: **content** theories ask WHAT motivates (Maslow's ladder, Herzberg's two factors – hygiene prevents dissatisfaction, motivators create satisfaction); **process** theories ask HOW motivation happens. **Expectancy theory** : force = $E \to P \times P \to O \times$ valence.



Force multiplies: any link near zero kills the drive - find the broken link, not the lazy worker

The definitions. Maslow: physiological → safety → belonging → esteem → self-actualisation (a lens, weakly evidenced). **Herzberg's two-factor theory** : pay and conditions are hygiene – their absence demotivates, their presence does not motivate; achievement, recognition, growth are motivators. **Goal-setting theory** : specific, difficult, accepted goals outperform "do your best" – the most replicated finding in OB.

The derivation. The tip-share scheme: if the crew believes effort leads to performance ($E \to P = 0.7$ – staffing makes it reachable), performance leads to the outcome ($P \to O = 0.5$ – tips are split evenly, so the link is diluted), and the outcome is valued (valence = 0.8), then force = $0.7 \times 0.5 \times 0.8 = \mathbf{0.28}$ ✓. Multiplication is the theory's discipline: a zero anywhere is a zero everywhere. The diluted 0.5 is the design lever – split tips by table served and the chain tightens.

> **Worked example 1 – find the broken link.** A quiet, skilled cook stops volunteering for specials. The manager assumes laziness; the chain says measure: $E \to P$ fine (he can cook them), $P \to O$ broken (specials earn team tips, not recognition), valence fine (he wants the sous-chef role). The repair is recognition with a name on it – a specials menu credit – not a pep talk. Process theory turns "motivate them" into "fix the link".

> **Worked example 2 – escalate to goals.** Two Saturday crews: one set "do your best on covers", the other "95 covers, no order over 12 minutes, checked at 2 pm and 6 pm". Stretch-goal crew: 14 of 20 shifts hit the targets; do-your-best crew: 8 of 20 hit the (unspecified) standard. Specific + difficult + feedback – and the acceptance condition is where it earns its honesty: a goal the crew did not accept is an instruction, and instructions do not carry Saturday energy.

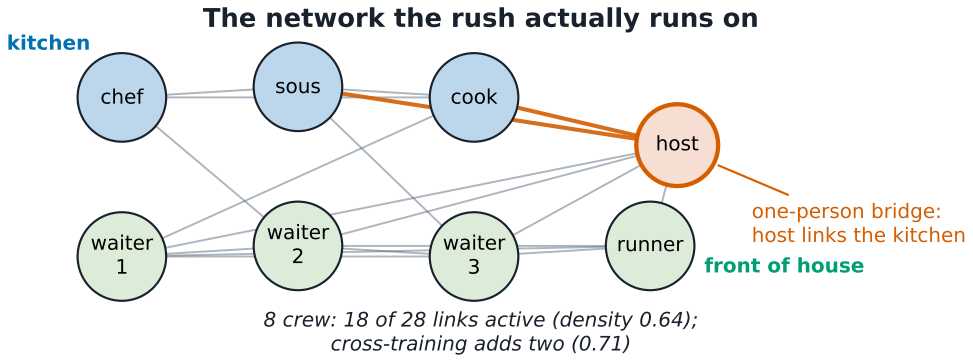
Common mistakes. Treating pay as a motivator when the evidence says hygiene. Setting vague goals ("improve service"). And diagnosing motivation problems as people problems before auditing the expectancy chain.

Summary. Content says what, process says how. Force = $E \to P \times P \to O \times V$ – find the broken link. Specific, difficult, accepted goals win.

Teams

Why it matters. Eight people staff the Saturday rush; the rush is won or lost as a system, not as eight individuals. Team design decides before the first ticket prints.

The idea, with a picture. The crew's sociogram: 18 of the 28 possible communication links are active (density 0.64) – and while a few waiter-to-cook shortcuts exist, both pass-critical links (host–cook and host–sous) run through one person. That bridge is a single point of failure wearing a smile.



Density 0.64 with a one-person bridge - the network the rush actually runs on

The definitions. Tuckman's forming–storming–norming–performing (a map, not a

law). **Role theory** (Belbin): teams need complementary roles, not clones. **Psychological safety** : the shared belief that raising problems is safe – the strongest single predictor of team effectiveness in Google's Project Aristotle. **Social loafing** : effort drops when output is unidentifiable.

