
WORKING PAPER

Not Good or Bad, But How Often

The Development of a Frequency-Based Behavioural Taxonomy for Organisational Culture Dialogue

Neil McGregor

Founder, Culture Science Cards

2025

ABSTRACT

The development of behavioural taxonomies for organisational culture work has been dominated by instruments designed for psychometric measurement — instruments whose item selection, validation, and scoring procedures are calibrated for diagnostic accuracy rather than conversational productivity. This paper describes a different kind of taxonomic development: one oriented explicitly toward facilitating dialogue rather than producing scores, and grounded in the proposition that the most important question to ask about any organisational behaviour is not whether it is good or bad, but what happens when it becomes normal.

Drawing on five years of iterative practitioner development across more than fifty organisations and two thousand participants, this paper presents the theoretical and empirical logic through which the Culture Science Cards sixteen-theme, eighty-item behavioural framework was constructed. The paper documents three distinct phases of development: the conceptual phase, in which the founding theoretical logic was established; the structural refinement phase, in which the taxonomy evolved from an early virtue-language model through iterative testing toward the current behaviourally specific framework; and the field validation phase, in which items were tested in real facilitation sessions and replaced when they failed to generate productive dialogue.

Four original contributions are advanced. First, the frequency insight is proposed as a theoretically grounded alternative to absolute good/bad classification in culture taxonomy development: behaviours are classified not by their intrinsic quality but by their effect when they become habitual patterns in organisational life. Second, the distinction between virtue language and behavioural language is identified as a critical design variable in culture facilitation instruments, with the former producing abstract agreement and the latter producing specific, discussable, designable conversation. Third, a three-criteria item selection framework is proposed for conversation-oriented instruments: recognisability, practical specificity, and productive ambiguity. Fourth, the iterative field validation methodology — in which items are refined based on observed conversational productivity rather than statistical properties — is proposed as the appropriate validation approach for facilitation instruments whose purpose is dialogue generation rather than psychometric measurement.

The paper concludes that the Culture Science Cards taxonomy represents a methodologically distinct contribution to organisational culture work — not a measurement instrument with different content, but a conversation instrument built on different epistemological foundations. Its development logic, item selection criteria, and validation methodology reflect and instantiate the behaviour-first, dialogue-oriented approach to culture change argued in the companion paper *Envision, Diagnose, Design* (McGregor, 2025).

Keywords: *organisational culture taxonomy, behavioural classification, frequency-based model, facilitation instrument development, productive ambiguity, iterative content validity, culture dialogue, behaviour-first, practitioner inquiry, item selection*

1. Introduction

Every taxonomy is a theory. The decision to classify organisational behaviours in one way rather than another encodes assumptions about what matters, what causes what, and what questions are worth asking. The taxonomies that have dominated organisational culture work for the past four decades have been shaped primarily by the requirements of psychometric measurement: they classify behaviours in ways that support reliable scoring, normative comparison, and statistical

analysis. These are legitimate requirements for diagnostic instruments. They are not the only requirements a behavioural taxonomy might serve.

This paper describes the development of a taxonomy built for a different purpose: facilitating honest dialogue about organisational culture in groups. The Culture Science Cards framework — sixteen themes, eighty behavioural items — was not designed to measure culture. It was designed to make culture visible, discussable, and designable in the room where people work together. That different purpose required different development logic, different item selection criteria, and a different approach to validation.

The central theoretical contribution of the taxonomy is what this paper terms the frequency insight: the proposition that the most useful question to ask about any organisational behaviour is not whether it is good or bad in the abstract, but what happens when it becomes the habitual pattern — the normal, repeated, taken-for-granted way of doing things in an organisation. This reframing shifts the taxonomy from moral classification to pattern recognition, from judgment to observation, and from prescription to dialogue. It is a reframing with significant theoretical grounding and significant practical consequences, both of which are explored in the sections that follow.

The paper is structured as follows. Section 2 reviews the relevant literature on culture taxonomy development and identifies the theoretical gap this paper addresses. Section 3 presents the founding theoretical logic of the framework. Section 4 describes the above and below line architecture and the frequency insight. Section 5 documents the three phases of development in detail. Section 6 presents the sixteen themes and their theoretical grounding. Section 7 describes the item selection criteria and field validation methodology. Section 8 examines the physical design evolution of the cards. Section 9 discusses implications and future research directions. Section 10 concludes.

A note on positionality. This paper adopts a practitioner-inquiry stance (Schön, 1983; Reason and Bradbury, 2008). The author developed the taxonomy described, facilitated its testing across more than fifty organisations, and generated the observations that drove its iterative refinement. The limitations of this positionality — including potential confirmation bias and the absence of independent validation — are acknowledged throughout. The paper claims the rigour of systematic reflective practice and the theoretical coherence of a framework grounded in converging research traditions, not the objectivity of detached empiricism.

2. Literature Review: Culture Taxonomy Development

2.1 The Dominant Tradition: Psychometric Classification

The dominant tradition in organisational culture taxonomy development is psychometric. Instruments in this tradition classify behavioural tendencies into categories, ask respondents to rate their frequency or desirability, aggregate responses across populations, and produce normative profiles that can be compared against benchmarks. The Competing Values Framework (Quinn and Rohrbaugh, 1983) and its associated Organisational Culture Assessment Instrument (Cameron and Quinn, 2011) exemplify this tradition, as do a range of other widely used instruments whose development followed similar paths.

The psychometric tradition has produced genuine knowledge about organisational culture. It has demonstrated that cultural orientations can be reliably measured, that they vary systematically across organisations and industries, and that certain cultural profiles are associated with better or worse organisational outcomes (Kotter and Heskett, 1992; Denison, 1990). These are significant contributions.

The tradition has also produced a characteristic set of limitations. Psychometric instruments are designed for inter-rater reliability — every respondent should interpret every item in the same way. This requirement pushes item design toward specificity and closure: items must mean the same thing to everyone. The result is instruments whose items are optimised for consistent measurement but not for productive conversation. An item designed to be interpreted identically by all respondents is, by design, an item with no productive ambiguity — and productive ambiguity, this paper argues, is the primary engine of genuine culture dialogue.

A second limitation of the psychometric tradition is its classification logic. Most instruments classify behaviours as positive or negative, constructive or defensive, effective or ineffective. This binary classification serves the diagnostic purpose — it allows organisations to be located relative to a normative standard — but it does not serve the dialogue purpose. When a behaviour is classified as inherently negative, participants who recognise it in their culture feel judged rather than understood. When a behaviour is classified as inherently positive, participants agree and move on. Neither response produces the quality of genuine, honest, contested dialogue that culture change requires.

2.2 Alternative Traditions

Several alternative traditions in culture and organisational behaviour research offer partial correctives to the psychometric dominance. Appreciative Inquiry (Cooperrider and Srivastva, 1987) challenged the deficit-focused logic of diagnostic instruments, arguing that organisations move in the direction of what they study and that culture change is better served by studying what works at its best than by cataloguing what is dysfunctional. This insight influenced the envisioning-first design of the Culture Science Cards methodology, though Appreciative Inquiry does not itself provide a behavioural taxonomy.

