



WHITE PAPER · CULTURE, CAPABILITY & FRAMEWORKS

THE PUNITIVE RULEBOOK

What Australian conduct documents cost us — and what
consequence perception reveals about the waste
Evidence from 215 documents

EXECUTIVE SUMMARY

Most Australian organisations govern behaviour with a document that tells people what not to do. The code of conduct, the employee handbook and the disciplinary policy together form the written constitution of working life, and in the majority of cases that constitution is written in the negative. This paper reports a content analysis of 215 publicly available Australian codes of conduct, employee handbooks and conduct-adjacent policies totalling approximately 2.2 million words, and sets the findings against the peer-reviewed evidence on psychological safety, perceived organisational support, role clarity and internal mobility.

The pattern is consistent. Across the corpus, a median 28.1 per cent of modal directives are phrased as prohibitions rather than positive standards. Punishment and consequence vocabulary outweighs recognition and reward vocabulary by a median ratio of 5.2 to 1. Inspirational framing is nearly universal, appearing in 98 per cent of documents, but the structures that would allow a person to act on that inspiration are not. Career pathing language appears in 43 per cent of documents and internal mobility language in only 5 per cent. We describe this pattern as **aspiration without architecture**: the document promises a culture it does not build.

The pattern is expensive. Australia carries a very large measurable burden from work-related injury and illness — Safe Work Australia and Deloitte Access Economics estimate the economy would be \$28.6 billion larger each year if it were eliminated — and organisations spend heavily on surveys, safety programmes and consulting to move it. What no published source can tell us is how much. There is no national figure for engagement or culture survey spend, no national figure for safety culture programme spend, and the Australian Bureau of Statistics has not run its employer training expenditure survey since 2001–02. We treat that absence as a finding in its own right.

It is also difficult to defend theoretically. Seventy years of motivation research — Maslow, Herzberg and McClelland, then Skinner's own analysis of aversive control, then the positive psychology programme of Seligman, Csikszentmihalyi and Peterson — converges on a single conclusion. Punishment is effective at suppressing behaviour and structurally incapable of producing the discretionary effort organisations say they want. Skinner, who understood aversive control better than anyone, argued against using it precisely because it works.

Between the document an organisation writes and the behaviour it gets sits a perceptual layer: **consequence perception**, the private appraisal of how likely an outcome is, how severe, how soon and whose fault. A prohibition-dominant document is a direct input into that appraisal. Where the appraisal turns pessimistic, the first casualty is not compliance — it is disclosure. A systematic review of United States studies found between 20 and 91 per cent of workers did not report an injury or illness to management or workers' compensation. In Australian aviation, 54 per cent of pilots surveyed reported having withheld safety information. When disclosure falls, every instrument the organisation has bought to measure its culture is reading a filtered signal.

28.1%

median share of modal directives phrased as prohibitions

5.2:1

punishment to recognition vocabulary, corpus median

98% / 5%

documents with inspirational framing versus internal mobility language

215

Australian conduct documents analysed, ≈2.2m words

WHAT THIS PAPER ARGUES

A conduct system built on prohibition can only ever establish a compliance floor. It can stop the worst behaviour. It cannot raise the ceiling, because it never describes what the ceiling looks like.

That limitation is not free. Organisations invest in measuring and improving safety culture, engagement and performance while leaving the governing document — the artefact that sets the perceived consequence environment — untouched. The investment is then asked to move a number that the document is quietly holding down.

The alternative is not a softer document. It is a differently engineered one: affirmative standards that are observable, a recognition architecture that sits alongside the consequence architecture, and a published growth pathway that makes advancement part of the conduct system rather than an unrelated HR process. We set out that alternative as the Culture, Capability and Frameworks System (CCFS).

Author. Patrick Sheehan, NSC Group. This paper is published to open a debate rather than to close one. Where the evidence is thin, we say so, and we identify five specific research gaps that the field has not filled.

1. THE COMPLIANCE FLOOR

Every organisation of any size writes down how it expects people to behave. The document has different names in different places, but its function is constant: it is the artefact an organisation points to when it needs to say that a standard existed and was known. That evidentiary purpose shapes the writing. Legal review sharpens the prohibitions, because prohibitions are what a tribunal can test. It rarely sharpens the affirmative content, because nobody is ever asked to prove that they described excellence clearly enough.

The result is a document that is optimised for defensibility and largely silent on development. Read as an employee rather than as a lawyer, a typical Australian code of conduct is a list of ways to be in trouble, prefaced by two paragraphs about values. The values are sincere. They are also structurally disconnected from everything that follows.

This matters because the conduct document is not a neutral container. Bowen and Ostroff (2004) argued that human resource systems influence behaviour not only through their content but through the strength of the signal they send, which depends on distinctiveness, consistency and consensus. A document that says the organisation values initiative, and then spends forty pages describing the consequences of error, is not sending a consistent signal. It is sending two, and employees are quite capable of working out which one carries the sanctions.

We wanted to know how widespread that inconsistency is in Australia, and whether it varies by sector. So we counted.

2. METHOD

We assembled 215 publicly available Australian documents that govern workplace conduct: codes of conduct, employee handbooks, codes of ethics, and conduct-adjacent policies covering discipline, grievance, speaking up and performance. The corpus totals approximately 2.2 million words and spans ten sector groupings: listed corporates (n = 94), the public sector (n = 29), education and research (n = 25), health, aged care and disability (n = 16), banking, insurance and superannuation (n = 10), mining, resources and energy (n = 9), local government (n = 8), transport and aviation (n = 8), not-for-profit (n = 8) and police and emergency services (n = 6). Four measures were computed per document.

