Workplace Wellbeing: Foundations

THE TALKING THERAPIES UK PRACTITIONER SERIES



Workplace Wellbeing: Foundations

NICE NG212 and a defensible alternative to generic mental health awareness training

Foundation level · 8.5 hours of continuing professional development

TALKING THERAPIES UK

First edition · 2026

Published by Talking Therapies UK, talkingtherapies.co.uk.

Copyright © 2026 Talking Therapies UK. All rights reserved.

The moral right of the author has been asserted. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations in a review, and except that a purchaser of the accompanying course may print a single copy for their own study.

First edition, 2026.

This manual accompanies the Workplace Wellbeing: Foundations course in the Talking Therapies UK Digital Classroom, from which its content is drawn and against which it is kept current.

Revision 6. This printing was set on 22 November 2026 from the course as it stood that day.

A note on the clinical material. Every case in this book is a fictional composite written for teaching. No real person appears in it, however disguised. Nothing here constitutes clinical supervision, and nothing in it replaces the judgement of a qualified practitioner working within their competence, their professional body's code and their local governance. Where guidance is cited it was current at the date of this edition; readers must check the standing of any guidance before relying on it.

Set in Poppins and Roboto, the Talking Therapies UK typefaces. Figures drawn by the Talking Therapies UK teaching-diagram engine.

Contents

How to use this manual **1**

Part 1 Why an evidence-led alternative to generic awareness training **4**

1. The limits of Mental Health First Aid's evidence base 4
2. NICE NG212 as a firmer evidence anchor 12
3. Reading Evidence Claims in Workplace Wellbeing: A Beginner's Guide to Knowing What Counts 18
4. The Job Demands-Resources model and the HSE Management Standards: the theoretical architecture behind NG212 26
5. Mental health as a continuum: flourishing, languishing and the limits of an illness-only model 33
6. The economic and business case: presenteeism, the Thriving at Work review, and what the numbers can and cannot tell you 40

How to use this manual

This foundation course offers an honest, evidence-led introduction to workplace mental wellbeing training, grounded in the recommendations of NICE guideline NG212 — the National Institute for Health and Care Excellence's dedicated guidance on mental wellbeing at work, published in 2022 — rather than in the generic mental health awareness model that the majority of providers continue to default to. It is designed for practitioners, trainers, HR professionals and organisational consultants who want to build a workplace offer they can defend on clinical and empirical grounds, and who are prepared to engage candidly with what the current evidence actually says about the most widely recognised brand in this space.

The course opens by examining why generic awareness training has become so dominant in UK workplaces and what its honest limitations are. You will study in detail the 2023 Cochrane systematic review of Mental Health First Aid, which analysed twenty-one randomised controlled trials involving more than twenty-two thousand participants and found only very low-certainty evidence for any effect on recipients' mental health and wellbeing, concluding that firm conclusions could not be drawn in either direction. Understanding that finding — what it means, what it does not mean, and how to communicate it professionally to commissioners and employers — is treated as a foundational competence throughout the programme. The module does not position Mental Health First Aid as worthless; it positions honest uncertainty as the starting point for a more rigorous conversation.

From that foundation, the course moves to what NICE NG212 actually recommends and why those recommendations represent a firmer anchor for a workplace wellbeing offer. You will study the guideline's emphasis on organisational-level assessment and action, on manager training in supportive conversations and early identification of distress, and on the structural conditions that either protect or erode employee mental health. NICE NG212 is not simply a checklist; it is a framework that locates individual-level support within a broader organisational ecology, and you will learn to articulate that distinction clearly when speaking with clients.

The second half of the course focuses on positioning. You will work through how to present a Talking Therapies UK workplace pathway honestly and compellingly — not by overstating outcomes or borrowing the authority of credentials the evidence does not support, but by making a transparent case grounded in what NICE recommends and what the research can and cannot yet confirm. This includes practical guidance on how to respond when a prospective client asks why they should choose this approach over a more familiar branded programme, and how to frame differentiation in terms of evidence quality rather than marketing claims alone.

No prior experience of workplace mental health delivery is assumed, and the course carries eight and a half hours of continuing professional development. It is suitable as a standalone introduction or as the entry point to the broader Talking Therapies UK workplace wellbeing pathway.

This manual is the written form of a course taught in the Talking Therapies UK Digital Classroom. It is set at foundation level and carries 8.5 hours of continuing professional development for practitioners who complete the accompanying course and its assessment.

It is written to be read in order. Each part opens with what it is for; each chapter teaches one idea to a working depth and closes with the points worth keeping and a practice task that converts reading into something you do. The tasks are the part most often skipped and the part that most reliably changes practice, and several of them are designed to be taken to supervision rather than completed alone.

The figures are teaching diagrams rather than decoration. Each is the conventional form of a model you will meet in the literature and in supervision, drawn in one visual language across the whole series so that the shape of an idea is recognisable wherever it reappears.

Where a chapter draws on published work, the reference is named in the text and gathered in the further reading at the back. Terms of art are defined in the glossary, which is worth reading straight through once before you begin.

What this manual will leave you able to do

- Critically evaluate the limitations of Mental Health First Aid training in light of recent Cochrane reviews.
- Apply NICE NG212 guidelines to formulate a proportionate, stepped response to employee distress.
- Distinguish between the roles of a line manager and a mental health professional in the workplace.
- Identify organisational factors that contribute to employee wellbeing and resilience.
- Articulate the boundaries of manager training in preventing mental ill health.

PART 1

Why an evidence-led alternative to generic awareness training

The honest limits of Mental Health First Aid's evidence, and NICE NG212 as a firmer anchor.

CHAPTER 1

The limits of Mental Health First Aid's evidence base

Mental Health First Aid (MHFA) has become ubiquitous in workplace wellbeing initiatives across the United Kingdom and beyond, often positioned as the gold standard for mental health literacy training. Its distinctive logo and curriculum are recognisable to many employees, and its training programmes have been delivered to millions of people worldwide since their inception in Australia in 2001. However, this lesson examines a critical question that is frequently overlooked in the enthusiasm for implementing MHFA: does its widespread adoption correspond to strong evidence of effectiveness in improving mental health outcomes for recipients? The answer, as revealed by rigorous systematic reviews, suggests a more complex picture than its popularity might indicate.

MHFA was developed by Betty Kitchener and Anthony Jorm as a public health intervention designed to improve mental health literacy and early intervention for mental health problems. The standard 12-hour course teaches participants to recognise the signs and symptoms of common mental health problems, provide initial help, and guide a person towards

appropriate professional help. Its workplace variant focuses specifically on mental health in professional settings, aiming to create mentally healthy workplaces by equipping employees with the skills to support colleagues experiencing mental health difficulties. The programme's intuitive appeal lies in its apparent simplicity and its promise of creating a network of 'first aiders' who can provide immediate support to those in distress.