Dialogic Organisation Development (Bushe and Marshak, 2015) offered a more fundamental challenge to the diagnostic paradigm, arguing that the generative conversations that produce culture change cannot be scripted or predicted in advance — that they emerge from the conditions of dialogue rather than from the content of diagnostic findings. The Culture Science Cards taxonomy is, in these terms, a resource for dialogic organisation development: it provides the shared language and structured prompts that make generative dialogue possible without prescribing its direction or content.

Weick's (1995) sensemaking framework provides theoretical grounding for the taxonomy's dialogue orientation. If culture is an ongoing process of collective meaning-making rather than a fixed property of organisations, then a culture instrument should support and enrich that process rather than interrupt it with external judgments. A taxonomy designed for sensemaking will look different from a taxonomy designed for measurement — and the Culture Science Cards framework is an attempt to build the former.

The positive organisational behaviour tradition (Luthans, 2002; Cameron, Dutton and Quinn, 2003) provides additional theoretical resources, demonstrating that the study of positive organisational states and behaviours produces different and in many respects more useful knowledge than the study of dysfunction. This tradition supports the performance-building orientation of the taxonomy's above-the-line themes and the frequency insight's reframing of below-the-line behaviours as patterns rather than pathologies.

2.3 The Theoretical Gap

The gap this paper addresses is the absence of a published account of behavioural taxonomy development for dialogue-oriented culture facilitation instruments. The psychometric tradition has produced extensive methodological literature on instrument development — item generation, expert review, pilot testing, factor analysis, reliability and validity assessment (DeVellis, 2016; Hinkin, 1998). No equivalent literature exists for facilitation instruments whose primary validation criterion is conversational productivity rather than measurement accuracy.

This paper is an attempt to begin filling that gap. It does not claim that the Culture Science Cards taxonomy has achieved formal psychometric validation — it has not, and the paper is explicit about this. It claims instead that the taxonomy has been developed through a rigorous process appropriate to its purpose, and that the development logic, item selection criteria, and field validation methodology described constitute a contribution to the methodological literature on dialogue-oriented culture instruments.

3. The Founding Theoretical Logic

3.1 The Performance-Through-People Insight

The Culture Science Cards taxonomy began with a question rather than a framework: what really drives organisational performance? The initial answer — that organisations need outcomes, delivery, execution, standards, accountability, and results — was quickly recognised as incomplete. Tasks do not complete themselves. Strategies do not execute themselves. Performance is achieved through people: their choices, habits, relationships, confidence, trust, energy, courage, and conversations.

This insight — that organisations achieve task performance through people — became the founding logic of the taxonomy. It established from the outset that culture was not going to be treated as something separate from performance: something soft, decorative, or secondary to the real work of the organisation. Culture was going to be understood as the pattern of behaviours that either helps or limits people in getting meaningful work done together. This framing connects to the sociotechnical systems tradition (Trist and Bamforth, 1951; Emery and Trist, 1960), which established that organisational performance is always jointly determined by the technical system — the tasks, tools, and processes through which work is done — and the social system — the relationships, norms, and behaviours through which people coordinate their effort.

The dual focus leadership research (Fleishman, 1953; Blake and Mouton, 1964) demonstrated empirically that the most effective leaders maintained simultaneous attention to task and people — that neither focus alone was sufficient for sustained high performance. The Culture Science Cards taxonomy encodes this insight structurally: eight performance-building themes describe the

people-oriented behaviours that make task performance possible, while the performance-limiting themes describe the patterns — control, fear, criticism, compliance, competition — that arise when task pressure drives out attention to the human conditions of performance.

Hackman and Oldham's (1976) job characteristics model provided additional theoretical grounding, establishing that intrinsic motivation — the engagement that produces sustained high performance — is a function of task variety, identity, significance, autonomy, and feedback. These five characteristics are not primarily technical properties of jobs. They are social and relational properties of the environment in which jobs are performed — properties that the behaviours of leaders, peers, and organisational systems either support or undermine. The taxonomy's performance-building themes — particularly Empowering, Supportive, and Authentic — map directly onto the conditions that Hackman and Oldham identified as prerequisites for intrinsic motivation.

3.2 Culture as Repeated Behavioural Pattern

The second founding insight was definitional: culture is not an attitude, a value, or a belief. Culture is what people repeatedly do, tolerate, reward, and avoid. This definition, consistent with Schein's (1985) account of culture as enacted assumption and Bourdieu's (1977) account of habitus as repeated practice, has a specific implication for taxonomy development. If culture is defined by repetition, then the relevant unit of analysis is not the behaviour itself but the behavioural pattern — the frequency and regularity with which a behaviour occurs in organisational life.

This definitional choice shaped the taxonomy's entire classification logic. Rather than asking whether a behaviour is good or bad, the taxonomy asks what happens when a behaviour becomes normal. Rather than classifying behaviours into fixed positive and negative categories, it distinguishes between behaviours that can be practised regularly and still build trust, ownership, and performance, and behaviours that may have a legitimate occasional role but become damaging when they become habitual. This distinction — the frequency insight — is the taxonomy's central theoretical contribution and is examined in detail in Section 4.

Duhigg's (2012) account of the habit loop — the cue-routine-reward cycle through which individual behaviours become automatic — provides a mechanism for understanding how individual behavioural choices aggregate into cultural patterns. When a behaviour is regularly triggered by organisational cues and regularly reinforced by organisational rewards — even when those rewards are social rather than formal — it becomes part of the organisation's operating system. The taxonomy's classification of behaviours as performance-building or performance-limiting is, in these terms, a classification of the cultural operating systems that repeated behaviour produces.

4. The Above and Below Line Architecture and the Frequency Insight

4.1 The Binary and Its Theoretical Basis

The above and below line distinction is the structural heart of the taxonomy. It provides the sorting frame through which groups engage with the behavioural items and the visual architecture through which the culture map is constructed. Its design requires careful theoretical justification because it is superficially similar to the good/bad binaries of existing culture instruments while being fundamentally different in its logic.

The above and below line is not a moral distinction. It is a performance distinction — specifically, a distinction about the pattern of effects that repeated behaviour produces in organisational life. Above-the-line behaviours are those that can be practised regularly and consistently strengthen the conditions for high performance: trust, openness, ownership, collaboration, learning, and adaptability. Below-the-line behaviours are those that may have a legitimate occasional role but consistently weaken those conditions when they become habitual.

This distinction draws on the positive organisational behaviour tradition's emphasis on the conditions that enable human flourishing in work settings (Seligman, 2011; Cameron, Dutton and Quinn, 2003) and on the extensive literature demonstrating that psychological safety — the condition most reliably associated with team learning and performance — is undermined by precisely the below-the-line behavioural patterns the taxonomy identifies: fear, control, criticism, approval-seeking, and conformity (Edmondson, 1999; Frazier et al., 2017).