- **Prohibitive directive share.** Every modal directive (must, shall, will, is required to, is expected to) was classified as prohibitive or affirmative, and the prohibitive share reported.
- **Punitive density.** Terms drawn from a sanction lexicon (disciplinary action, termination, breach, misconduct, penalty, dismissal, investigation, sanction) normalised per 1,000 words.
- **Recognition density.** The same normalisation applied to a recognition lexicon (recognise, reward, celebrate, acknowledge, commend, thank).
- **Growth architecture signals.** Binary presence of four constructs: inspirational framing, role and performance clarity, career pathing and development, and internal mobility or promotion.

Two coders independently classified a 20 per cent sample to check the lexicons; disagreements were resolved by rewriting the ambiguous patterns rather than by adjudication. All associations reported are Spearman rank correlations, appropriate to the skewed distributions involved.

STATED LIMITATIONS

This is a descriptive content analysis of documents, not a study of behaviour. It tells us what organisations write, not what they do, and not what employees experience. Sector cells are small — six of the ten groupings have fewer than eleven documents — so sector medians should be read as indicative rather than as population estimates. Publicly available documents may differ systematically from internal ones.

The theoretical material in Section 5 and the economic and consequence-perception material in Sections 6 and 7 are drawn from external published sources and are not linked to our corpus by any analysis. Where a commonly quoted figure could not be confirmed against its primary source, we say so rather than repeat it. No causal claim is made or implied anywhere in this paper.

3. WHAT THE CORPUS SHOWS

Prohibition is the default voice. The corpus median for negatively framed directives is 28.1 per cent, but the distribution is wide. Listed corporates sit at 34.6 per cent, mining, resources and energy at 34.2 per cent and local government at 30.8 per cent, while the public sector sits at 17.6 per cent and not-for-profits at 18.9 per cent. The top decile of documents exceeds 45 per cent: nearly one directive in two is a statement of what is forbidden.




Figure 1. Prohibitive share of modal directives by sector grouping (n = 215). Coral bars sit at or above the corpus median. Sector cells are small and should be read as indicative.




Figure 2. Punishment vocabulary relative to recognition vocabulary, by sector grouping. Transport and aviation is the extreme case at 15.6 to 1, followed by mining, resources and energy at 10.0 to 1 and local government at 7.9 to 1 — sectors where a sanction register is arguably justified, and where the reporting evidence in Section 4 is also the most confronting.




Figure 3. Aspiration without architecture. Nearly every document inspires; almost none explains how a person moves.

Recognition is not the counterweight. The median document uses punishment and consequence vocabulary 5.2 times more often than recognition vocabulary. The transport and aviation sub-sample reaches 15.6 to 1, and local government 7.9 to 1. Only fourteen documents in the corpus specify a recognition mechanism with the same procedural detail applied to the disciplinary mechanism — who nominates, who decides, on what timeframe, with what record.

Aspiration without architecture. Ninety-eight per cent of documents contain inspirational framing: a values statement, a purpose paragraph, a line about being an employer of choice. Seventy-two per cent contain role and performance clarity language, though most of it is generic. Career pathing appears in 43 per cent. Internal mobility — any statement about how a person moves from one role to a better one — appears in 5 per cent. Nineteen documents out of 215.




Figure 4. Associations with punitive density (Spearman).

Control language and growth language do not co-occur. Punitive density is negatively associated with internal mobility language ($\rho = -0.187$, $p < .01$) and with role and performance clarity ($\rho = -0.174$, $p < .05$), and positively but weakly associated with inspirational framing ($\rho = +0.161$, $p < .05$). Career pathing shows no reliable association ($\rho = -0.043$, ns). The coefficients are modest and the design is cross-sectional; we report them as pattern, not as cause.

4. THE EVIDENCE BASE

Psychological safety. Frazier and colleagues (2017), meta-analysing the psychological safety literature, placed psychological safety in a substantial relationship with task performance, work engagement and information sharing. A document whose dominant register is sanction is a document that raises the perceived cost of admitting a mistake, and admitting mistakes is the observable behaviour that psychological safety predicts.

Perceived organisational support. Kurtessis and colleagues (2017), meta-analysing 558 studies, found perceived organisational support corrected correlations of .69 with affective commitment, .28 with organisational citizenship behaviour and -.50 with turnover intention. Employees infer support partly from formal signals. A governing document that catalogues consequence and omits recognition is a formal signal.

Internal mobility. Bidwell (2011), analysing personnel records in a United States investment bank between 2003 and 2009, found that internally promoted workers outperformed external hires for their first two years, exited at lower rates both voluntarily and involuntarily, and were paid roughly 18 per cent less than external hires at entry. This is a single-firm study and should not be over-generalised. It does, however, make the 5 per cent internal mobility figure in our corpus look expensive.

Role clarity. The relationship between role ambiguity and performance has been meta-analysed twice, by Jackson and Schuler (1985) and by Tubre and Collins (2000), both finding negative associations between ambiguity and performance. We flag a verification limit: we were unable to obtain the primary tables of either meta-analysis, and we therefore report direction rather than magnitude. Secondary sources cite a corrected correlation in the region of -.16 to -.24 for the Tubre and Collins analysis, which we have not been able to confirm against the source and do not rely on.

Safety behaviour and safety climate. Christian and colleagues (2009) meta-analysed 282 studies and found psychological safety climate corrected at .49 with safety performance behaviour and group-level safety climate at .51, with safety knowledge (.61) and safety motivation (.57) the strongest proximal predictors. The same analysis found much weaker relationships with actual accidents and injuries — psychological safety climate at -.14 at the individual level. Climate predicts what people do far better than it predicts the rare events that get counted. That gap is itself informative: it is where under-reporting lives.