Despite its intuitive appeal and widespread adoption, the evidence base for MHFA's effectiveness has come under increasing scrutiny. The most comprehensive evaluation to date is the 2023 Cochrane review by Richardson and colleagues, which synthesised findings from twenty-one randomised controlled trials involving over twenty-two thousand participants. This review represents the gold standard for evidence synthesis in healthcare, employing rigorous methodology to minimise bias and maximise the reliability of its conclusions. The reviewers examined MHFA's effects on recipients' mental health and well-being, mental health first aid knowledge, stigmatising attitudes, and helping behaviours. Their findings were striking: for the most important outcomes—mental health and well-being of recipients—the evidence was deemed to be of very low certainty, meaning the reviewers had 'very little confidence in the effect estimate'.

The 'very low-certainty' rating assigned by the Cochrane review is not a minor qualification but represents a fundamental limitation in the evidence base. In the GRADE system (Grading of Recommendations Assessment, Development and Evaluation), which was used to assess the evidence quality, 'very low certainty' indicates that further research is very likely to change the estimate of effect and that the true effect may be substantially different from the estimated effect. This rating reflects multiple concerns across the included studies, including risk of bias in study design, inconsistency of results, imprecision of estimates, and potential publication bias. The reviewers were unable to draw firm conclusions about whether MHFA improves mental health outcomes for recipients, stating that 'the evidence is very uncertain about the effects of MHFA training on mental health first aid knowledge, stigmatising attitudes, or helping behaviours'.

A crucial distinction emerges when examining the evidence more closely: while the Cochrane review found very low certainty evidence for recipient outcomes, some studies did show improvements in trainees' knowledge and

confidence. This represents a fundamental disconnect between what participants report feeling after training and what actually changes for the people they are intended to support. Several studies have demonstrated that MHFA training can increase participants' mental health literacy, improve their recognition of mental health problems, and boost their confidence in providing help. These self-reported improvements are valuable in themselves, as they may encourage more people to engage with mental health issues rather than avoiding them. However, the critical question for workplace wellbeing interventions is whether these improvements translate into better mental health outcomes for those who receive support from MHFA-trained colleagues.

The distinction between trainee outcomes and recipient outcomes is particularly important in workplace contexts where organisations invest significant resources in training programmes expecting measurable improvements in employee wellbeing, productivity, and retention. MHFA training typically requires employees to dedicate substantial time—often two full days—to complete the course, plus ongoing time commitments for maintaining their skills. When organisations implement MHFA based on the assumption that it will improve workplace mental health outcomes, they may be making an investment without robust evidence that it achieves this primary goal. This is not merely an academic concern but has practical implications for how organisations allocate limited wellbeing resources and design comprehensive mental health strategies.

The methodological limitations identified in the MHFA evidence base deserve closer examination. Many of the included trials suffered from high risk of bias, with inadequate randomisation procedures, lack of blinding (which is challenging in behavioural interventions but still important to assess), and incomplete outcome data. The heterogeneity of interventions and outcomes across studies further complicated the analysis, making it difficult to draw consistent conclusions. Additionally, many studies relied on self-report measures rather than objective clinical assessments, introducing potential bias in how outcomes were measured. The follow-up periods in most studies were relatively short, typically ranging from one week to six months, raising questions about the sustainability of any effects that might have been observed.

Another significant limitation is the lack of evidence regarding MHFA's

impact on key workplace outcomes such as absenteeism, presenteeism, productivity, and staff turnover. While these outcomes are not the primary focus of MHFA training, they are often the metrics that organisations care about most when investing in wellbeing initiatives. The Cochrane review found no evidence addressing these important organisational outcomes, leaving a critical gap in our understanding of MHFA's real-world impact in workplace settings. Without this evidence, organisations cannot make informed decisions about whether MHFA represents a good return on investment compared to other wellbeing interventions that might have stronger evidence of improving these outcomes.

The popularity of MHFA despite these evidence limitations raises important questions about how evidence-based practice is implemented in workplace wellbeing. Several factors may explain MHFA's widespread adoption despite the modest evidence base. Its branding and marketing have been highly effective, creating a sense of legitimacy and recognition that other evidence-based interventions lack. The programme's simplicity and apparent ease of implementation make it attractive to organisations seeking quick solutions to complex problems. Additionally, the positive experiences of trainees—who often report feeling more knowledgeable and confident—create a powerful testimonial effect that can overshadow the lack of evidence for recipient outcomes.

An alternative approach to workplace mental health training that has emerged with a stronger evidence base is the Mental Health First Aid England (MHFA England) programme, which differs from the original Australian programme in several important ways. MHFA England has conducted more rigorous evaluations and has developed specific workplace-focused training that addresses some of the limitations identified in the Cochrane review. However, even this more developed version has not yet produced the robust evidence of effectiveness for recipient outcomes that would be expected for a workplace wellbeing intervention with such widespread adoption.

A more promising evidence-based approach to workplace mental health is the development of comprehensive, multi-component wellbeing strategies that address organisational factors as well as individual skills. Interventions such as those based on the Job Demands-Resources model, which focuses on reducing job demands while increasing resources and support, have

shown stronger evidence of improving mental health outcomes in workplace settings. Similarly, interventions that address psychosocial risks through risk assessment and targeted interventions, such as those recommended in the Management Standards approach developed by the Health and Safety Executive, have demonstrated effectiveness in creating mentally healthy workplaces.

Consider a concrete example of how the limitations of MHFA might manifest in a workplace setting. A medium-sized professional services firm implements MHFA training as part of its wellbeing strategy, investing in training twenty employees as mental health first aiders. After the training, surveys show that these first aiders report increased confidence in recognising mental health problems and providing support. However, when examining organisational data over the following year, there is no significant reduction in self-reported mental health problems among employees, no decrease in absenteeism related to mental health, and no improvement in staff retention rates. Meanwhile, the organisation has invested significant resources in the training without achieving the expected improvements in key wellbeing metrics. This scenario illustrates the disconnect between trainee confidence and actual outcomes that the Cochrane review identified, highlighting the importance of looking beyond self-reported improvements to evaluate the real impact of wellbeing interventions.

In contrast, an evidence-based alternative might involve a comprehensive workplace mental health strategy that includes risk assessment, leadership training, flexible working arrangements, and access to evidence-based psychological therapies. Such an approach might show more substantial improvements in mental health outcomes because it addresses multiple factors that contribute to workplace mental health rather than focusing solely on individual skills. For example, a randomised controlled trial of such a multi-component intervention in the UK public sector found significant reductions in common mental disorder symptoms and improved work-related outcomes compared to a control group (Mental Health Foundation, 2020). This more holistic approach aligns with the growing recognition that effective workplace mental health requires addressing organisational factors as well as individual capabilities.

The implications of these evidence limitations for workplace wellbeing practice are significant. Organisations should approach MHFA as one

potential component of a comprehensive mental health strategy rather than as a standalone solution. When implementing MHFA, it should be accompanied by robust evaluation mechanisms that track both trainee outcomes and recipient outcomes, as well as key organisational metrics. Organisations should also consider investing in interventions with stronger evidence of effectiveness for improving mental health outcomes, particularly those that address organisational factors and provide systemic support for employee wellbeing.