4.2 The Frequency Insight

The frequency insight is the taxonomy's most significant theoretical contribution. It emerged from a specific observation during early development: that many of the behaviours that appeared in the below-the-line category were not universally or unconditionally damaging. Following the rules may be important in regulated environments. Checking decisions with others can reflect appropriate humility. Avoiding unnecessary conflict can preserve relationships. Being critical can identify genuine risk. Taking control can be necessary in a genuine crisis. Staying on top of details can prevent failure.

The question was therefore not: is this behaviour always right or always wrong? The more useful question was: what happens when this behaviour is used too often? When does a behaviour that serves a legitimate occasional purpose become a cultural pattern that limits performance?

This reframing produced the taxonomy's central classification principle. Above-the-line behaviours are characterised by their sustainability: they can be practised regularly, even habitually, without diminishing the conditions for high performance. Indeed, the more consistently they are practised, the stronger those conditions become. Authenticity practised consistently builds trust. Empowerment practised consistently builds capability and ownership. Collaboration practised consistently builds collective intelligence.

Below-the-line behaviours are characterised by their contextual legitimacy but habitual toxicity. Approval-seeking used occasionally reflects social sensitivity. Used habitually, it reduces honesty and constrains genuine dialogue. Control exercised occasionally reflects appropriate accountability. Exercised habitually, it reduces ownership and initiative. Criticism applied occasionally improves quality. Applied habitually, it reduces creativity and psychological safety. Competition encouraged occasionally drives performance. Encouraged habitually, it undermines collaboration and trust.

This frequency-based classification logic connects to several converging theoretical traditions. Aristotle's account of virtue as the habitual practice of excellence — virtue is not a single act but a pattern of repeated choice — provides a philosophical antecedent. More proximate theoretical support comes from social norms theory (Cialdini and Trost, 1998), which establishes that repeated behaviour creates descriptive norms — implicit standards of what is normal and expected — that then constrain and shape the behaviour of others. A behaviour that begins as an individual choice

becomes, through repetition, a cultural norm that applies social pressure to everyone in the environment. The taxonomy's below-the-line behaviours are precisely those whose normalisation produces the most damaging social norms.

Argyris's (1990) defensive routines — the habitual patterns of behaviour through which organisations protect themselves from the discomfort of honest examination — map directly onto the below-the-line themes of the taxonomy. Conventional, Validating, Fearful, and Dependent behaviours are, in Argyris's terms, the building blocks of organisational defensive routines: the repeated patterns of avoidance, compliance, and approval-seeking that protect the status quo at the cost of learning and adaptation.

4.3 The Line as a Facilitation Device

The above and below line distinction serves not only as a classification principle but as a facilitation device — a structure that makes the sorting task manageable and the conversation generative. By providing a binary rather than a scale, it forces commitment: every card must go somewhere. There is no safe middle ground, no ability to rate something as 'somewhat positive' and avoid the conversation. The binary forces a position, and positions invite dialogue.

The line is also a literal physical feature of the facilitation space — a rope, a piece of tape, an imaginary threshold across the table — that gives the culture map a spatial dimension. Cards above the line are physically above the line. Cards below are physically below. The resulting map is visible, movable, and reworkable in ways that a scored profile on a printed report cannot be. This physicality is not incidental to the taxonomy's design. It is integral to the facilitation methodology that the taxonomy serves.

5. The Development Process

5.1 Phase One: Conceptual Development

The development of the Culture Science Cards taxonomy proceeded through three distinct phases over approximately five years of iterative practitioner inquiry. The first phase was conceptual: the establishment of the founding theoretical logic described in Sections 3 and 4. This phase produced the performance-through-people insight, the definition of culture as repeated behavioural pattern, and the above and below line architecture grounded in the frequency insight.

The conceptual phase also produced the first significant structural decision: that the taxonomy would distinguish between performance-building and performance-limiting themes rather than constructive and defensive styles, positive and negative orientations, or any of the other binary classifications used in existing frameworks. This terminology was chosen deliberately. Performance-building and performance-limiting are functional descriptions rather than moral ones — they describe what the behavioural patterns do to organisational performance rather than what kind of people exhibit them. This functional framing is consistent with the behaviour-first philosophy of the methodology and with the aspiration to create a taxonomy that produces systemic attribution rather than individual blame.

5.2 Phase Two: The 16 Cs Model and Its Limitations

The second phase of development began with the generation of candidate themes and items. An early iteration of the taxonomy was organised around a 16 Cs structure — a mnemonic architecture that provided simplicity and memorability. The performance-building Cs included Clarity, Contribution, Courage, Creativity, Collaboration, Care, Compassion, and Capability. The performance-limiting Cs included Compliance, Caution, Critical, Competitive, Control, and Conventional, among others.

The 16 Cs model served an important developmental function. Its structural elegance — the alliterative consistency of the theme names — made the taxonomy easy to hold in mind and communicate in workshop settings. Several of the Cs survived into the final model in recognisable form: Collaboration became the Collaborative theme, Creativity became the Creative theme, Care and Compassion merged into the Empathetic theme, Capability informed both the Goal Directed and Confident themes, Control became the Controlling theme, Critical and Competitive survived intact.

However, testing the 16 Cs model across real facilitation sessions revealed a fundamental limitation: the C-named themes were predominantly virtues and qualities rather than behaviours. Courage is an admirable quality. Compassion is a valued disposition. Clarity is a desirable characteristic. But none of these terms describes an observable, discussable, designable behaviour. You cannot see Courage. You can see someone backing themselves when the organisational culture would prefer compliance. You can see someone raising an uncomfortable truth in a meeting where silence would be safer. The virtue names pointed at the right territory but did not give groups the specific behavioural language they needed to have the conversations the methodology required.

This observation produced one of the most important insights of the development process: the distinction between virtue language and behavioural language as a critical design variable in culture facilitation instruments. Virtue language produces abstract agreement — everyone values Courage and Compassion in principle, and everyone can say so without committing to anything specific. Behavioural language produces specific, contestable, discussable conversation — because observable behaviours can be agreed with, challenged, reinterpreted, and placed above or below a line in ways that abstract virtues cannot.

This distinction has theoretical grounding in the implementation intentions literature (Gollwitzer, 1999), which demonstrates that specific behavioural intentions — when X happens, I will do Y — produce significantly higher follow-through than abstract goal intentions — I will be more courageous. The taxonomy's shift from virtue language to behavioural language is, in these terms, a shift from abstract goal intention to specific behavioural intention at the collective level — a shift that the implementation intentions research predicts will produce higher commitment and follow-through.

5.3 Phase Three: Field Validation and Iterative Refinement

The third phase of development was field validation: the testing of candidate items in real facilitation sessions and the iterative refinement of items based on observed conversational productivity. This phase produced the final sixteen-theme, eighty-item framework through a process of generation, testing, observation, and replacement that constitutes what this paper terms iterative content validity — the repeated cycle of empirical testing against the instrument's primary performance criterion.

For a psychometric instrument, the primary performance criterion is measurement accuracy: does the item reliably measure the construct it is designed to measure? For a facilitation instrument, the primary performance criterion is conversational productivity: does the item generate the quality and depth of dialogue that the methodology requires? These are different criteria and they produce different validation processes.