Punitive climate and reporting. Feeser and colleagues (2021) documented emergency clinicians altering or suppressing safety event reports where the reporting system was perceived as punitive. Probst and Estrada (2010) found that psychological safety climate and supervisor enforcement moderate accident under-reporting, and Probst, Barbaranelli and Petitta (2013) replicated a link between job insecurity and under-reporting across two countries. In Australian aviation, McMurtrie and Molesworth found that 57 per cent of Australian pilots surveyed trusted their airline's just culture, against 84 per cent of European pilots, and that 54 per cent of Australian pilots reported having withheld safety information, against 30 per cent of Europeans, with fear of reprisal the leading reason. Their follow-up work found limited evidence of actual reprisal, which sharpens rather than weakens the point: the suppression is driven by the perceived posture of the system, and the posture is set by what the system writes down.

A NOTE ON EFFECT SIZES WE DID NOT USE

Several claims that would have suited this argument could not be verified to primary source and are therefore excluded: a pooled effect size for punitive climate and reporting suppression (none exists); a numeric discount rate for delayed workplace consequences (none exists); and a workplace-specific effect size for diffusion of responsibility (none exists). The exact moderation coefficients in Probst and Estrada (2010) and Probst et al. (2008) were not obtainable, so we report direction only.

5. WHAT THE MOTIVATION THEORISTS SAID

None of this is new information. The case against governing behaviour by prohibition was made, in different vocabularies, by every major motivation theorist of the twentieth century — including, most inconveniently for the punitive position, by the behaviourist who mapped aversive control most precisely. The conduct document is the one artefact of organisational life that appears to have been drafted without reference to any of it.




Figure 5. Seventy years of convergent argument. Three traditions, different methods and incompatible epistemologies, arriving at the same practical conclusion: sanction sets a floor and cannot build a ceiling.

The classic theorists: fear buys compliance, not contribution

Maslow (1943, 1954) placed safety needs immediately above the physiological — not because safety is noble, but because an unsatisfied safety need is prepotent: it dominates attention until it is met. Esteem and self-actualisation, the levels at which discretionary effort actually lives, only operate once safety is assumed rather than negotiated. A document that makes the terms of one's continued employment the most vivid thing an employee reads keeps attention parked at the deficiency levels. We note that the strict hierarchical ordering has never held up well empirically (Wahba and Bridwell, 1976). The deficiency-versus-growth distinction has worn considerably better, and it is the part that matters here.

Herzberg (1968) made the same point with a sharper instrument. His two-factor model separated hygiene factors — policy, supervision, working conditions, job security — from motivators: achievement, recognition, the work itself, responsibility, advancement. Hygiene factors can produce dissatisfaction when absent, but supplying them does not produce motivation. Herzberg's term for management by threat was KITA, and his verdict on it was that it produces *movement* rather than motivation: the person moves because you moved them, and stops when you stop. The list of true motivators is worth reading against our corpus. Recognition, responsibility and advancement are three of the six, and they are exactly the three that Australian conduct documents most reliably omit. Herzberg's critical-incident method has been challenged for decades (Vroom, 1964), and we cite the framework as an argument rather than as a measurement.

McClelland (1961, 1985) shifted the question from needs everyone shares to motives that vary between people. Achievement-motivated employees are not driven by consequence; they are driven by moderately difficult goals and by concrete, prompt feedback on how they are going. A prohibition-dominant document supplies neither. What it does supply, in abundance, is material for the fear-of-failure disposition that McClelland's own work distinguished from achievement striving — the motive that produces cautious target-setting and defensive behaviour. His distinction between personalised and socialised power is equally pointed for anyone drafting a disciplinary policy: institutional power exercised for the organisation's benefit builds commitment, and the same power exercised to demonstrate control does not.

Skinner: aversive control works, and that is the problem

It is common to hear the behaviourist tradition invoked in defence of the punitive document, on the reasonable-sounding basis that consequences shape behaviour. Skinner's own position was close to the opposite, and it is worth stating precisely, because the terminology is almost always used incorrectly.

A TERMINOLOGICAL CORRECTION WORTH MAKING

Negative reinforcement is not punishment. It is the removal of something aversive when the desired behaviour occurs — the pressure lifts once you comply. It *increases* behaviour. **Punishment** adds an aversive consequence, or removes something valued, in order to *decrease* behaviour.

Both are forms of what Skinner called aversive control, and both are widely used in conduct documents. Negative reinforcement is the more insidious of the two: the behaviour it reliably strengthens is whatever makes the pressure stop, which in a reporting culture is silence rather than safety.

Skinner (1953) was explicit that aversive control is effective. That was never in dispute. His objection was to its by-products, which he catalogued at length: escape and avoidance of the controlling agent, emotional responses including anxiety and aggression, and counter-control — the ingenuity people apply to evading the contingency rather than meeting the standard. Punished behaviour, he argued, is suppressed rather than extinguished; it returns when surveillance lifts. In *Beyond Freedom and Dignity* (1971) he pressed the design implication: the objection to aversive control is not sentimental but technical, and the alternative is to arrange positive reinforcement well enough that aversive control becomes unnecessary.

Read against our corpus, the technical failure is obvious. Australian conduct documents apply Skinner's own design principles — contingency, specificity, immediacy, a named schedule — with real rigour to the sanction side, and almost never to the reinforcement side. A typical document can tell you exactly who investigates a breach, on what timeframe, with what right of reply and with what consequence. The same document cannot tell you who notices when the work is done well. That is not an argument against behaviourism. It is behaviourism applied in one direction only.