As we move towards a more evidence-based approach to workplace wellbeing, it is essential to maintain a balanced perspective. MHFA is not without value; it increases mental health literacy, reduces stigma, and provides a framework for supporting people with mental health problems. These are important outcomes in themselves. However, the evidence limitations identified in the Cochrane review suggest that MHFA should be implemented with realistic expectations about its impact on recipient mental health outcomes. Organisations that adopt MHFA as part of a broader, evidence-based wellbeing strategy are more likely to achieve meaningful improvements in employee mental health and organisational performance than those that rely solely on MHFA training.

In conclusion, while Mental Health First Aid has achieved remarkable popularity as a workplace mental health intervention, its evidence base does not support the strength of outcomes that its widespread adoption might suggest. The 2023 Cochrane review found very low certainty evidence for MHFA's effects on recipient mental health outcomes, despite some evidence of improvements in trainee knowledge and confidence. This distinction between trainee and recipient outcomes is crucial for organisations seeking to make informed decisions about wellbeing investments. As we advance workplace mental health practice, it is essential to maintain a commitment to evidence-based approaches while recognising the value of interventions like MHFA within a comprehensive wellbeing strategy. By doing so, organisations can create mentally healthy workplaces that are both supportive and effective, ultimately improving outcomes for both employees and the organisations that employ them.

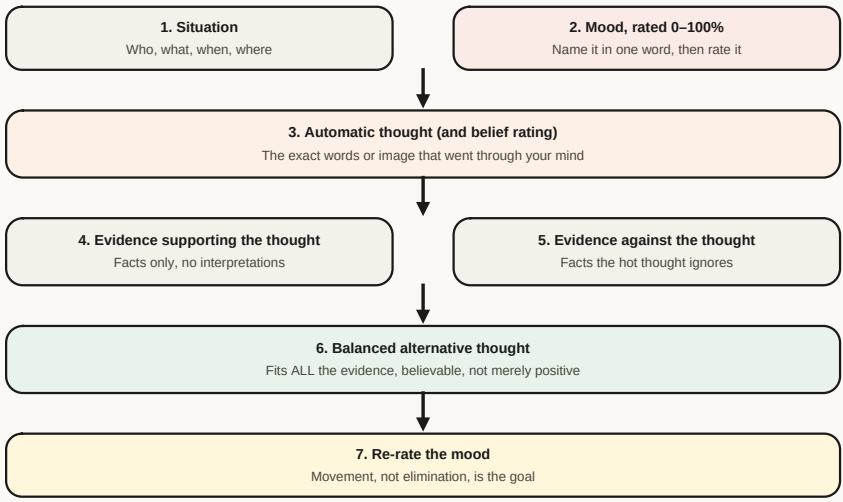


FIGURE 1.1 Working through a thought record

The classic seven-column thought record as a sequence: capture the situation and mood, catch the automatic thought, weigh the evidence on both sides, build a balanced alternative, and re-rate the mood. The learning is in the discipline of evidence before reassurance.

Points worth keeping

- MHFA's ubiquity does not equal strong outcome evidence
- The Cochrane review: 21 trials, 22,000+ participants, very low-certainty findings (Richardson et al., 2023)
- Trainee confidence and knowledge versus recipient outcomes: a crucial distinction
- Stating the evidence limitation plainly rather than smoothing over it

PRACTICE TASK

Write two sentences you could say to a prospective organisational client that state the MHFA evidence position accurately and fairly, without either dismissing or overselling it.

CHAPTER 2

NICE NG212 as a firmer evidence anchor

When organisations decide to invest in workplace wellbeing, the question of what actually works is rarely answered with the rigour it deserves. Generic awareness training has dominated the market for the better part of two decades, and its appeal is understandable: it is easy to commission, easy to deliver to large groups in a single session, and easy to report upward as evidence of organisational commitment. The difficulty is that the evidence base for such approaches is, at best, modest. Awareness alone does not reliably change behaviour, reduce sickness absence, or improve the psychological safety of teams. It is against this backdrop that the National Institute for Health and Care Excellence published its guideline on mental wellbeing at work, catalogued as NG212, in 2022. For practitioners and organisations serious about building a credible, defensible pathway, NG212 represents a firmer evidence anchor than anything the generic awareness training market has previously offered, and understanding what it actually recommends is therefore the proper starting point for this course.

NG212 was developed through the standard NICE process, which means it draws on systematic reviews of the available research literature, weighs the quality of that evidence using established grading frameworks, and produces recommendations that are explicitly tied to the strength of the evidence behind them. This matters enormously in a field where confident-sounding claims are abundant and rigorous evaluation is rare. When a recommendation appears in NG212, it is not there because a training provider found it commercially convenient, nor because a well-meaning occupational health team believed it to be sensible. It is there because a committee of independent experts, including clinicians, researchers, people with lived experience of mental health difficulties at work, and employer representatives, examined the evidence and judged that the recommendation was warranted. Practitioners who ground their workplace wellbeing work in NG212 are therefore standing on considerably more solid ground than those who rely on the accumulated folklore of the training industry.

One of the most important structural features of NG212 is its insistence on

organisational-level assessment and action as a precondition for effective individual-level intervention. This is a point that deserves to be stated plainly, because it runs counter to the dominant model in which organisations skip straight to training their staff. The guideline is clear that employers should first assess the factors within the working environment that are contributing to poor mental wellbeing. These factors include workload and work demands, the degree of control employees have over how they do their work, the quality of relationships between colleagues and between staff and managers, clarity of role, and the extent to which change within the organisation is managed in a way that feels fair and transparent. These are not peripheral concerns. They are the structural determinants of psychological health at work, and no amount of individual-level training will compensate for an environment in which they are chronically misaligned.

The practical implication of this emphasis on organisational-level assessment is that a responsible workplace wellbeing programme begins with a diagnostic phase rather than a training phase. Organisations need to gather data about their own working conditions, using tools such as staff surveys, sickness absence data, and structured conversations with employee representatives, before they can make sensible decisions about where to direct their resources. NG212 does not prescribe a single assessment instrument, which is appropriate given the diversity of workplaces, but it does establish the principle that action should follow assessment rather than precede it. For the practitioner advising an organisation, this means that one of the most valuable contributions you can make at the outset is to slow the process down, to resist the pressure to commission training immediately, and to help the organisation understand what its own data is telling it about the specific pressures its workforce is facing.

Once that organisational-level picture has been established, NG212 turns its attention to the role of line managers, and here the guideline makes recommendations that are both specific and practically significant. The evidence reviewed by the NICE committee consistently pointed to the centrality of the manager-employee relationship in determining whether mental health difficulties at work are identified early and responded to well. Managers are not therapists, and NG212 does not ask them to become therapists. What the guideline does ask is that managers develop the skills necessary for supportive conversations and early identification of difficulty.

These are distinct and learnable competencies, and they are quite different from the broad mental health literacy content that generic awareness training typically delivers.