The derivation. The meeting audit: in twelve pre-service meetings, three distinct voices spoke – 3 of the 8 crew (37.5%). After the round-robin change ("each name, one blocker, thirty seconds"), voices per meeting rose to seven of the eight and stock-outs fell from three to one a month. The mechanism: the loafing fix (identifiable contribution) IS the safety intervention – a structure that makes speaking routine.

> **Worked example 1 – design before personality.** Two strong personalities collide on plating. The resolution is not mediation first; it is role clarity: one owns tickets, one owns pass quality, the interface is a checklist. Belbin's point: conflict often signals two people cast for the same role – redesign the cast, not the characters.

> **Worked example 2 – escalate to the bridge.** The bridge person takes a festival Saturday off; the crew mis-routes nine orders. The network fix: cross-train one waiter and one cook as a second bridge (adds two links, density to $20/28 = 0.71$), and the checklist makes the hand-off explicit. Networks are infrastructure – they need redundancy like any other infrastructure.

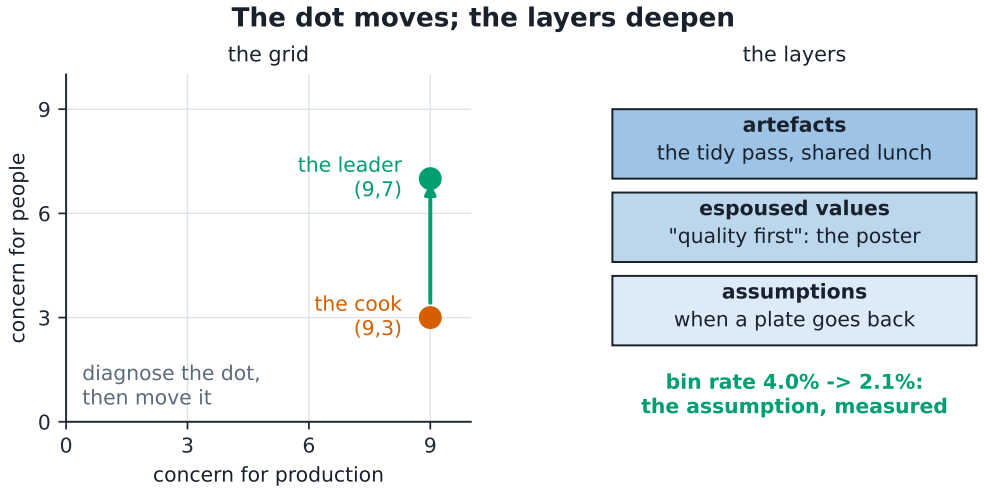
Common mistakes. Building teams of clones ("all stars") and losing the complements. Mistaking harmony for safety – safe teams disagree out loud. Ignoring social loafing until the annual review discovers the invisible shirker.

Summary. Design roles complementary, links redundant. Make contribution identifiable – loafing and silence share a cure. Safety is measured in voices, not vibes.

Leadership and culture

Why it matters. The founder can cook, sell and schedule – and at twelve staff, the job changes from doing to enabling. Leadership theory is the map of that transition, honestly graded by evidence.

The idea, with a picture. Theory families as a grid: task versus relationship behaviours (the managerial grid), then the modern evidence: **transformational** leadership (vision, individual consideration, intellectual stimulation) outperforms transactional bookkeeping – with effect sizes that are real but modest, and the caveat that most studies are correlational.



The dot moves; the layers deepen - culture is what the crew does when the founder is off

The definitions. Managerial grid: concern for production × concern for people; the evidence favours high–high, situational leadership says match the follower's readiness. **Culture** in three layers: **artefacts** (the tidy pass, the shared lunch), **espoused values** ("quality first"), **basic assumptions** (what the crew assumes without asking). Culture is read at the deepest layer that shows.

The derivation. The founder's audit: she scores 9,3 on the grid (nine task, three relationship) – the cook's score. The crew's engagement survey says the strongest driver is "someone notices my development": the dot must move toward 9,7. The behaviour priced concretely: one development conversation a fortnight, diarised like a supplier order – because what gets scheduled gets led.