One qualification belongs here, because it cuts against the simple reading. Deci, Koestner and Ryan (1999), meta-analysing 128 experiments, found that tangible rewards made contingent on performance can undermine intrinsic motivation, while verbal recognition and positive feedback generally enhance it. The design implication is not "add a bonus scheme". It is that recognition should be informational — telling a person their work was seen and why it mattered — rather than controlling. Podsakoff and colleagues (2006), meta-analysing leader reward and punishment behaviour, found the contingency is what matters: contingent recognition relates positively to performance and attitudes, while non-contingent punishment — sanction that arrives unpredictably — relates negatively to satisfaction, commitment and perceived fairness. A document that lists thirty prohibitions without base rates is, from the employee's side of the desk, indistinguishable from non-contingent punishment.

Positive psychology: the deliberate counterweight

The positive psychology programme began as an explicit correction to a discipline that had spent fifty years cataloguing pathology. Seligman and Csikszentmihalyi (2000) argued that psychology had become a science of repair and had almost nothing to say about what makes life worth living. Organisational conduct documents are in precisely that position: forty pages on damage and two paragraphs on purpose.

The most directly relevant finding is the oldest. Seligman and Maier (1967) demonstrated learned helplessness: exposure to aversive events perceived as uncontrollable produces passivity that persists even after control is restored. The organisational analogue is not exotic. An employee who concludes that outcomes are determined by factors they cannot influence — who gets sanctioned, whose report is believed, who advances — stops attempting to influence them. On an engagement survey, that appears as disengagement. It is more accurately described as a rational response to a perceived contingency structure.

Csikszentmihalyi (1990) supplies the positive specification. Flow — the absorbed, high-performing state organisations claim to want — has well-described antecedents: clear proximate goals, immediate and unambiguous feedback, and a challenge level matched to skill. Set that list beside a prohibition inventory and the mismatch is total. The goals are negative and distal, the feedback arrives only on failure, and challenge is not addressed at all. An organisation cannot instruct people into flow, but it can build the three conditions into the way work and standards are described. Almost none do.

Peterson and Seligman (2004) closed the loop by building the artefact the argument implied. The *Character Strengths and Virtues* classification was written deliberately as the counterpart to the diagnostic manual of mental disorders: a structured, measurable description of what is right with people, offered because the field had a detailed taxonomy of what is wrong and none of what is strong. The parallel to the conduct document is exact, and it is the reason the classic theorists appear in this paper at all. Most organisations have a rigorous, well-drafted, legally reviewed taxonomy of employee failure. Very few have anything with equivalent structure describing employee excellence. The corpus quantifies the imbalance; Peterson and Seligman show what fixing it looks like as a piece of design work.

WHAT THE THREE TRADITIONS AGREE ON

They disagree about almost everything else. Maslow's humanism, Skinner's radical behaviourism and the empirical positive psychology programme are incompatible accounts of why people do things. They converge on the practical point anyway.

- 1.** Aversive control reliably suppresses specified behaviour. Nobody disputes its short-run effectiveness.
- 2.** It does not generate the behaviour nobody specified — initiative, disclosure, discretionary effort — because suppression has no positive target.
- 3.** It produces side effects that are costly and predictable: avoidance, concealment, defensive caution and passivity.
- 4.** The conditions that do generate contribution are specifiable — clear goals, prompt informational feedback, calibrated challenge, contingent recognition, a visible route forward — and can be written down as deliberately as prohibitions are.

That fourth point is the whole design argument. The theory has been settled for decades. It has simply never been applied to the document.

The uncomfortable conclusion is that the punitive conduct document is not a theoretical error. It is a theoretical absence. No serious motivational account recommends it. It persists because it is drafted for a different purpose entirely — evidentiary defensibility — by people who are not asked to consider its motivational consequences, and because nobody in the organisation owns the question of what the document is doing to behaviour.

6. THE MEASUREMENT PARADOX

Australian organisations spend a great deal of money finding out how their people feel and how safely they work. They run engagement surveys, culture surveys, pulse checks, safety climate assessments, leadership programmes and behavioural safety interventions. They then take corrective action against the results. Very little of that expenditure is directed at the document that sets the perceived consequence environment in the first place.




Figure 6. What can and cannot be established about Australian expenditure. Sources: Safe Work Australia / Deloitte Access Economics (2022) for the \$28.6bn annual economic value and \$37.6bn cumulative health-system cost; IBISWorld (2026) for management consulting industry revenue. The three grey columns are not zero — they are unmeasured. No national published figure exists for engagement or culture survey spend, for safety culture programme spend, or for total HR and learning and development expenditure. The Australian Bureau of Statistics discontinued its Employer Training Expenditure and Practices survey after 2001-02.

The honest answer to “how much is Australia spending?” is that nobody publishes it. That is worth stating plainly, because the question is usually answered with a confident number lifted from a vendor market-sizing report. We looked for an authoritative figure and did not find one. What is verifiable is the size of the problem the spending is aimed at, and the size of the industry that sells the remedy.

The burden is well measured. Safe Work Australia’s 2022 study with Deloitte Access Economics modelled the economy-wide effect of eliminating work-related injury and illness: gross domestic product \$28.6 billion higher each year, 185,500 additional full-time equivalent jobs and average wages 1.3 per cent higher. The same work put cumulative health-system costs from 6.9 million work-related injuries and illnesses between 2008 and 2018 at \$37.6 billion.

The remedy market is large. IBISWorld estimates Australian management consulting industry revenue at \$40.8 billion for 2025–26. That figure covers far more than culture and safety work and cannot be read as a culture-spend proxy, but it establishes the order of magnitude of the advisory economy that organisational problems support.