Supportive conversations in this context means something specific. It means the capacity to notice when a member of staff appears to be struggling, to create a private and unhurried opportunity to check in with that person, to listen without immediately problem-solving or minimising, to ask open questions that invite the employee to describe their own experience, and to respond in a way that communicates genuine concern without crossing into clinical territory. This is not a trivial skill set. Many managers find these conversations deeply uncomfortable, partly because of their own uncertainties about mental health, partly because of a well-founded anxiety about saying the wrong thing, and partly because the cultures of many organisations have historically rewarded task-focus over relational attentiveness. Training that is genuinely focused on supportive conversations needs to address all of these dimensions, and it needs to give managers the opportunity to practise the relevant skills in realistic scenarios rather than simply absorbing information about mental health conditions.

Early identification of difficulty is the complementary competency. The evidence base reviewed in NG212 supports the view that outcomes are significantly better when mental health difficulties are identified and responded to before they escalate to the point of sickness absence or crisis. Managers are often the people best placed to notice early signs of difficulty, because they see their team members regularly and have a baseline sense of how each person normally presents. What they frequently lack is a framework for interpreting what they observe and a clear sense of what to do with that observation. Training focused on early identification gives managers that framework. It helps them distinguish between the ordinary fluctuations in mood and performance that are part of any working life and the more sustained changes in behaviour, engagement, or presentation that may indicate that someone is struggling in a way that warrants a conversation. It also equips them with a sense of the organisational resources available, so that when a conversation does take place, they can signpost appropriately rather than leaving the employee with nowhere to go.

The third major structural feature of NG212 that this course will return to repeatedly is its advocacy for a proportionate, stepped response rather than

a single blanket training session treated as sufficient in itself. The stepped care logic that underpins so much of the NICE evidence base in mental health is applied here to the organisational context in a way that is both intellectually coherent and practically useful. The idea is that different employees will have different levels of need, and that the organisation's response should be calibrated accordingly. Some employees will benefit from access to good-quality information about mental health and self-care strategies. Others will need a supportive conversation with their manager and perhaps a referral to an employee assistance programme. Others still will need more structured support, whether through occupational health, a return-to-work plan, or onward referral to clinical services. A single blanket session delivered to all staff simultaneously cannot possibly address this range of need. It is, by definition, a one-size-fits-all response to a problem that is inherently varied in its presentation and severity.

The proportionate, stepped response model also has implications for how training itself is structured. Rather than a single event, NG212 points towards a layered approach in which different groups within the organisation receive different kinds of support. Senior leaders need to understand the business case for mental wellbeing investment and their own responsibilities under the guideline. Line managers need the specific skills training around supportive conversations and early identification already described. All staff benefit from access to clear information about available support and from working in an environment where psychological safety is actively cultivated. And those with particular responsibilities, such as mental health first aiders or occupational health professionals, need more specialist input that goes beyond what a general awareness session can provide. Designing a programme that reflects this layered logic requires more thought and more organisational engagement than commissioning a single training day, but the evidence reviewed by NICE suggests strongly that it is this kind of structured, proportionate approach that actually moves the dial on workforce mental health outcomes.

To make this concrete, consider the example of a medium-sized logistics company that approached a workplace wellbeing consultant following a noticeable increase in stress-related sickness absence over the preceding eighteen months. The consultant's first step, consistent with the NG212 framework, was not to commission training but to conduct an organisational-level assessment. This involved analysing sickness absence

data by team and by role, running a brief staff survey focused on the key working conditions identified in the guideline, and holding structured conversations with a sample of line managers and employee representatives. The assessment revealed that the increase in absence was concentrated in two specific operational teams, both of which had experienced significant management changes and an increase in workload following a contract expansion. The broader workforce was not experiencing the same level of difficulty. Armed with this information, the consultant was able to recommend a proportionate response: targeted work with the two affected teams, including specific training for their managers in supportive conversations and early identification, alongside a review of workload allocation and a clearer communication process for future organisational changes. The rest of the organisation received a lighter-touch offer, including access to updated information about the employee assistance programme and a brief session for all managers on the principles of psychologically safe team cultures. This is what a NICE-informed, evidence-led approach looks like in practice: specific, proportionate, grounded in organisational data, and structured around the competencies that the evidence identifies as genuinely effective.

For practitioners working in this field, NG212 offers something that has been conspicuously absent from workplace wellbeing for too long: a credible, independently developed evidence anchor that can be used to challenge the dominance of generic awareness training and to make the case for approaches that are more carefully designed, more specifically targeted, and more likely to produce meaningful outcomes. The guideline does not provide a complete blueprint for every workplace, and it does not remove the need for skilled professional judgement in applying its recommendations to specific organisational contexts. What it does provide is a principled framework, grounded in the best available evidence, that distinguishes between what is likely to help and what is likely merely to give the appearance of helping. In a field where the appearance of action has too often substituted for action itself, that distinction is not a small thing. It is the foundation on which genuinely effective workplace wellbeing practice must be built.

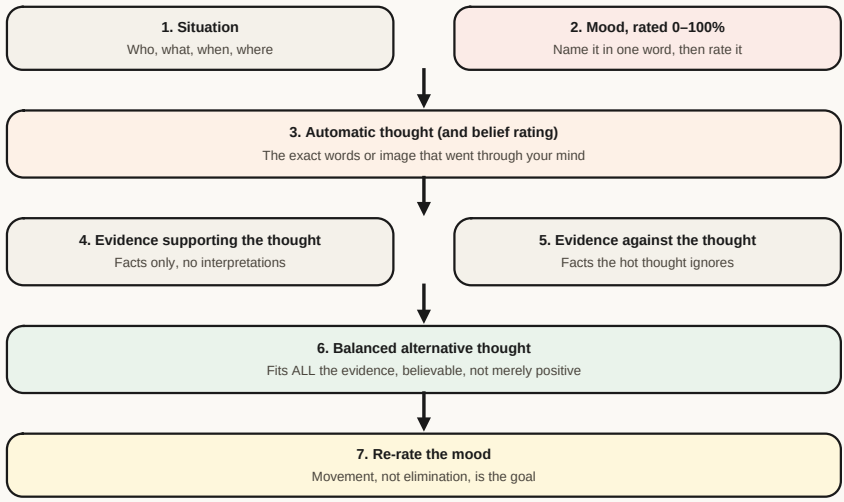


FIGURE 2.1 Working through a thought record

The classic seven-column thought record as a sequence: capture the situation and mood, catch the automatic thought, weigh the evidence on both sides, build a balanced alternative, and re-rate the mood. The learning is in the discipline of evidence before reassurance.

Points worth keeping

- NICE NG212: organisational-level assessment and action
- Manager training focused on supportive conversations and early identification
- A proportionate, stepped response rather than one blanket session
- NG212 as the firmer evidence anchor for this pathway

PRACTICE TASK

List the three NICE NG212 elements you consider most relevant to an organisation with no existing wellbeing provision at all.