Items that consistently generated rich, contested, honest dialogue were retained. Items that consistently produced easy consensus without genuine discussion — items where groups agreed quickly and moved on without the kind of specific, story-driven, perspective-challenging conversation that culture change requires — were identified as conversationally underperforming and replaced. This iterative replacement process continued until the deck reached a stable state in which the great majority of items were reliably generating productive dialogue across diverse groups and organisational contexts.

The replacement of 'cooperate with others' with 'collaborate with others' exemplifies the iterative refinement process and deserves detailed examination. In facilitation sessions, some participants who encountered 'cooperate with others' interpreted it through the lens of their organisational experience in ways the item's designers had not anticipated. In many workplace contexts, cooperation carries a power asymmetry: the junior person cooperates with the senior person's wishes, the compliant person cooperates with the dominant person's preferences. Cooperation, in this reading, is compliance with a collegial face.

This interpretation — which participants were not misreading but reading accurately through their experience — was diverting conversation from the intended topic of joint working toward the unintended topic of subservience and compliance. The card was generating conversation, but not the right conversation. Collaboration, by contrast, implies mutual contribution: two or more parties bringing something to a shared endeavour as equals, neither subsuming their perspective into the other's. The replacement item restored the intended meaning while preserving the productive ambiguity — because 'collaborate with others' can still be contested: what does genuine collaboration look like in this organisation? When do we fall short of it? What gets in the way? — but redirected that ambiguity toward a more generative conversation.

This example illustrates what this paper terms the observed meaning-in-use principle of item revision. Drawing on Wittgenstein's (1953) argument that meaning is use — that words mean what they do in the contexts in which they are used rather than what dictionaries say they mean — the principle holds that item revision should be guided by observed participant interpretation rather than intended item meaning. Where participants consistently construct meanings that divert conversation away from the behaviour the item was designed to surface, the item should be revised to better align the lived meaning with the conversational purpose. The taxonomy's items are not fixed texts with authoritative interpretations. They are signs whose meaning is produced in the facilitation context — and item development is, in part, the process of refining those signs until the meanings they produce in context are the meanings that serve the methodology's purpose.

6. The Sixteen Themes and Their Theoretical Grounding

6.1 The Architecture of Eight Plus Eight

The final taxonomy comprises sixteen themes — eight performance-building and eight performance-limiting — each containing five behavioural items, for a total of eighty items across the deck. The decision to balance eight performance-building against eight performance-limiting themes encodes the task-people balance insight at the structural level of the taxonomy itself. It signals, through the architecture of the deck, that the model is not primarily a catalogue of dysfunction — it is an equal mapping of what high performance requires and what limits it.

This structural balance also addresses a specific problem identified in early testing: when groups were asked to sort cards above and below the line, they tended to place significantly more cards below the line. The limitation patterns were highly visible to people with organisational experience — they could easily recognise the fear, control, criticism, conformity, and approval-seeking that characterised their culture. The aspiration patterns were less immediately present. An unbalanced taxonomy — with more limiting than building themes — risked making the conversation feel heavy, overwhelming, and demoralising rather than generative and forward-looking.

The eight-eight balance ensures that the culture map that emerges from a card sort is visually and proportionally balanced: equal territory mapped above and below the line. This visual balance has a psychological effect in the facilitation room. It signals that the process is as interested in what the organisation wants to build as in what it wants to change. It makes the below-the-line conversation feel proportionate rather than dominant. And it ensures that the aspiration gap — the distance between the performance-building behaviours the group aspires to and the performance-limiting patterns they currently experience — is visible in the physical layout of the sorted deck.

6.2 The Performance-Building Themes

The eight performance-building themes describe behavioural orientations that can be practised regularly and consistently strengthen the conditions for high performance. Each theme is grounded in converging research traditions in organisational behaviour, positive psychology, and leadership theory.

Theme	Theoretical Grounding	Sample Items
Goal Directed	McClelland (1961) achievement motivation; Locke and Latham (1990) goal setting theory; Duhigg (2012) habit formation	Look for opportunities to be better; Plan ahead; Set challenging goals; Work for a sense of achievement; Consider options before acting
Creative	Amabile (1996) creativity and intrinsic motivation; Csikszentmihalyi (1990) flow; Argyris (1990) double-loop learning	Challenge current conventions; Experiment; Seek out feedback; Use conflict to generate ideas; Be curious
Authentic	Avolio and Gardner (2005) authentic leadership; Brown and Mitchell (2010) ethical leadership; Edmondson (1999) psychological safety	Be transparent; Tell it like it is; Maintain personal integrity; Share feelings and thoughts openly; Be honest even when uncomfortable
Confident	Bandura (1997) self-efficacy; Thaler and Sunstein (2008) choice architecture; Deci and Ryan (1985) self-determination	Be flexible and adaptive; Take considered risks; Back themselves; Be enthusiastic; Quality over quantity

Theme	Theoretical Grounding	Sample Items
Em-powering	Spreitzer (1995) psychological empowerment; Hackman and Oldham (1976) job characteristics; Vygotsky (1978) scaffolding	Help people grow and develop; Help people think for themselves; Involve people in decisions that affect them; Trust others; Share responsibility
Supportive	Edmondson (1999) psychological safety; Cooperider and Srivastva (1987) appreciative inquiry; Deci and Ryan (1985) relatedness	Encourage people to be their best; Recognise people's efforts; Motivate others through friendliness; Make time for people; Ask challenging questions
Collaborative	Habermas (1984) communicative action; Putnam (1995) social capital; Pierce, Kostova and Dirks (2003) psychological ownership	Collaborate with others; Use diplomacy to engage; Work with others in friendly ways; Focus on shared goals; Have unconditional respect for others
Empathetic	Goleman (1995) emotional intelligence; Noddings (1984) ethics of care; Rogers (1961) unconditional positive regard	Show concern for the needs of everyone; Be open and warm towards others; Put relationships ahead of results; Seek to understand; Be aware of people's emotions

6.3 The Performance-Limiting Themes

The eight performance-limiting themes describe behavioural orientations that may have a legitimate occasional role in organisational life but consistently weaken the conditions for high performance when they become habitual patterns. Each theme is grounded in research traditions that have documented its damaging effects at the organisational level.