The engagement position has not moved. Gallup’s country-level data puts Australian employee engagement at 21 per cent on a three-year rolling average to 2025, against 23 per cent the year before and 21 per cent the year before that. Globally, Gallup estimates the cost of low engagement at US\$8.8 trillion, or about 9 per cent of global gross domestic product. Gallup does not disaggregate that figure by country, and we decline to do it for them: an Australian dollar figure derived by applying the global percentage to Australian GDP would be arithmetic dressed as evidence.

Put those observations together and a specific paradox appears. The burden is measured precisely. The remedy industry is large. The outcome measure has been flat for three years. And the one artefact that every employee actually reads — the conduct document — is almost never part of the intervention.




Figure 7. How the investment is absorbed. Each step is individually reasonable; the sequence is what wastes the money.

We are not arguing that surveys and safety programmes do not work. We are arguing something narrower and more testable: that an organisation which improves its measurement while leaving its governing document prohibition-dominant has left the input variable untouched, and should not be surprised when the output variable does not move. The corpus shows how common that configuration is. It cannot show what it costs, and we do not claim to know.

There is a version of this point that the motivation literature already made. Herzberg’s distinction between movement and motivation is, in budget terms, the difference between an intervention that runs while it is funded and one that changes the conditions under which people work. An organisation that measures annually, exhorts quarterly and leaves the contingency structure exactly as written has bought movement.

WHAT WE COULD NOT ESTABLISH

There is no published Australian national figure for engagement or culture survey spend, for safety culture programme spend, or for HR and learning and development expenditure. There is no Gallup-published Australian dollar cost of disengagement. We have not substituted an estimate for any of these. Readers who have seen such a figure quoted should check what it is sourced to.

7. CONSEQUENCE PERCEPTION: THE MISSING LAYER

The mechanism connecting a written document to an unmeasured behaviour is perceptual. People do not respond to the consequence structure that exists; they respond to the consequence structure they believe exists. Consequence perception is the appraisal that produces that belief, and it has four separable facets — likelihood, severity, proximity and attribution (Influential Research, 2026).



Figure 8. Consequence perception as the layer between the governing document and observable behaviour. A prohibition-dominant, recognition-silent document is an input to all four facets. Facet definitions follow the Consequence Perception thought paper (Influential Research, 2026), which draws on Slovic (1987) on risk perception and Trope and Liberman (2010) on construal-level theory.

Likelihood. How probable the person believes the downstream outcome to be. Slovic's (1987) psychometric work established that lay risk appraisal is driven by dread and unfamiliarity rather than by statistical frequency. A document that lists thirty prohibitions and no base rates gives an employee nothing to calibrate against, so the estimate is formed from rumour and from the last case anybody remembers.

Severity. How serious the outcome would be if it landed. Here the punitive document is unusually precise, and precise in one direction only: it describes what happens when a person fails in detail, and what happens when a person excels not at all. The asymmetry is the point, and it is the asymmetry Skinner would have identified immediately: severity is well specified on the loss side and undefined on the gain side.

Proximity. How close the consequence feels in time and in psychological distance. Trope and Liberman's (2010) construal-level theory holds that distal outcomes are represented more abstractly and weighted less concretely than proximal ones. An immediate disciplinary risk is concrete; a cumulative safety harm three years away is abstract. This is the same immediacy principle the reinforcement literature insists on, operating against the organisation rather than for it. We note that no peer-reviewed study quantifies a discount rate for occupational safety consequences specifically. The mechanism is well established; the workplace magnitude is a genuine gap.

Attribution. Who the person believes will be held responsible. This is the facet the corpus speaks to most directly. A document that names an investigator, a decision-maker and a timeframe for misconduct, but names nobody for improvement, allocates accountability in one direction. Gyekye (2010) found workers' causal attributions for accidents shape both compliance and perceived organisational responsibility. We searched for a meta-analytic effect size for diffusion of responsibility in occupational settings and did not find one.

The observable consequence: disclosure falls

Consequence perception is not an abstraction. Where the appraisal turns pessimistic, the behaviour that changes first is disclosure, and disclosure is what every measurement system in the organisation depends on. It is also, in Skinner's terms, exactly what a negatively reinforced environment should be expected to produce: the response that most reliably removes the aversive prospect is not to report.




Figure 9. Under-reporting where it has been measured. Sources: Kim et al. (2023) systematic review; Scherzer et al. (2005); McMurtrie and Molesworth (2018).

A 2023 systematic review of twenty United States studies found that between 20 and 91 per cent of workers did not report injuries or illnesses to management or to workers' compensation, with distrust of reporting consequences among the contributing factors. In one study of Las Vegas hotel room cleaners, only 33 per cent of those with pain or discomfort in the previous twelve months had reported it to the employer, and 26 per cent had filed a claim. The Australian aviation figures cited earlier sit in the same territory.

WHY THIS COMPOUNDS THE MEASUREMENT PARADOX

Under-reporting does not simply hide incidents. It corrupts the instruments the organisation is paying for. An engagement survey run inside a low-attribution, high-severity perception environment collects the answers people believe are safe to give. A safety climate score in the same environment reads the reporting behaviour, not the hazard. Christian and colleagues (2009) found climate correlated with safety behaviour at around .49 to .51 but with actual accidents at -.14 to -.39 — precisely the pattern you would expect if the outcome data were being filtered on its way to the dashboard.

This is the strongest practical argument in the paper, and it is the one we hold most carefully. The mechanism is well evidenced in single studies and theoretically coherent. It has not been demonstrated end-to-end — document grammar to consequence perception to under-reporting to measurement error — in a single study, in Australia or anywhere else. That study is worth doing.