CHAPTER 3

Reading Evidence Claims in Workplace Wellbeing: A Beginner's Guide to Knowing What Counts

One of the most practically useful things a practitioner entering the workplace wellbeing field can learn is how to read an evidence claim critically. This is not an abstract academic skill. It is something you will need almost immediately, because the marketplace for workplace mental health products is crowded, commercially motivated, and frequently characterised by confident assertions that do not survive scrutiny. Vendors will tell you their programme is 'evidence-based'. Commissioners will ask whether your offer has 'a strong evidence base'. Colleagues will cite statistics from white papers produced by organisations with a financial interest in the conclusions. Without a working understanding of what evidence actually means in this context, a practitioner is poorly equipped to make honest recommendations, to push back on poor-quality products, or to explain to a client organisation why one approach is preferable to another. This lesson provides that working understanding at a level appropriate to someone beginning their professional engagement with the field.

The phrase 'evidence-based' has become so widely used that it has lost much of its precision. In healthcare and psychology, the term has a reasonably well-established meaning rooted in the hierarchy of evidence developed through the evidence-based medicine movement of the 1990s, associated most prominently with the work of David Sackett and colleagues at McMaster University. Sackett's formulation — that clinical decisions should integrate the best available research evidence with clinical expertise and patient values — was never intended to mean that only randomised controlled trials matter, but the hierarchy he and others articulated does place different study designs at different levels of confidence. At the apex sit systematic reviews and meta-analyses of well-conducted randomised controlled trials, because these aggregate findings across multiple studies and control for the biases that afflict individual experiments. Below these sit individual randomised controlled trials, then controlled observational studies,

then uncontrolled observational studies, then expert opinion and case reports. This hierarchy exists because each step down introduces more opportunity for the results to reflect something other than the true effect of the intervention.

Why does this matter for workplace wellbeing? Because a great deal of what is sold and delivered in this space rests on evidence from the lower rungs of the hierarchy, and practitioners need to be able to recognise this without dismissing the evidence entirely. A well-conducted cohort study showing that employees who received a particular intervention reported lower levels of psychological distress at three months is genuinely informative, even if it cannot establish causation with the confidence of a randomised trial. The question is not whether such evidence is worthless — it is not — but whether the claims made on its basis are proportionate to what the evidence can actually support. When a provider asserts that their training 'reduces absenteeism by thirty per cent', the practitioner should immediately ask: compared to what, in which population, over what time period, measured how, and by whom? These are not pedantic questions. They are the minimum required to assess whether the claim is meaningful.

A particularly important concept for practitioners in this field is the distinction between efficacy and effectiveness. Efficacy refers to whether an intervention works under controlled, often ideal, conditions — typically in a research trial with carefully selected participants, trained and supervised deliverers, and close monitoring. Effectiveness refers to whether it works in real-world conditions, with the full range of participants who present in practice, delivered by practitioners with varying levels of skill and support, in organisations with competing priorities and imperfect implementation. Many interventions demonstrate reasonable efficacy in trials but show substantially weaker effects when rolled out at scale. This gap is well documented in psychological intervention research generally and is particularly relevant in workplace wellbeing, where the conditions of delivery vary enormously between a small professional services firm and a large public sector organisation. When evaluating a product or approach, the practitioner should ask not only whether it has been tested, but whether it has been tested in conditions resembling those in which it will actually be used.

A related concept is that of the active ingredient. Psychological and

educational interventions are complex: they involve multiple components, and it is rarely obvious which components are responsible for any observed benefit. A training programme might include psychoeducation about mental health, skills practice, peer discussion, and follow-up resources. If participants improve, it is not clear whether the improvement stems from the psychoeducation, the skills practice, the social connection generated by group discussion, or simply the experience of their employer investing time and attention in their wellbeing — what researchers call a non-specific or common factor effect. This matters because if the active ingredient is the non-specific effect of feeling valued, then the specific content of the training is less important than the quality of its delivery and the organisational context in which it sits. NICE guideline NG212, which this module has already introduced, reflects this complexity by recommending a whole-organisation approach rather than isolated training events, precisely because the evidence suggests that context and culture are themselves active ingredients in workplace mental health.

Practitioners should also be alert to the difference between statistical significance and clinical or practical significance. A study may find that an intervention produces a statistically significant reduction in a measure of psychological distress, meaning that the result is unlikely to have occurred by chance. But the size of that reduction — the effect size — may be small enough to be of limited practical importance. Effect sizes are conventionally described using Cohen's d, where 0.2 is considered small, 0.5 medium, and 0.8 large. Many workplace wellbeing interventions produce effect sizes in the small range. This does not mean they are useless, particularly if they are low-cost and scalable, but it does mean that claims of transformative impact should be treated with scepticism. A practitioner who understands effect sizes can have a more honest conversation with a commissioner about what a programme is realistically likely to achieve.

To make this concrete, consider the following worked example. A practitioner is approached by a commissioner at a medium-sized logistics company who has seen an advertisement for a digital mental health platform. The platform's website states that it is 'clinically proven to reduce anxiety and depression in the workplace' and cites a published study showing a statistically significant improvement in scores on the Generalised Anxiety Disorder-7 scale among employees who used the platform for eight weeks. The practitioner examines the study and notes several things. First, it

was conducted by researchers employed by the company that sells the platform, introducing a potential conflict of interest. Second, there was no control group: all participants received the intervention, so there is no way to know whether they would have improved anyway over eight weeks without any intervention. Third, the effect size was d = 0.23, which is small. Fourth, the sample consisted entirely of employees at a single technology company who volunteered to participate, raising questions about whether the findings generalise to a logistics workforce with different demographics, working patterns, and stressors. The practitioner is now in a position to explain to the commissioner, in plain language, that the evidence for this platform is preliminary and limited, that the claim of clinical proof overstates what the study can support, and that before committing a significant budget to the platform it would be worth seeking independent evidence or piloting the intervention with a proper evaluation design. This is not obstructionism. It is the honest professional service that a commissioner deserves and that distinguishes a practitioner grounded in evidence literacy from one who simply defers to marketing.

It is worth noting that the absence of strong evidence is not the same as evidence of absence of effect. Some interventions have not been rigorously tested simply because the research has not yet been done, or because the populations or settings in which they are used are difficult to study. A practitioner should not refuse to recommend anything that lacks a randomised controlled trial, because that standard would rule out a great deal of legitimate clinical and organisational practice. The appropriate stance is one of calibrated uncertainty: being willing to act on the best available evidence while being transparent about its limitations, and being open to updating one's position as better evidence emerges. This is precisely the stance that NICE guidelines attempt to embody, which is one reason why they provide a more reliable anchor than the promotional literature of individual providers.

Finally, practitioners should be aware of publication bias, which is the tendency for studies with positive findings to be published more readily than those with null or negative findings. This means that the published literature on any given intervention is likely to overrepresent positive results. Systematic reviews attempt to address this by searching for unpublished studies and using statistical techniques to detect and correct for publication bias, which is one reason why systematic reviews sit at the top of the

evidence hierarchy. When a provider cites a body of research in support of their product, a useful question is whether that body of research has been subject to systematic review, and if so, what the review concluded. If the answer is that no systematic review exists, or that the review found weaker effects than the individual studies suggested, this is important information.