Theme	Theoretical Grounding	Sample Items
Dependent	Argyris (1990) defensive routines; Seligman (1975) learned helplessness; Langer (1989) mindlessness	Accept goals without challenging them; Loyalty above all else; Check decisions with others; Do what's expected; Keep those in authority happy
Conventional	Asch (1951) conformity; Janis (1972) groupthink; Bourdieu (1977) habitus and social reproduction	Always follow the rules; Do not rock the boat; Focus on making a good impression; If it ain't broke don't fix it; Don't upset other people
Validating	Cialdini (1984) social proof and liking; Festinger (1954) social comparison; Baumeister (1991) need to belong	Be everyone's friend; Change priorities to please others; Do things for others' approval; Go along with others; Never disagree with others
Fearful	Edmondson (1999) psychological safety; Lazarus (1991) stress and coping; Argyris (1990) defensive routines	Don't take chances; Pass responsibilities to others; Push decisions up; Take cover when things get tough; Avoid conflict at all costs
Controlling	McClelland (1975) power motivation; Kipnis (1972) use of power; French and Raven (1959) bases of social power	Be forceful; Be stern but fair; Demand people are loyal; Keep control of everything; Play politics
Critical	Argyris (1990) single-loop learning; Dweck (2006) fixed mindset; Baumeister et al. (2001) bad is stronger than good	Never settle for good enough; Find flaws; Focus on mistakes; Oppose anything new and different; Question decisions made by others
Competitive	Deutsch (1949) competitive interdependence; Kohn (1986) no contest; Axelrod (1984) evolution of co-operation	Be better than everyone else; Be seen and noticed; Put individual success ahead of team success; Never be wrong; There are only winners and losers

Theme	Theoretical Grounding	Sample Items
Demand-ing	Maslach and Leiter (1997) burnout; Karasek (1979) job demands-control model; Cooper, Dewe and O'Driscoll (2001) organisational stress	Directly take care of everything; Ensure work is the number one priority; Errors are unacceptable; Stay on top of everything; Work very long hours

7. Item Selection Criteria and Field Validation

7.1 The Three Selection Criteria

Item selection for the Culture Science Cards taxonomy followed a two-stage process. In the first stage, candidate items were generated from multiple sources — brainstorming sessions, organisational behaviour literature, facilitation experience, and AI-assisted exploration — and curated against three criteria. In the second stage, selected items were tested in real facilitation sessions and evaluated against their primary performance criterion: the quality and depth of dialogue they generated.

The three first-stage criteria were developed iteratively through the development process and reflect the specific requirements of a conversation-oriented rather than measurement-oriented instrument.

7.1.1 Recognisability

The first criterion was recognisability: would practitioners immediately connect the item to observable workplace behaviour that they had personally witnessed or experienced? An item that requires explanation before it can be engaged with is an item that interrupts the flow of the card sort and redirects attention from culture to vocabulary. Items were selected for their capacity to produce immediate recognition — the response of 'yes, I know exactly what that looks like' that activates the personal organisational experience the methodology depends on.

Recognisability is connected to what cognitive psychologists call availability (Tversky and Kahneman, 1973): the ease with which relevant examples come to mind. An item with high recognisability produces immediate availability of specific examples from the participant's experience. Those examples are the raw material of the conversation. An item with low recognisability produces confusion or abstraction — neither of which serves the methodology's purpose.

7.1.2 Practical Specificity

The second criterion was practical specificity: was the behaviour described specific enough to be observable, discussable, and designable? This criterion drove the shift from virtue language to behavioural language described in Section 5.2. Items were evaluated for their capacity to generate specific, concrete conversation about what the behaviour looks like in practice, when it occurs, and what would need to change for it to become more or less frequent.

Practical specificity connects to Gollwitzer's (1999) implementation intentions research and to the behaviour change literature more broadly (Michie, van Stralen and West, 2011): specific behavioural intentions are more actionable than abstract goal intentions. An item that names a specific, observable behaviour gives groups something concrete to design toward in the methodology's design phase. An item that names an abstract virtue or quality gives groups only an aspiration — which is where culture change most often stalls.

7.1.3 Productive Ambiguity

The third criterion was productive ambiguity: did the item carry sufficient interpretive openness to generate genuine dialogue rather than easy consensus? This criterion reflects the taxonomy's fundamental design principle: items should be specific enough to recognise but open enough to contest. The ideal item is one that different participants, drawing on different organisational experiences and different professional perspectives, will place differently above or below the line — and whose divergent placement produces a conversation that illuminates the culture in ways that neither participant's individual perspective could produce alone.

Productive ambiguity is theoretically grounded in the disequilibrium principle of constructivist learning theory (Piaget, 1954; Festinger, 1957): genuine learning requires encountering something that does not fit neatly into existing cognitive categories, creating a productive discomfort that motivates deeper engagement. Items that produce easy consensus — where everyone immediately agrees on the placement — may still serve an anchoring function in the card sort, but they do not themselves generate the disequilibrium that produces the richest learning.

The productive ambiguity criterion also draws on Wittgenstein's (1953) meaning-as-use principle: the most productive items are those whose meaning is sufficiently open that different participants can bring their own organisational experience to bear on them, producing different interpretations that, when discussed, generate collective insight. An item with fixed meaning produces consensus. An item with open meaning produces conversation.

7.2 The Iterative Field Validation Process

The second stage of item selection was field validation: the testing of selected items in real facilitation sessions and the replacement of items that failed to meet the conversational productivity criterion. This process constitutes what this paper terms iterative content validity — a form of validation appropriate to facilitation instruments whose primary performance criterion is conversational productivity rather than measurement accuracy.

Items were evaluated in facilitation sessions through systematic observation of the quality and depth of dialogue they generated. An item was identified as conversationally underperforming when it consistently produced one of three failure modes. The first was easy agreement: participants placed the card quickly and without genuine discussion, the placement reflecting social consensus rather than considered individual judgment. The second was disengagement: participants were uncertain what the item meant in practice and moved on without generating specific, experience-grounded conversation. The third was misdirection: participants engaged with the item but the conversation it generated was about something other than the behaviour the item was designed to surface — as in the cooperate/collaborate example described in Section 5.3.

When an item was identified as conversationally underperforming, it was replaced through a process of candidate generation and curatorial selection, with the replacement item evaluated against the same three criteria and then field-tested in subsequent sessions. This replacement cycle continued until the deck reached a stable state in which the great majority of items were reliably generating productive dialogue. The deck's current composition reflects this iterative refinement: items that survived multiple rounds of testing across diverse groups and organisational contexts.

The iterative field validation process has a specific implication for how the taxonomy's validity should be understood and claimed. The Culture Science Cards items have not been subjected to

formal psychometric validation — no factor analysis has been conducted, no inter-rater reliability coefficients have been calculated, no normative database has been established. This is not an oversight. It reflects the deliberate choice to develop a conversation instrument rather than a measurement instrument, and to validate it against the criterion appropriate to its purpose.

The taxonomy's claim to validity is a claim of content validity in the specific sense: the items validly represent the behavioural domain they are designed to surface, as evidenced by their consistent capacity to generate rich, honest, productive dialogue about that domain across diverse organisational contexts. This is a form of validity that the psychometric literature recognises (Messick, 1995; Kane, 2006) but that has been inadequately developed for facilitation instruments. The methodology described here represents a contribution to that development.

7.3 The Five Items Per Theme

The decision to include four to six items per theme — standardised at five for consistency across the final deck — reflects both practical and theoretical considerations. Practically, five items per theme produces a deck of eighty cards that sits at the upper boundary of what group sorting tasks can handle without cognitive overload. Sweller's (1988) cognitive load theory establishes that working memory can process approximately seven plus or minus two discrete items simultaneously. Five items per theme means groups engage with five cards at a time within a thematic conversation — comfortably within working memory limits. The thematic grouping of the deck reduces the cognitive load of the overall sort by providing an organising structure that makes eighty items manageable.