> **Worked example 1 – situational honesty.** The newest waiter needs directing (new task, low readiness); the sous-chef needs delegating (seasoned, high readiness). Applying the founder's 9,3 everywhere over-directs the sous-chef into quiet disengagement and under-directs the new waiter into anxiety. Situational leadership's prescription (getting the readiness diagnosis right) is common sense made explicit – its evidence base is thinner than its popularity, and the honest sentence says so.

> **Worked example 2 – escalate to culture arithmetic.** Espoused value: "quality first". Artefact: waste bin audit shows 4% of mains binned for defects; assumption test – what happens when a waiter sends back a late plate? If the chef blames, the assumption is "speed hides errors" whatever the poster says; if the chef thanks, the assumption is catching. Culture change starts at assumptions by changing what gets praised – the bin rate is a culture metric, and its trend (4.0% → 2.1% over a quarter) is the change measured.

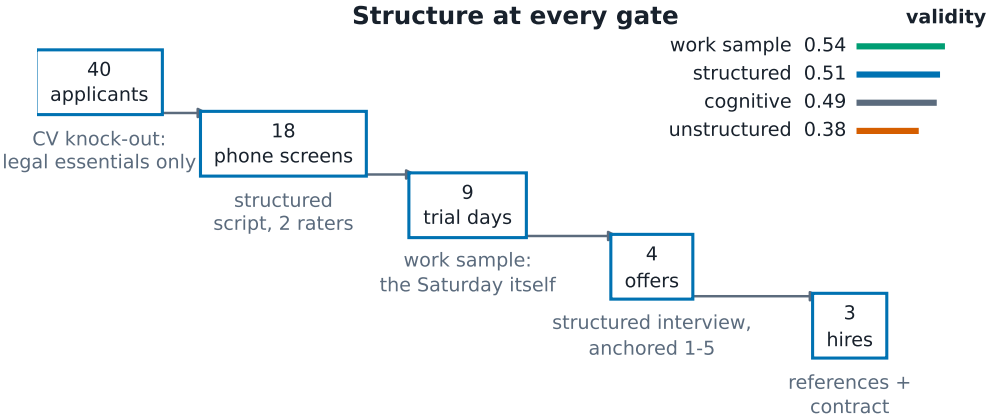
Common mistakes. Anointing one theory as the answer (each explains a slice). Reading culture from the poster (artefacts) instead of the behaviour under pressure. And leadership training without scheduled behaviour change – inspiration decays like milk.

Summary. Grid for the diagnosis, situational for the match, transformational for the evidence (modest, correlational). Culture = artefacts, values, assumptions – change what is praised.

Recruitment and development

Why it matters. Every people problem downstream is cheaper to solve at selection than at correction. Nimbus's first structured hire builds the funnel that the whole course can audit.

The idea, with a picture. The funnel: 40 applicants → 18 phone screens → 9 trial days → 4 offers → 3 hires. Each arrow is a rate, and the funnel's honesty depends on the criteria at each gate – structure beats instinct by half again in every meta-analysis.



Structure beats instinct at every gate - and the work sample is the strongest single predictor

The definitions. Predictor validity (correlation with job performance): work samples ≈ 0.54, structured interviews ≈ 0.51, cognitive tests ≈ 0.49, unstructured interviews ≈ 0.38 (the illusion of insight). **Base rate**: the share of strong performers in the pool (0.2 here). **Reliability** (consistency) precedes validity (a ruler that changes length measures nothing).

The derivation. With base rate 0.20 and a test of validity $r = 0.4$, the top-quartile selection lifts the expected share of strong performers to about 0.35 – the maths of why a mediocre test plus a cut still beats instinct. The funnel's design choice: work-sample trial day (validity ≈ 0.5+) as the widest expensive gate, CVs only as a knock-out for legal requirements. Cost arithmetic: trial day costs £90 a candidate × 9 = £810; one mis-hire costs a season of Saturday chaos (recruitment rule of thumb: 30–50% of salary) – the funnel is an investment with a number.

> **Worked example 1 – structure the interview.** Old question: "Are you good under pressure?" (everyone says yes; answers legally risky). A **structured interview** asks: "Tell me about a Saturday that went wrong – what did you do first?" – same behaviour probe for every candidate, scored on a 1–5 anchored scale by two interviewers independently. The validity jump (0.38 → 0.51) is bought with a question sheet, not a personality.