Developing evidence literacy is a gradual process, and no practitioner beginning their engagement with workplace wellbeing is expected to be a research methodologist. What is expected, at this foundational level, is an awareness that evidence claims require scrutiny, a familiarity with the basic concepts that enable that scrutiny — hierarchy of evidence, efficacy versus effectiveness, effect size, publication bias, conflict of interest — and a professional disposition that prioritises honest appraisal over the comfort of accepting reassuring claims at face value. The clients and organisations that practitioners serve are entitled to that honesty, and the credibility of evidence-led practice in this field depends upon it.

Returning to the worked example of the digital mental health platform is instructive here, because the practitioner's conversation with the commissioner does not end with the initial appraisal. Suppose the commissioner responds positively to the critique and asks what a proper evaluation design might look like if the organisation wished to pilot the platform before committing to a full contract. This is precisely the moment at which evidence literacy becomes generative rather than merely critical, and it is worth following the practitioner's reasoning through that next step.

The practitioner might begin by explaining that the most important single improvement on the original study would be the inclusion of a comparison condition. In an organisational setting, a fully randomised controlled trial is rarely feasible: it is difficult to randomise individuals within the same workplace without contamination, and managers are understandably reluctant to withhold a potentially beneficial resource from one group of employees while offering it to another. A more practical design in this context would be a waitlist-controlled pilot, in which one department or site receives access to the platform immediately whilst another of comparable size and composition receives it after a defined period, typically three to six months. This does not eliminate all confounds, but it does provide a contemporaneous comparison group and allows the practitioner to say something meaningful about whether change in the intervention group

exceeds what might be expected from the passage of time alone. The practitioner should also advise that outcome measures are selected before the pilot begins, that they are collected at baseline and at the end of the active period, and that the measures chosen have established psychometric properties in working-age populations rather than being bespoke scales developed by the platform provider.

A common pitfall at exactly this stage is what might be called the implementation optimism trap. Having rightly pushed back on weak evidence, practitioners sometimes overcorrect by designing an evaluation that is so methodologically ambitious that the organisation cannot sustain it, and the evaluation collapses partway through, leaving no usable data at all. A pilot that collects two time points of data using a validated measure and a waitlist comparison, even imperfectly, is considerably more informative than a gold-standard protocol that is abandoned after six weeks because the HR team ran out of capacity to administer it. The practitioner's role is to help the commissioner identify the most rigorous design that the organisation can actually execute, not the most rigorous design in the abstract. This requires a frank conversation about the resources available for data collection, the degree of senior sponsorship for the evaluation, and the realistic likelihood that employees will complete follow-up measures. Calibrating ambition to context is itself an evidence-informed skill.

This principle carries directly into supervision and case review. When practitioners bring workplace wellbeing work to supervision, the discussion of evidence should not be confined to the initial selection of an intervention. Supervisors working in this field are well placed to ask, at review, whether the practitioner has been able to observe any discrepancy between what the evidence predicted and what they are seeing in delivery. If an intervention with a modest evidence base appears to be producing meaningful change in a particular organisational context, the supervisor might encourage the practitioner to think carefully about what contextual factors could be accounting for that, and whether those factors are replicable or idiosyncratic. Conversely, if an intervention with stronger evidence appears to be producing little observable benefit, the supervisor might explore whether fidelity to the model is being maintained, whether the organisational conditions are sufficiently supportive, or whether the population being served differs in important ways from those in whom the evidence was generated. These are not rhetorical questions. They are the mechanism by

which practitioners develop the kind of refined, contextually sensitive judgement that no amount of pre-reading can substitute for, and they are the reason that supervision remains indispensable even for practitioners who have already developed a solid grounding in evidence literacy.

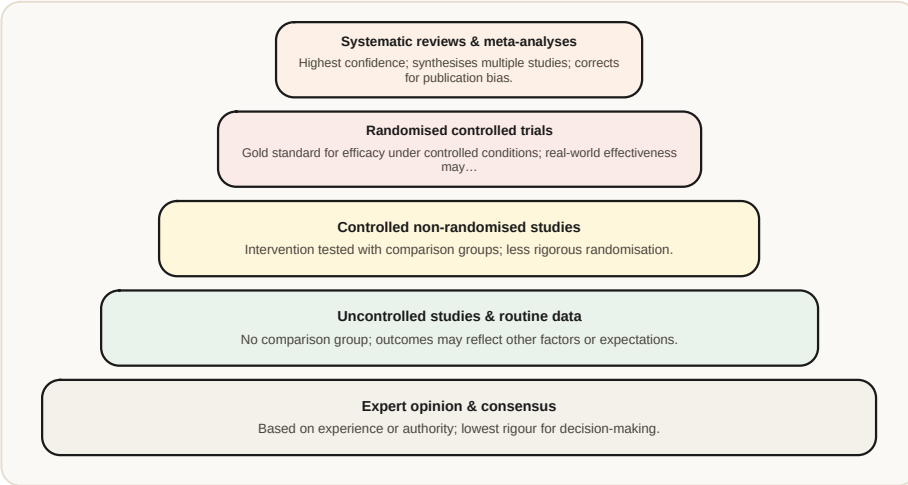


FIGURE 3.1 The evidence hierarchy for workplace wellbeing claims
This figure ranks study designs from strongest to weakest evidence. It helps practitioners evaluate the quality and applicability of wellbeing claims proportionately.

Points worth keeping

- Evidence claims in workplace wellbeing vary enormously in quality, and practitioners need a working understanding of the evidence hierarchy — from systematic reviews and randomised controlled trials down to uncontrolled studies and expert opinion — to evaluate them proportionately.
- The distinction between efficacy (what works under trial conditions) and effectiveness (what works in real-world delivery) is critical: many workplace wellbeing interventions show weaker effects when scaled beyond the controlled conditions in which they were originally tested.
- Statistical significance and practical significance are not the same thing; practitioners should attend to effect sizes and be sceptical of claims that overstate the magnitude of benefit a programme is likely to deliver.
- Publication bias means the published literature tends to overrepresent positive findings; systematic reviews, which attempt to correct for this, sit at the top of the evidence hierarchy and provide a more reliable basis for

recommendations than individual studies cited in promotional materials.

PRACTICE TASK

Locate the promotional or marketing materials for one workplace mental health training product or digital platform currently available in the UK market. Identify every evidence claim made — for example, statements about reducing absenteeism, improving wellbeing scores, or being 'clinically proven'. For each claim, note: what type of study is cited or implied; whether a control group was used; whether the effect size is reported; and whether there is any indication of who funded the research. Write a brief paragraph (150–200 words) summarising what the evidence actually supports and what it does not, using the concepts introduced in this lesson. Bring this paragraph to your next supervision or peer consultation session for discussion.

CHAPTER 4

The Job Demands-Resources model and the HSE Management Standards: the theoretical architecture behind NG212

The previous lessons in this module established that NICE NG212 represents a firmer evidence anchor than generic awareness training, and that its insistence on organisational-level assessment is not an arbitrary preference but a structural feature grounded in evidence. What has not yet been made explicit is where that structural emphasis actually comes from. NG212 did not invent the idea that working conditions themselves determine psychological health; it synthesised and endorsed a body of occupational psychology research that had already established this long before the guideline was written. Two frameworks in particular sit behind the guideline's reasoning, and a practitioner who understands them will be able to explain NG212's recommendations in terms of their underlying mechanism rather than simply reciting them as guidance to be followed. The first is the Job Demands-Resources model, developed by Evangelia Demerouti, Arnold Bakker and colleagues in the early 2000s. The second is the Management Standards approach developed and published by the UK's Health and Safety Executive. Both predate NG212 by many years, and both are, in a meaningful sense, the theoretical scaffolding on which the guideline's organisational-level recommendations rest.