Theoretically, five items per theme provides sufficient coverage of the theme's behavioural territory without collapsing the theme into a single dimension. Each set of five items is designed to cover different aspects, levels, or expressions of the theme — from more obvious to more subtle, from more interpersonal to more organisational, from more individual to more collective. This internal diversity within each theme ensures that the five items generate five distinct conversations rather than five versions of the same conversation.

The five-item structure also ensures that no single item dominates the theme. A group that finds one item in a set unrecognisable or inapplicable to their context still has four other items through which to engage with the theme. The redundancy built into the five-item structure is a resilience feature: it ensures that the theme is accessible to groups with diverse organisational contexts and diverse facilitation dynamics.

8. The Physical Design Evolution of the Cards

8.1 The Card as Material Artefact

The Culture Science Cards are physical objects — cards held in hands, dealt across tables, placed above and below lines, reworked and debated and moved. The physical design of the cards is not separate from the taxonomy's development. It is part of it. Changes to the physical design of the cards have produced observable changes in how groups interact with them and what those interactions produce. Two design evolutions are particularly significant: the introduction of sequential numbering and colour coding, and the consequent activation of the card's back as a facilitation resource.

8.2 From Random to Sequential Numbering

In the initial version of the deck, cards were numbered randomly — a deliberate choice intended to prevent participants from treating the deck as a sequence to be worked through in order and to preserve the randomness of the dealing mechanism. The random numbering was consistent with the methodology's design principle of preventing preparation and strategic engagement with the cards.

However, facilitation observation revealed that participants were not using the backs of the cards at all. The back of the card — which carried the random number and theme category — was being ignored. Participants engaged entirely with the front, where the behavioural item was printed, and the back served no function in the facilitation dynamic.

This observation prompted a design revision. The cards were renumbered sequentially — one through eighty, corresponding to a consistent ordering across the deck — and colour-coded by theme category. This change activated the back of the card as a facilitation resource, creating what the methodology terms a second entry point for quality discussions.

8.3 The Second Entry Point

The sequential numbering and colour coding of the card backs produces at least four distinct facilitation affordances that the randomly numbered deck did not provide.

The first is facilitator navigation. A facilitator listening at the meta level — attending to the cultural disclosures that emerge in how cards are discussed rather than to their placement — can note card numbers rather than having to remember or transcribe full card text. 'Cards 23, 47, and 61 all produced similar responses' is a navigational reference that allows the facilitator to return efficiently to specific moments in the session without losing the group's thread.

The second is cross-session continuity. In the multi-session model that the methodology has evolved toward, participants can reference specific cards across sessions: 'we placed card 34 above the line last time but I'm not sure we'd do that now.' This cross-session referencing enables the longitudinal tracking of cultural movement — the visible record of what has shifted and what has not — that the iterative Envision-Diagnose-Design cycle depends on.

The third is visual pattern recognition. The colour coding of cards by theme category means that the physical layout of the sorted deck produces a visual culture map in which thematic patterns are immediately apparent. A concentration of red cards below the line — if red denotes the Controlling theme — is visible at a glance without requiring the facilitator or group to read every card. This visual pattern is itself a discussion prompt: 'I notice a lot of the same colour below the line — what do you make of that?' The colour system transforms the sorted deck from a collection of individual card placements into a coherent visual representation of the culture's thematic profile.

The fourth is participant orientation. The combination of number and colour on the back of the card allows any participant to reorient the group to a specific card at any moment in the session without having to locate and re-read the card text. 'Can we go back to the blue card, number 49?' is a more efficient navigational reference than 'can we go back to the one about keeping control of everything?' This efficiency reduces the friction of returning to cards mid-session and increases the group's capacity to pursue the lines of inquiry that the most productive conversations open up.

The physical design evolution of the card backs illustrates a broader principle of the development process: that the tool was refined through observation of how it was actually used rather than through prior specification of how it should be used. The random numbering was a reasonable design choice given the methodology's principles. The observation that participants were ignoring the card backs revealed that the design was leaving a resource unutilised. The sequential numbering and colour coding was the response — a design change that activated an existing physical feature of the card for new facilitation purposes. This is Schön's (1983) reflective practice in its most concrete form: knowledge generated from the tool's use, fed back into the tool's design.

9. Implications and Future Research

9.1 Implications for Taxonomy Development

The development logic described in this paper has implications for how behavioural taxonomies for culture facilitation instruments should be developed more broadly. The three-criteria item selection framework — recognisability, practical specificity, and productive ambiguity — provides a systematic basis for item evaluation that is appropriate to the dialogue-oriented purpose of facilitation instruments and that could be applied to the development of other instruments in this category.

The iterative content validity methodology — field testing items against the conversational productivity criterion and replacing underperforming items — provides a validation approach that is both rigorous and appropriate for instruments whose purpose is dialogue generation rather than psychometric measurement. The observed meaning-in-use principle of item revision provides a specific guide for identifying and addressing item failure: when participants consistently construct meanings that divert conversation away from the behaviour an item was designed to surface, the item should be revised to align lived meaning with conversational purpose.

The frequency insight provides a classification logic that could inform the development of other culture taxonomies seeking to move beyond absolute good/bad classification toward more nuanced, context-sensitive, pattern-oriented frameworks. The question 'what happens when this behaviour becomes normal?' is a more sophisticated and more practically useful classification question than 'is this behaviour good or bad?' — and its application to taxonomy development could produce frameworks with greater facilitation value than the binary classifications that currently dominate the field.

9.2 Implications for Practice

For practitioners of organisational culture facilitation, the taxonomy's development story has several practical implications. The shift from virtue language to behavioural language is immediately applicable: facilitators who describe desired culture in terms of observable behaviours rather than abstract values or qualities will consistently produce more specific, actionable, and honest conversations than those who work with virtue language. The difference between 'we value courage' and 'we want people to back themselves even when the organisation would prefer compliance' is the difference between an aspiration and a behavioural commitment.

The three item selection criteria provide a practical framework for evaluating the conversational quality of any prompt used in culture facilitation — not only card items but workshop questions,

survey items, and discussion prompts of all kinds. Does the prompt produce immediate recognition? Does it name a specific, observable behaviour? Does it carry sufficient interpretive openness to generate genuine dialogue? Items that fail any of these criteria are likely to produce the easy consensus, disengagement, or misdirection that constitute conversational underperformance.

The physical design evolution of the cards illustrates the broader principle that facilitation tools should be observed in use and refined based on what is observed. Every facilitation session is a data source. The facilitator who attends not only to the content of the conversation but to how the tool is being used — which features are engaged with, which are ignored, which produce the richest responses — has access to a continuous stream of design feedback that can inform iterative improvement.

9.3 Future Research Directions

The taxonomy described in this paper has not been subjected to formal psychometric validation. This is appropriate given its dialogue-oriented purpose, but it does not preclude formal validation of a different kind. Several research directions are identified as priorities.

First, a formal assessment of the taxonomy's content validity by independent expert reviewers — organisational behaviour scholars, experienced culture practitioners, and facilitation specialists — would provide an external check on the item selection judgments described in this paper. Do independent experts agree that the items validly represent the behavioural themes they are assigned to? Do they identify items that should be added, removed, or revised? This expert review process would constitute a significant step toward establishing the taxonomy's formal content validity.