> **Worked example 2 – escalate to development ROI.** The stock-out problem traces to a training gap. Design: two hours on the ordering system, then measure the gap closed – stock-outs fall from three a month to one, each worth about £40 of lost margin: £80 a month, £960 a year against £180 of training time. ROI honestly stated: the transfer conditions (practice on real Saturdays, a checklist that survives the rush) are where training dies; the evaluation asks whether the BEHAVIOUR changed, not whether the room was warm.

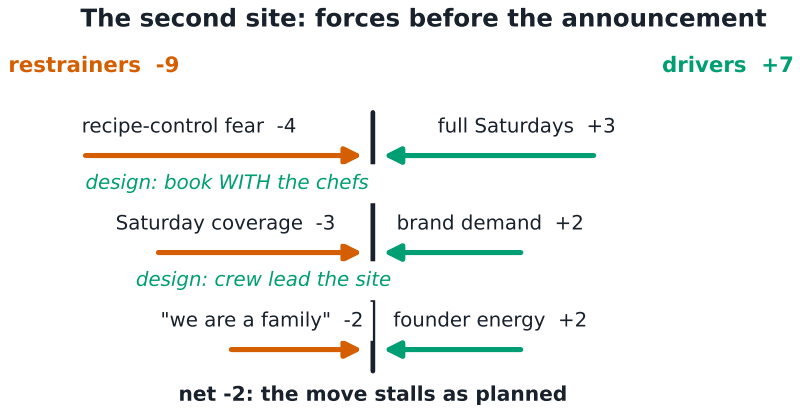
Common mistakes. Unstructured interviews defended by gut ("I know in five minutes" – the five-minute illusion). Validity claimed without a criterion to correlate against. And training evaluated by satisfaction sheets while the behaviour never moves.

Summary. Funnel with structure at every gate; work samples lead validity. Base rate × validity decides the cut's value. Development ROI = measured behaviour change.

Performance and change

Why it matters. Performance appraisal is where fairness becomes visible – or not. And the move to a second site is the founder's first real change programme: resistance guaranteed, information guaranteed too.

The idea, with a picture. Appraisal with measurable criteria: each crew member rated on 3–6 agreed measures (covers served, defect rate, development goal) by two raters independently – the design's fairness is in its redundancy. For change: **force-field analysis** – driving forces against restraining forces, listed, weighted and drawn.



Equilibrium is a balance of forces - the change plan either adds drivers or removes restrainers, never both by accident

The definitions. Rating **errors**: halo, central tendency, recency, leniency – controlled by anchored scales and multiple raters. **BARS**: behaviours anchored to numbers. Change: Lewin's **unfreeze–change–refreeze** ; resistance as information about forces the plan missed, not as insubordination.

The derivation. The second-site force-field: drivers – full Saturday demand (+3), brand pull (+2), founder energy (+2); total +7. Restrainers – recipe-control fear (−4), Saturday coverage worry (−3), "we're a family" identity (−2); total −9. Net −2: the move stalls unless the plan removes restrainers. The designed move: write the recipe book WITH the chefs (removes −4 by converting control to authorship), promise the founding crew the second site's leads (removes −3). Net becomes +5 – and the arithmetic says why "push harder" alone fails: adding +2 of founder pressure against unchanged fears still nets zero.

> **Worked example 1 – the recency rescue.** The waiter's year: strong spring, midsummer wobble, strong autumn. Recency-biased raters score the autumn; the anchored diary (monthly one-line ratings) scores the year. The difference changed one rating band – and one pay decision. Measures defend the rater from themselves.

> **Worked example 2 – escalate to resistance as data.** The sous-chef's resistance to the recipe book turns out to be knowledge: unwritten variations the book would erase. The design changes – the book gains a "house variations" page with his name on it – and the strongest restrainer becomes a driver. Resistance audited is free consulting; resistance crushed is the same information purchased later at the price of trust.