8. TWO CONDITIONS

The useful way to read all of this is as a comparison between two coherent designs, rather than as a grading of individual documents. Condition A is the punitive, closed system that the corpus median describes. Condition B is the affirmative, open system that the Culture, Capability and Frameworks System is built to produce. Both are internally consistent. They differ in what they are engineered to achieve.




Figure 10. The two conditions. Condition A establishes a compliance floor; Condition B raises the performance ceiling by adding the architecture that makes the stated standard actionable.

Design dimension	Condition A — punitive / closed	Condition B — affirmative / CCFS
Directive grammar	Obligations written as prohibitions. In the corpus, a median 28.1% of modal directives are negatively framed; the top decile exceeds 45%.	Obligations written as observable positive standards. Prohibition is reserved for genuine bright lines, not used as the default voice.
Consequence and recognition	Punishment vocabulary outweighs recognition vocabulary 5.2:1 at the corpus median, and 15.6:1 in the transport and aviation sub-sample.	Consequence architecture and recognition architecture are specified in the same document, with named mechanisms for both.
Motivational logic	Aversive control applied with rigour in one direction. Movement, in Herzberg's sense: behaviour persists while the contingency is visible and stops when it is not.	Contingent, informational recognition with the same specificity as sanction. Motivation in the sense the literature actually supports: achievement, feedback, advancement.
Standard of performance	Expectations are asserted but rarely operationalised. Role and performance clarity language is present in 72% of documents but is largely generic.	Behavioural standards are defined per role and per level, observable, and tied to the capability framework used in development conversations.
Growth architecture	Career pathing language appears in 43% of documents; internal mobility language in 5%. Advancement is invisible in the governing document.	Pathways, mobility rules and progression criteria are published and referenced from the code, making advancement a visible part of the conduct system.
Consequence perception	Severity is precisely specified on the loss side only; attribution points one way; proximity of any positive outcome is undefined. The appraisal that results discourages disclosure.	All four facets are addressed deliberately, and measured. Likelihood is calibrated with base rates, severity is specified in both directions, proximity is shortened by fast feedback, and attribution is named for improvement as well as for breach.
Psychological posture	Employees read the document as a risk register about themselves. Speaking up carries perceived downside; sustained, the posture is the organisational form of learned helplessness.	Employees read the document as a description of good practice. Reporting and challenge are framed as contributions.
Measurement	Compliance is measured by breach counts and training completion — both of which are sensitive to reporting behaviour.	Compliance is measured alongside culture signal: divergence between stated standards and lived experience, tracked over time by cohort.

Note what Condition B does not do. It does not remove consequence. Safety-critical sectors need prohibitions, and the aviation data is a reminder that the issue is not the existence of sanction but the absence of anything else. Condition B keeps the sanction and adds the four things the corpus shows are usually missing: an observable positive standard, a recognition mechanism with the same specificity as the disciplinary mechanism, a published route from here to somewhere better, and a deliberate read on how consequence is actually being perceived.

9. FIVE DESIGN MOVES

Test every prohibition against an affirmative rewrite. Can the obligation be expressed as an observable positive standard without losing the bright line? Most can. “Employees must not withhold information relevant to a safety risk” becomes “Employees report safety risks as soon as they are identified, and are supported for doing so.” The obligation is unchanged. The register is not. Prohibitions that survive the test — fraud, harassment, impairment on site — are kept as prohibitions and gain force from being fewer.

Build the recognition architecture in the same document. If the disciplinary process is specified to the level of who investigates, within what timeframe and with what right of reply, then the recognition process should be specified to the same level. Most organisations can describe termination in four hundred words and cannot describe commendation in forty. Contingency, specificity and immediacy are the reinforcement conditions the literature has been clear about since Skinner; conduct documents already demonstrate that the organisation can meet them when it wants to.

Define behaviour by role and level. Generic values statements are the reason 72 per cent of documents contain role clarity language that nobody can act on. CCFS defines what each value looks like in the behaviour of a frontline employee, a team leader and a senior leader, in observable terms, and uses those same definitions in development conversations, selection and progression decisions. This is the Peterson and Seligman move applied to an organisation: build the taxonomy of what good looks like with the same rigour already spent on the taxonomy of failure.

Publish the pathway. The single largest gap in the corpus is internal mobility. Making progression criteria visible costs an organisation very little and changes how the conduct document reads, because it converts the document from a description of how to avoid loss into a description of how to accumulate something. Advancement is on Herzberg’s motivator list and absent from 95 per cent of the corpus.

Measure the perception, not just the outcome.

The four facets of consequence perception are measurable. Reading likelihood, severity, proximity and attribution separately tells you which deficit you have, which matters because they need different interventions: weak likelihood responds to calibration and base-rate exposure; weak severity to direct customer and community contact; weak proximity to shorter feedback loops and after-action reviews; weak attribution to explicit decision rights and named accountability. A single undifferentiated culture programme treats four different problems with one instrument, which is one reason the numbers so often refuse to move.

The measurement discipline matters as much as the drafting. CCFS instruments the gap between the stated standard and the lived experience, and tracks the divergence over time by cohort. An organisation that rewrites its code and measures nothing has bought a new document. An organisation that measures the divergence has bought an instrument.

10. IMPLICATIONS

For boards and executives, the practical question is not whether the code of conduct is legally sound. It almost certainly is. The question is whether the organisation has ever audited the ratio between what the document forbids and what it enables, and whether anyone owns the second half of that ratio.

There is a budgetary version of the same question. An organisation that runs an annual engagement survey, funds a safety culture programme and retains external advisers, while its conduct document sits at the corpus median, is spending on the measurement and the remediation of an environment it is simultaneously writing. That is not an accusation of bad faith. It is an accounting observation. The document is usually owned by legal or governance; the engagement number is usually owned by human resources; the safety number is usually owned by operations. Nobody owns the join.