Second, a systematic study of conversational productivity across items — using structured observation protocols applied by independent researchers across a large sample of facilitation sessions — would provide empirical evidence for the item quality claims made in this paper. Which items consistently generate the richest dialogue? Which underperform? Do patterns of conversational productivity vary across organisational contexts, hierarchical levels, or cultural backgrounds? These questions have significant implications for both the taxonomy's development and the methodology's practice.

Third, the cross-cultural validity of the taxonomy deserves specific investigation. The items were developed primarily in New Zealand organisational contexts. Hofstede's (1980, 2001) research on cultural dimensions establishes that behavioural norms vary significantly across national cultures — what is recognised as approval-seeking in a low power distance culture may be experienced as appropriate respect in a high power distance one. Cross-cultural testing of the taxonomy's items would establish the boundaries of their applicability and identify items that require cultural adaptation.

Fourth, the relationship between the taxonomy's sixteen themes and the established theoretical constructs in the organisational behaviour literature — psychological safety, self-determination, achievement motivation, defensive routines, and others — deserves formal investigation. Do groups that place more items above the line in the performance-building categories score higher on validated measures of these constructs? This convergent validity investigation would situate the taxonomy within the broader empirical literature and provide evidence for its theoretical claims.

10. Conclusion

The Culture Science Cards taxonomy began with a question — what really drives organisational performance? — and arrived, through five years of iterative practitioner development, at an answer expressed in eighty specific, observable, discussable behavioural items organised into sixteen themes. The journey from question to taxonomy was not a straight line. It passed through the performance-through-people insight, the above and below line architecture, the frequency re-frame, the 12 Cs model and its virtue-language limitations, the shift to behavioural specificity, the field testing and iterative replacement of underperforming items, and the physical design evolution that activated the back of the card as a facilitation resource.

Each stage of that journey produced a specific theoretical insight. The frequency insight — that the most useful question about any organisational behaviour is what happens when it becomes normal — reframes culture taxonomy from moral classification to pattern recognition. The virtue-to-behaviour distinction identifies behavioural specificity as the critical design variable that separates conversation instruments from values instruments. The three-criteria item selection framework provides a systematic basis for developing items that are recognisable, practically specific, and productively ambiguous. The iterative content validity methodology provides a validation approach appropriate to facilitation instruments whose purpose is dialogue generation. And the observed meaning-in-use principle of item revision provides a specific guide for aligning lived meaning with conversational purpose.

Together, these contributions constitute a methodological framework for developing dialogue-oriented culture instruments that is distinct from the psychometric tradition that has dominated the field. The Culture Science Cards taxonomy is not a better measurement instrument. It is a different kind of instrument entirely — one built on the conviction that the most important thing a culture instrument can do is not produce a score but open a conversation. The development logic described in this paper is the logic of building an instrument that does exactly that.

Culture is what people repeatedly do, tolerate, reward, and avoid. Because those repetitions can be named, discussed, and chosen, culture can be changed. The taxonomy exists to make those repetitions visible — and in making them visible, to make them available for the collective reflection and deliberate redesign that genuine culture change requires.

References

- Amabile, T.M. (1996) *Creativity in Context*. Boulder, CO: Westview Press.
- Argyris, C. (1990) *Overcoming Organizational Defenses: Facilitating Organizational Learning*. Boston: Allyn and Bacon.
- Aristotle (350 BCE) *Nicomachean Ethics*. Translated by W.D. Ross. Oxford: Oxford University Press, 1998.
- Asch, S.E. (1951) 'Effects of group pressure upon the modification and distortion of judgments,' in H. Guetzkow (ed.) *Groups, Leadership and Men*. Pittsburgh: Carnegie Press, pp. 177-190.
- Avolio, B.J. and Gardner, W.L. (2005) 'Authentic leadership development: Getting to the root of positive forms of leadership,' *The Leadership Quarterly*, 16(3), pp. 315-338.

- Axelrod, R. (1984) *The Evolution of Cooperation*. New York: Basic Books.
- Bandura, A. (1997) *Self-Efficacy: The Exercise of Control*. New York: Freeman.
- Baumeister, R.F., Bratslavsky, E., Finkenauer, C. and Vohs, K.D. (2001) 'Bad is stronger than good,' *Review of General Psychology*, 5(4), pp. 323-370.
- Baumeister, R.F. and Leary, M.R. (1995) 'The need to belong: Desire for interpersonal attachments as a fundamental human motivation,' *Psychological Review*, 117(3), pp. 497-529.
- Blake, R.R. and Mouton, J.S. (1964) *The Managerial Grid*. Houston: Gulf Publishing.
- Bourdieu, P. (1977) *Outline of a Theory of Practice*. Cambridge: Cambridge University Press.
- Brown, M.E. and Mitchell, M.S. (2010) 'Ethical and unethical leadership: Exploring new avenues for future research,' *Business Ethics Quarterly*, 20(4), pp. 583-616.
- Bushe, G.R. and Marshak, R.J. (eds.) (2015) *Dialogic Organization Development: The Theory and Practice of Transformational Change*. San Francisco: Berrett-Koehler.
- Cameron, K.S., Dutton, J.E. and Quinn, R.E. (eds.) (2003) *Positive Organizational Scholarship: Foundations of a New Discipline*. San Francisco: Berrett-Koehler.
- Cameron, K.S. and Quinn, R.E. (2011) *Diagnosing and Changing Organizational Culture*. 3rd edn. San Francisco: Jossey-Bass.
- Cialdini, R.B. (1984) *Influence: The Psychology of Persuasion*. New York: William Morrow.
- Cialdini, R.B. and Trost, M.R. (1998) 'Social influence: Social norms, conformity and compliance,' in D.T. Gilbert, S.T. Fiske and G. Lindzey (eds.) *The Handbook of Social Psychology*. 4th edn. New York: McGraw-Hill, pp. 151-192.
- Cooper, C.L., Dewe, P.J. and O'Driscoll, M.P. (2001) *Organizational Stress: A Review and Critique of Theory, Research, and Applications*. Thousand Oaks, CA: Sage.
- Cooperrider, D.L. and Srivastva, S. (1987) 'Appreciative inquiry in organizational life,' *Research in Organizational Change and Development*, 1, pp. 129-169.
- Csikszentmihalyi, M. (1990) *Flow: The Psychology of Optimal Experience*. New York: Harper and Row.
- Deci, E.L. and Ryan, R.M. (1985) *Intrinsic Motivation and Self-Determination in Human Behavior*. New York: Plenum.
- Denison, D.R. (1990) *Corporate Culture and Organizational Effectiveness*. New York: Wiley.
- Deutsch, M. (1949) 'A theory of cooperation and competition,' *Human Relations*, 2(2), pp. 129-152.
- DeVellis, R.F. (2016) *Scale Development: Theory and Applications*. 4th edn. Thousand Oaks, CA: Sage.
- Duhigg, C. (2012) *The Power of Habit: Why We Do What We Do in Life and Business*. New York: Random House.
- Dweck, C.S. (2006) *Mindset: The New Psychology of Success*. New York: Random House.