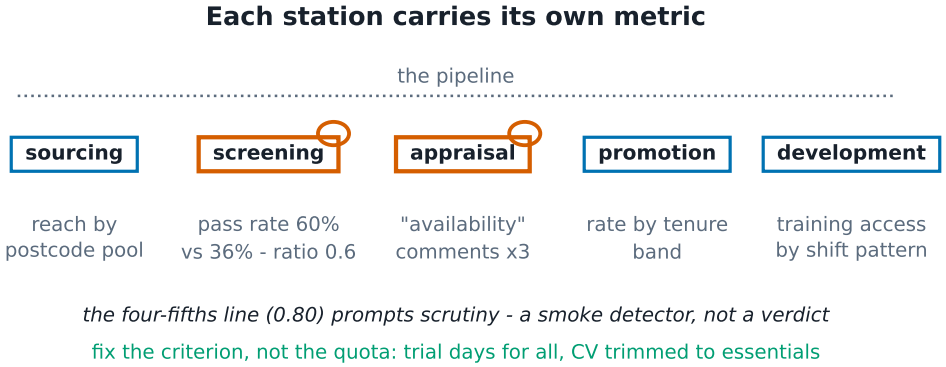
Common mistakes. One rater, one memory, one band. Rating traits instead of behaviours. Leading change with communication alone while the force-field is net negative – inspiration does not alter forces.

Summary. Anchored criteria, multiple raters, the year not the fortnight. Force-field before the announcement: remove restrainers, then push. Resistance is the plan's missing information.

Inclusion and ethics

Why it matters. People processes decide who gets hired, praised, paid and promoted – and every one of them can carry bias invisibly. The audit is what makes inclusion a practice instead of a poster.

The idea, with a picture. The process audit map: each people process (sourcing, screening, appraisal, promotion) checked against its own data for group differences that the criteria cannot explain.



Each station carries its own metric - the audit flags what the criteria cannot explain

The definitions. Adverse impact : a selection rate for a group less than four-fifths of the top group's rate – the standard's trip line, a prompt for scrutiny rather than a verdict by itself. **Bias audit** : comparing process outcomes across groups BEFORE adjusting for legitimate criteria. Ethical frames: utilitarian (greatest good), rights, justice – named, because most HR dilemmas are frames colliding.

The derivation. The screening audit: of applicants from two postcodes, one group passes the CV knock-out at 60%, the other at 36% – a ratio of 0.60, below the four-fifths line ($0.80 \times 0.60 = 0.48$; $36\% < 48\%$). Investigation: the knock-out listed "fine-dining experience", which the one postcode's casual-dining workforce could never show – a criterion sounding neutral, functioning as a gate. Fix: trial-day work samples for all, CV reduced to legal essentials – and the pass-rate gap closes to 52% vs 46%. The audit's arithmetic found what the intent never wrote down.

Worked example 1 – the appraisal distribution. Ratings cluster: the two crew members with caring responsibilities receive "availability concerns" comments at three times the rate of others – with identical measured covers and defect rates. The comment field is carrying the bias the numbers do not hold. Fix: comments must cite a measured criterion or be deleted – making the appraisal's prose as auditable as its scores.

Worked example 2 – escalate to the ethical collision. A key Saturday chef has a second job that sometimes collides. Utilitarian: flex him (the rush needs him); rights: the contract is the contract; justice: flexibility offered to one must be offerable to all. The designed answer: a swap-board with transparent rules – the flexibility is real, the justice condition is met by access, not by exception. HR ethics is mostly the design of systems where the frames stop colliding.

Common mistakes. Declaring processes unbiased because intent was clean – bias lives in outcomes, not intentions. Auditing once and framing it. Treating the four-fifths rule as either proof or permission – it is a smoke detector.

Summary. Audit outcomes by group; fix the criterion, not the quota. The four-fifths line prompts scrutiny. Ethical frames named = dilemmas designed, not ducked.

The crew sheet (reproducible)

The eight-member front-of-house crew, month 12:

- Attribution clip: 40 staff, 25 person (62%) / 15 situation (38%); the chef's evidence ledger: consistency high, distinctiveness low, consensus low – with the printer-jam fact withheld.
- Motivation: $E \to P$ 0.7, $P \to O$ 0.5, valence 0.8 → force 0.28; goal crews 14/20 versus 8/20.
- Team: 8 members, 18/28 links (density 0.64), one bridge; voices 3/12 meetings → 7/8 after the round-robin; stock-outs 3 → 1 a month.
- Leadership: grid 9,3 → target 9,7; bin rate 4.0% → 2.1%.
- Selection: 40 → 18 → 9 → 4 → 3; base rate 0.20; screen pass rates 60% vs 36% (ratio 0.60) → 52% vs 46% after the fix.
- Change: drivers +7, restrainers −9, net −2 → +5 after the two designs; training ROI £960 a year on £180.

Every figure in this handout recomputes from these values.