There is also a defence that will not survive contact with the literature. It is sometimes argued that conduct documents are simply not motivational instruments — that they set boundaries and that motivation happens elsewhere. Skinner’s analysis disposes of that: a contingency structure does not become motivationally inert because its authors did not intend it to be read that way. The document is the most explicit statement an organisation makes about what happens to people, and it is read as such.

For regulators and for boards in safety-critical sectors, the reporting evidence deserves particular weight. The Australian aviation findings — 54 per cent of pilots withholding safety information despite limited evidence of actual reprisal — establish that perceived posture, not actual practice, governs disclosure. Every

organisation that relies on voluntary reporting for its safety picture is exposed to that gap, and the governing document is one of the few levers that acts on it directly.

FIVE RESEARCH GAPS WE COULD NOT CLOSE

1. Clause-ratio evidence. No peer-reviewed study quantifies prohibitive versus affirmative clause ratios in codes of conduct at scale and links them to outcomes. The code-of-conduct effectiveness literature (Kaptein, 2011; Erwin, 2011) measures code quality and code embedding, not directive grammar.

2. No pooled effect size for punitive climate and reporting suppression. The direction is well supported by single studies and by the just-culture literature. A meta-analytic estimate does not exist.

3. No general Australian evidence outside aviation. Australian peer-reviewed research on punitive policy language and its effects is confined to aviation just-culture studies.

4. No national expenditure series. Australia does not publish what it spends on engagement, culture or safety-culture work. The ABS training expenditure survey has not run since 2001-02.

5. No workplace magnitude for temporal discounting or diffusion of responsibility. Both mechanisms are well established in general psychology and neither has been quantified in an occupational safety context. We would add a sixth if we were being strict: no study has tested the classic motivational predictions against document grammar directly.

We state these gaps deliberately. A paper that argued from the corpus straight to a claim about performance, or from three verified national statistics straight to a wasted-billions headline, would be doing exactly what we are criticising: asserting a standard without building the architecture to support it. The clause-ratio question in particular is tractable. The corpus methodology described here could be paired with engagement, turnover or incident-reporting data in a single large organisation to test whether directive grammar predicts anything at all. We would be glad to see that study done, including if it disconfirms us.

THE ARGUMENT IN ONE LINE

A rulebook can prevent the worst version of an organisation. Only an architecture can produce the best one — and 95 per cent of Australian conduct documents do not contain a single sentence about how a person moves forward.

Patrick Sheehan, NSC Group. nscgroup.au

ABOUT THIS WORK

NSC Group works with boards and executive teams on organisational high performance: the design of culture, capability and conduct systems that hold together under pressure. The Culture, Capability and Frameworks System is our applied model for that work, and this paper describes the diagnostic evidence that sits behind it.

We publish research of this kind to be argued with. If your organisation holds internal conduct documents, engagement data or incident-reporting data that could test the associations reported here, we are interested in hearing from you, including on a de-identified basis.

WHERE TO START

Run the ratio on your own document. Count the modal directives, split them into positive and negative, and count the sentences that describe how a person progresses. Then ask four questions of your workforce: how likely, how severe, how soon, and whose fault. The first exercise takes an afternoon. The second takes a week. Together they usually settle the question of whether you have a rulebook or an architecture.