- Edmondson, A.C. (1999) 'Psychological safety and learning behavior in work teams,' *Administrative Science Quarterly*, 44(2), pp. 350-383.
- Emery, F.E. and Trist, E.L. (1960) 'Socio-technical systems,' in C.W. Churchman and M. Verhulst (eds.) *Management Science, Models and Techniques*. Volume 2. Oxford: Pergamon, pp. 83-97.
- Festinger, L. (1954) 'A theory of social comparison processes,' *Human Relations*, 7(2), pp. 117-140.
- Festinger, L. (1957) *A Theory of Cognitive Dissonance*. Stanford: Stanford University Press.
- Fleishman, E.A. (1953) 'The description of supervisory behavior,' *Journal of Applied Psychology*, 37(1), pp. 1-6.
- Frazier, M.L., Fainshmidt, S., Klinger, R.L., Pezeshkan, A. and Vacheva, V. (2017) 'Psychological safety: A meta-analytic review and extension,' *Personnel Psychology*, 70(1), pp. 113-165.
- French, J.R.P. and Raven, B. (1959) 'The bases of social power,' in D. Cartwright (ed.) *Studies in Social Power*. Ann Arbor: University of Michigan Press, pp. 150-167.
- Goleman, D. (1995) *Emotional Intelligence: Why It Can Matter More Than IQ*. New York: Bantam.
- Gollwitzer, P.M. (1999) 'Implementation intentions: Strong effects of simple plans,' *American Psychologist*, 54(7), pp. 493-503.
- Hackman, J.R. and Oldham, G.R. (1976) 'Motivation through the design of work: Test of a theory,' *Organizational Behavior and Human Performance*, 16(2), pp. 250-279.
- Habermas, J. (1984) *The Theory of Communicative Action*, Volume 1. Boston: Beacon Press.
- Hinkin, T.R. (1998) 'A brief tutorial on the development of measures for use in survey questionnaires,' *Organizational Research Methods*, 1(1), pp. 104-121.
- Hofstede, G. (1980) *Culture's Consequences: International Differences in Work-Related Values*. Beverly Hills, CA: Sage.
- Hofstede, G. (2001) *Culture's Consequences: Comparing Values, Behaviors, Institutions and Organizations Across Nations*. 2nd edn. Thousand Oaks, CA: Sage.
- Janis, I.L. (1972) *Victims of Groupthink*. Boston: Houghton Mifflin.
- Kane, M.T. (2006) 'Validation,' in R.L. Brennan (ed.) *Educational Measurement*. 4th edn. Westport, CT: American Council on Education/Praeger, pp. 17-64.
- Karasek, R.A. (1979) 'Job demands, job decision latitude, and mental strain: Implications for job redesign,' *Administrative Science Quarterly*, 24(2), pp. 285-308.
- Kipnis, D. (1972) 'Does power corrupt?,' *Journal of Personality and Social Psychology*, 24(1), pp. 33-41.
- Kohn, A. (1986) *No Contest: The Case Against Competition*. Boston: Houghton Mifflin.
- Kotter, J.P. and Heskett, J.L. (1992) *Corporate Culture and Performance*. New York: Free Press.
- Langer, E.J. (1989) *Mindfulness*. Reading, MA: Addison-Wesley.
- Lazarus, R.S. (1991) *Emotion and Adaptation*. New York: Oxford University Press.

- Locke, E.A. and Latham, G.P. (1990) *A Theory of Goal Setting and Task Performance*. Englewood Cliffs, NJ: Prentice Hall.
- Luthans, F. (2002) 'The need for and meaning of positive organizational behavior,' *Journal of Organizational Behavior*, 23(6), pp. 695-706.
- Maslach, C. and Leiter, M.P. (1997) *The Truth About Burnout*. San Francisco: Jossey-Bass.
- McClelland, D.C. (1961) *The Achieving Society*. Princeton: Van Nostrand.
- McClelland, D.C. (1975) *Power: The Inner Experience*. New York: Irvington.
- McGregor, N. (2025) 'Envision, Diagnose, Design: A Behaviour-First Dialogue Methodology for Organisational Culture Change.' Unpublished manuscript. culturesciencecards.com.
- Messick, S. (1995) 'Validity of psychological assessment,' *American Psychologist*, 50(9), pp. 741-749.
- Michie, S., van Stralen, M.M. and West, R. (2011) 'The behaviour change wheel: A new method for characterising and designing behaviour change interventions,' *Implementation Science*, 6(1), p. 42.
- Noddings, N. (1984) *Caring: A Feminine Approach to Ethics and Moral Education*. Berkeley: University of California Press.
- Piaget, J. (1954) *The Construction of Reality in the Child*. New York: Basic Books.
- Pierce, J.L., Kostova, T. and Dirks, K.T. (2003) 'The state of psychological ownership: Integrating and extending a century of research,' *Review of General Psychology*, 7(1), pp. 84-107.
- Putnam, R.D. (1995) 'Bowling alone: America's declining social capital,' *Journal of Democracy*, 6(1), pp. 65-78.
- Quinn, R.E. and Rohrbaugh, J. (1983) 'A spatial model of effectiveness criteria: Towards a competing values approach to organizational analysis,' *Management Science*, 29(3), pp. 363-377.
- Reason, P. and Bradbury, H. (eds.) (2008) *Handbook of Action Research*. 2nd edn. London: Sage.
- Rogers, C.R. (1961) *On Becoming a Person*. Boston: Houghton Mifflin.
- Schön, D.A. (1983) *The Reflective Practitioner: How Professionals Think in Action*. New York: Basic Books.
- Seligman, M.E.P. (1975) *Helplessness: On Depression, Development, and Death*. San Francisco: Freeman.
- Seligman, M.E.P. (2011) *Flourish: A Visionary New Understanding of Happiness and Well-Being*. New York: Free Press.
- Spreitzer, G.M. (1995) 'Psychological empowerment in the workplace: Dimensions, measurement, and validation,' *Academy of Management Journal*, 38(5), pp. 1442-1465.
- Sweller, J. (1988) 'Cognitive load during problem solving: Effects on learning,' *Cognitive Science*, 12(2), pp. 257-285.
- Thaler, R.H. and Sunstein, C.R. (2008) *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New Haven: Yale University Press.

Trist, E.L. and Bamforth, K.W. (1951) 'Some social and psychological consequences of the Longwall method of coal-getting,' *Human Relations*, 4(1), pp. 3-38.

Tversky, A. and Kahneman, D. (1973) 'Availability: A heuristic for judging frequency and probability,' *Cognitive Psychology*, 5(2), pp. 207-232.

Vygotsky, L.S. (1978) *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA: Harvard University Press.

Weick, K.E. (1995) *Sensemaking in Organizations*. Thousand Oaks, CA: Sage.

Wittgenstein, L. (1953) *Philosophical Investigations*. Oxford: Blackwell.

— *End of Paper* —

© 2025 Neil McGregor / Culture Science Cards

culturesciencecards.com

Page 1 of 37