The Job Demands-Resources model begins from a simple but powerful distinction. Every job, in every sector, can be characterised in terms of two broad categories of feature. Job demands are the aspects of work that require sustained physical, cognitive or emotional effort, and that therefore carry a cost if they are not matched by adequate recovery: workload, time pressure, emotional labour, physically demanding tasks, role conflict, and exposure to difficult or distressing content are all demands. Job resources are the aspects of work that help an employee achieve their goals, reduce the cost of demands, or support personal growth: autonomy over how work is done, support from colleagues and managers, clear feedback, opportunities for development, and a sense of meaning in the work itself are

all resources. The model's central claim is that these two categories operate through largely distinct psychological pathways. High job demands, unmitigated by sufficient recovery or resources, predict exhaustion and, over time, burnout. Rich job resources, independent of the level of demand, predict engagement, motivation and the energised, absorbed quality of work that organisational psychologists call work engagement. Crucially, resources also buffer the effect of demands: a role with high workload but strong manager support and real autonomy over how the work gets done produces a markedly different psychological outcome than the same workload delivered with neither.

This distinction matters enormously for how a practitioner thinks about workplace wellbeing, because it explains why training alone, delivered without attention to the demands and resources actually present in a given role, has such limited traction. A training session cannot alter an employee's workload, their degree of control over their own tasks, or the quality of support available from their line manager; it can only, at best, help the individual cope somewhat better with demands that remain unchanged, or help a manager notice distress that structural conditions are actively producing. The Job Demands-Resources model therefore gives a practitioner a vocabulary for diagnosing why a given workplace is struggling, which is quite different from simply observing that staff report high stress. Two organisations can report identical staff stress scores for entirely different reasons: one because demands are genuinely excessive relative to any realistic resourcing, the other because demands are moderate but resources, particularly managerial support and autonomy, are badly depleted. The appropriate intervention in each case is different, and a training-only response addresses neither.

The second framework, the HSE Management Standards, approaches the same underlying problem from a regulatory and risk-assessment angle rather than an academic-psychology one, and the two frameworks are complementary rather than competing. The Health and Safety Executive, as the UK's workplace health and safety regulator, developed the Management Standards specifically to help employers meet their existing legal duty under the Management of Health and Safety at Work Regulations to assess and manage risks to employee health, including psychological health, arising from work. The Standards identify a set of working conditions, commonly summarised as six areas, that research consistently associates with

work-related stress when poorly managed: the demands of the job, including workload and working patterns; the degree of control employees have over how they do their work; the support available from managers and colleagues; the quality of relationships at work, including the management of unacceptable behaviour; role clarity, including whether people understand their responsibilities and how they fit into the organisation; and how organisational change is managed and communicated. Readers familiar with the Job Demands-Resources model will notice the clear family resemblance: several of the six areas map directly onto job resources as that model defines them, while demands and change management map onto the demands side of the equation. The Standards were never designed as an academic model in the way the Job Demands-Resources framework was; they were designed as a practical risk-assessment tool that employers, including those with no occupational psychology expertise, could use to identify where their organisation's working conditions might be generating preventable harm. That practical, risk-assessment orientation is precisely why NG212 draws on this tradition when it insists that organisational-level assessment must precede training: the Management Standards provide employers with an existing, legally grounded method for conducting exactly that assessment, and NG212 is, in effect, endorsing an established regulatory tool rather than inventing a new requirement from nothing.

Understanding both frameworks together gives a practitioner something practically valuable: a shared vocabulary that connects the language of psychological research to the language of legal compliance that many organisational buyers, particularly those in HR and health and safety functions, already use fluently. When a practitioner can explain that a proposed organisational assessment is, in substance, an application of the HSE's own Management Standards framework to the specific working conditions in a client's organisation, and that doing so addresses risk factors that the Job Demands-Resources research identifies as causally implicated in burnout and disengagement, the practitioner is speaking a language that a health and safety manager, an HR director and an occupational psychologist can all recognise as credible, even though each of those audiences might otherwise engage with the material through a different professional lens.

It is worth being precise about the limits of both frameworks, in keeping with the evidence-literacy habits introduced elsewhere in this module. The Job Demands-Resources model is a general theoretical framework supported by

a substantial body of correlational and some longitudinal research; it is not a diagnostic instrument with fixed cut-off scores, and applying it to a specific organisation requires judgement about which demands and resources are most salient in that particular context rather than a mechanical checklist exercise. The HSE Management Standards similarly do not prescribe a single assessment tool; organisations can and do use varying methods, from staff surveys to structured interviews, to gather data against the six areas, and the quality of the resulting picture depends heavily on how well that data collection is designed and how honestly staff feel able to respond. Neither framework removes the need for a skilled practitioner's judgement in interpreting what the data shows and recommending a proportionate response; what both frameworks provide is a disciplined structure for asking the right questions, rather than a mechanical answer.

A worked example illustrates how the two frameworks combine in practice. A mid-sized architecture practice approaches a workplace wellbeing consultant after several senior staff resign within a short period, citing burnout in their exit interviews. The consultant's instinct, informed by the Job Demands-Resources model, is to ask not simply how stressed people feel but what the specific demands and resources look like for the roles affected. Structured conversations reveal that project deadlines have become significantly tighter following the loss of two senior project managers who were never replaced, meaning remaining staff have absorbed additional workload (a demand) while simultaneously losing the day-to-day mentoring and workload-triaging support those project managers used to provide (a resource). Using the HSE Management Standards as a structured frame for a brief staff survey, the consultant confirms that scores on the demands and support dimensions have moved sharply in the twelve months since the departures, while scores on role clarity and relationships remain comparatively stable. This combined picture, which training alone would never have surfaced, points towards a recommendation that is substantially about resourcing and interim management support rather than psychoeducation: the practice needs either additional project management capacity or a temporary redistribution of workload, alongside targeted support for the remaining team while that resourcing gap is addressed. Training for managers in supportive conversations, which this pathway addresses elsewhere, remains a valuable complementary measure, but it is positioned honestly as one part of a response whose primary driver is a

demands-resources imbalance that no amount of conversational skill can fully offset.