REFERENCES

- Australian Bureau of Statistics. (2002). *Employer Training Expenditure and Practices, Australia* (cat. no. 6362.0). ABS. [Series discontinued after 2001-02.]
- Bidwell, M. (2011). Paying more to get less: The effects of external hiring versus internal mobility. *Administrative Science Quarterly*, 56(3), 369–407. doi:10.1177/0001839211433562
- Bowen, D. E., & Ostroff, C. (2004). Understanding HRM–firm performance linkages: The role of the “strength” of the HRM system. *Academy of Management Review*, 29(2), 203–221. doi:10.5465/AMR.2004.12736076
- Christian, M. S., Bradley, J. C., Wallace, J. C., & Burke, M. J. (2009). Workplace safety: A meta-analysis of the roles of person and situation factors. *Journal of Applied Psychology*, 94(5), 1103–1127. doi:10.1037/a0016172
- Csikszentmihalyi, M. (1990). *Flow: The psychology of optimal experience*. Harper & Row.
- Deci, E. L., Koestner, R., & Ryan, R. M. (1999). A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin*, 125(6), 627–668. doi:10.1037/0033-2909.125.6.627
- Erwin, P. M. (2011). Corporate codes of conduct: The effects of code content and quality on ethical performance. *Journal of Business Ethics*, 99(4), 535–548. doi:10.1007/s10551-010-0667-y
- Feeser, V. R., Jackson, A. K., Savage, N. M., Layng, T. A., Senn, R. K., Dhindsa, H. S., Nguyen, C. Q., & Hays, W. B. (2021). When safety event reporting is seen as punitive: “I’ve been PSN-ed!” *Annals of Emergency Medicine*, 77(4), 449–458. doi:10.1016/j.annemergmed.2020.06.048
- Frazier, M. L., Fainshmidt, S., Klinger, R. L., Pezeshkan, A., & Vracheva, V. (2017). Psychological safety: A meta-analytic review and extension. *Personnel Psychology*, 70(1), 113–165. doi:10.1111/peps.12183
- Gallup. (2026). *State of the Global Workplace: Australia country-level data*. Gallup.
- Gyekye, S. A. (2010). Occupational safety management: The role of causal attribution. *International Journal of Psychology*, 45(6), 405–416. doi:10.1080/00207594.2010.501337
- Herzberg, F. (1968). One more time: How do you motivate employees? *Harvard Business Review*, 46(1), 53–62.
- Herzberg, F., Mausner, B., & Snyderman, B. B. (1959). *The motivation to work* (2nd ed.). Wiley.
- IBISWorld. (2026). *Management consulting in Australia* (ANZSIC M6962a). IBISWorld industry report.
- Influential Research. (2026). *Consequence perception: A thought paper*. Influential Solutions.
- Jackson, S. E., & Schuler, R. S. (1985). A meta-analysis and conceptual critique of research on role ambiguity and role conflict in work settings. *Organizational Behavior and Human Decision Processes*, 36(1), 16–78.
- Kaptein, M. (2011). Toward effective codes: Testing the relationship with unethical behavior. *Journal of Business Ethics*, 99(2), 233–251. doi:10.1007/s10551-010-0652-5
- Kim, M., et al. (2023). Underreporting of workers’ injuries or illnesses and contributing factors: A systematic review. *BMC Public Health*, 23, 558. doi:10.1186/s12889-023-15487-0
- Kurtessis, J. N., Eisenberger, R., Ford, M. T., Buffardi, L. C., Stewart, K. A., & Adis, C. S. (2017). Perceived organizational support: A meta-analytic evaluation of organizational support theory. *Journal of Management*, 43(6), 1854–1884. doi:10.1177/0149206315575554
- Maslow, A. H. (1943). A theory of human motivation. *Psychological Review*, 50(4), 370–396. doi:10.1037/h0054346
- Maslow, A. H. (1954). *Motivation and personality*. Harper & Row.
- McClelland, D. C. (1961). *The achieving society*. Van Nostrand.
- McClelland, D. C. (1985). *Human motivation*. Scott, Foresman.
- McMurtrie, K., & Molesworth, B. R. C. (2018). Australian flight crews’ trust in voluntary reporting systems and just culture policies. *Aviation Psychology and Applied Human Factors*, 8(1), 11–21. doi:10.1027/2192-0923/a000131
- McMurtrie, K., & Molesworth, B. R. C. (2022). Confidence and trust in the “just culture” construct. *Transportation Research Procedia*, 66, 137–145. doi:10.1016/j.trpro.2022.12.022 (conference proceedings)
- Nagin, D. S. (2013). Deterrence: A review of the evidence by a criminologist for economists. *Annual Review of Economics*, 5, 83–105. doi:10.1146/annurev-economics-072412-131310
- Peterson, C., & Seligman, M. E. P. (2004). *Character strengths and virtues: A handbook and classification*. American Psychological Association / Oxford University Press.
- Podsakoff, P. M., Bommer, W. H., Podsakoff, N. P., & MacKenzie, S. B. (2006). Relationships between leader reward and punishment behavior and subordinate attitudes, perceptions, and behaviors: A meta-analytic review of existing and new research. *Organizational Behavior and Human Decision Processes*, 99(2), 113–142. doi:10.1016/j.obhdp.2005.09.002
- Probst, T. M., Barbaranelli, C., & Petitta, L. (2013). The relationship between job insecurity and accident under-reporting: A test in two countries. *Work & Stress*, 27(4), 383–402. doi:10.1080/02678373.2013.850756
- Probst, T. M., Brubaker, T. L., & Barsotti, A. (2008). Organizational injury rate underreporting: The moderating effect of organizational safety climate. *Journal of Applied Psychology*, 93(5), 1147–1154.
- Probst, T. M., & Estrada, A. X. (2010). Accident under-reporting among employees: Testing the moderating influence of psychological safety climate and supervisor enforcement of safety practices. *Accident Analysis & Prevention*, 42(5), 1438–1444. doi:10.1016/j.aap.2009.06.027
- Safe Work Australia. (2022). *Safer, healthier, wealthier: The economic value of reducing work-related injuries and illnesses* (Deloitte Access Economics). Safe Work Australia.
- Scherzer, T., Rugulies, R., & Krause, N. (2005). Work-related pain and injury and barriers to workers’ compensation among Las Vegas hotel room cleaners. *American Journal of Public Health*, 95(3), 483–488.
- Seligman, M. E. P., & Csikszentmihalyi, M. (2000). Positive psychology: An introduction. *American Psychologist*, 55(1), 5–14. doi:10.1037/0003-066X.55.1.5
- Seligman, M. E. P., & Maier, S. F. (1967). Failure to escape traumatic shock. *Journal of Experimental Psychology*, 74(1), 1–9. doi:10.1037/h0024514
- Skinner, B. F. (1953). *Science and human behavior*. Macmillan.
- Skinner, B. F. (1971). *Beyond freedom and dignity*. Knopf.
- Slovic, P. (1987). Perception of risk. *Science*, 236(4799), 280–285. doi:10.1126/science.3563507
- Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, 117(2), 440–463. doi:10.1037/a0018963
- Tubre, T. C., & Collins, J. M. (2000). Jackson and Schuler (1985) revisited: A meta-analysis of the relationships between role ambiguity, role conflict, and job performance. *Journal of Management*, 26(1), 155–169. doi:10.1177/014920630002600104
- Vroom, V. H. (1964). *Work and motivation*. Wiley.
- Wahba, M. A., & Bridwell, L. G. (1976). Maslow reconsidered: A review of research on the need hierarchy theory. *Organizational Behavior and Human Performance*, 15(2), 212–240. doi:10.1016/0030-5073(76)90038-6
- Corpus data.** NSC Group Australian conduct-document content analysis, 2026. N = 215 documents, ≈2.2 million words. Method and lexicons available on request.