It is worth tracing the Job Demands-Resources model back one step further, because doing so helps explain why it displaced its predecessor as the dominant framework in this area and gives a practitioner a fuller sense of the intellectual lineage they are drawing on when they cite it. The model most directly builds on Robert Karasek's demand-control model, first published in 1979, which proposed that psychological strain at work arises specifically from the combination of high demands and low control, a pairing Karasek termed high-strain work, while high demands combined with high control, which he termed active work, were comparatively protective despite the equally high workload. Karasek's model was influential and empirically productive, generating a large body of research linking high-strain jobs to cardiovascular and psychological morbidity, but it was also limited in scope, treating control as the only resource worth measuring and leaving other plausible protective factors, such as social support, largely outside the formal model. Demerouti and Bakker's Job Demands-Resources framework can be understood as a deliberate generalisation of Karasek's insight: rather than treating control as the sole buffering resource, it opens the category of resources to anything that genuinely helps an employee meet demands or achieve goals, which allows the model to accommodate workplaces where, for instance, strong colleague support or clear feedback matters more than formal autonomy over task sequencing. For a practitioner explaining this lineage to a sophisticated buyer, perhaps one with their own background in occupational psychology or health and safety, being able to locate the Job Demands-Resources model as an extension of Karasek's earlier, narrower demand-control tradition, rather than presenting it as though it emerged from nowhere, signals exactly the kind of grounded familiarity with the field's actual history that distinguishes a credible practitioner from one reciting model names without appreciating how the underlying ideas developed.

For a practitioner building a credible, evidence-led workplace wellbeing offer, these two frameworks are not academic decoration. They are the mechanism that explains why NG212's organisational-level emphasis is correct, they give practical, defensible structure to what an organisational assessment should actually examine, and they equip the practitioner to translate psychological research into language that resonates with the compliance and risk-management concerns that often sit alongside, or even

ahead of, an organisation's humanitarian motivation for investing in employee wellbeing in the first place.

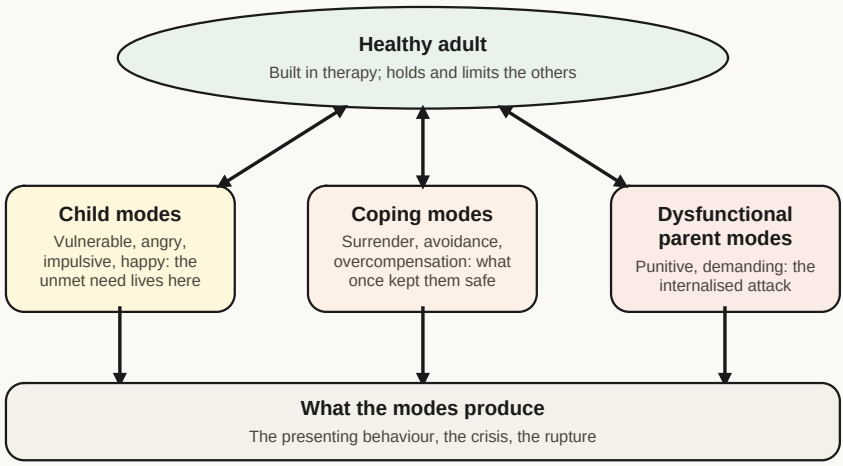


FIGURE 4.1 Schema modes

Schema therapy describes a mode as the emotional state and coping response active at a given moment. Grouping the modes makes the work in the room legible: the child modes carry the unmet need, the coping modes are what once kept the person safe, the internalised parent modes carry the attack, and the healthy adult, built over the course of therapy, is the part that eventually holds and limits the other three. What presents in the consulting room is the product of the three mode classes, which is why the crisis is read as a mode and not as a diagnosis.

Points worth keeping

- The Job Demands-Resources model distinguishes job demands (which predict exhaustion when unmatched by recovery) from job resources (which predict engagement and buffer the cost of demands), giving a mechanistic account of why working conditions, not individual resilience alone, determine psychological health at work.
- The HSE Management Standards translate a closely related set of working-condition factors — demands, control, support, relationships, role, and change — into a practical, legally grounded risk-assessment framework that NG212's organisational-assessment requirement draws on.

- Neither framework is a mechanical diagnostic instrument; both require skilled judgement to apply to a specific organisational context, and neither removes the need for proportionate, context-sensitive recommendations.
- Fluency in both frameworks lets a practitioner speak credibly to HR, health and safety, and occupational psychology audiences simultaneously, connecting research language to the compliance language many organisational buyers already use.

PRACTICE TASK

Take an organisation you know reasonably well, whether a current or former employer or a client. Using the six HSE Management Standards areas as headings, write one or two sentences under each describing what you believe the working conditions actually look like there, based on your own observation. Then identify which of those six areas you suspect would show up as a job resource and which as a job demand if you mapped them onto the Job Demands-Resources model, and note where you think the two pictures agree or diverge.

CHAPTER 5

Mental health as a continuum: flourishing, languishing and the limits of an illness-only model

Much of the generic awareness training this pathway positions itself against operates, often without saying so explicitly, on an implicit binary model of mental health: a person either has a mental health problem or they do not, and training exists to help colleagues recognise which category someone currently falls into. This binary framing is intuitive, and it is not entirely without foundation, since diagnosable mental health conditions are real and clinically meaningful categories. But the binary model is also seriously incomplete as a basis for workplace wellbeing practice, and this lesson introduces a different way of thinking about mental health that has a genuine evidence base and that will recur throughout this pathway: the dual-continuum model, most closely associated with the American sociologist and psychologist Corey Keyes.

Keyes' central proposal, developed across a substantial body of research beginning in the late 1990s and continuing through the 2000s, is that mental health and mental illness are better understood as two related but distinct continua rather than as opposite ends of a single spectrum. On this account, a person can be positioned anywhere along an illness continuum, from no diagnosable condition through to severe and persistent illness, and independently positioned anywhere along a wellbeing continuum, from what Keyes termed flourishing at one end to languishing at the other. Flourishing describes a state of high emotional, psychological and social wellbeing: a person who is flourishing experiences positive emotion with reasonable frequency, feels their life has purpose and direction, experiences a sense of mastery and autonomy, and feels genuinely connected to others and to their community. Languishing, at the other end, describes an absence of these qualities without necessarily involving diagnosable illness: a person who is languishing may report feeling empty, stagnant, or simply going through the motions, without meeting criteria for depression, anxiety or any other recognised condition.

The practical significance of treating these as two separate continua rather than one is considerable, and it has direct implications for workplace practice. Because the continua are related but distinct, a person can be entirely free of diagnosable mental illness and still be languishing: showing up to work, meeting their basic obligations, exhibiting no symptoms that would register on a clinical screening tool, and yet experiencing none of the vitality, engagement or sense of purpose that flourishing would involve. Keyes' research, replicated across multiple large population samples, consistently found that a substantial proportion of any given population sits in precisely this quadrant: no diagnosable illness, but also no flourishing, a state Keyes described as moderate mental health or, more starkly, as languishing without illness. Equally, and this is the finding that most challenges the binary model, a person can have a diagnosed and actively managed mental illness while still experiencing meaningful flourishing on the wellbeing continuum: someone living with a chronic depressive disorder that is well managed through therapy and medication can nonetheless report strong social connection, a sense of purpose, and reasonable day-to-day positive emotion, occupying a position that the illness-only model has no vocabulary to describe.

For workplace wellbeing practice, this reframing matters because of what it implies about where a sensible intervention should focus its attention. A training programme built entirely around the illness-only model will naturally focus on recognising symptoms and signposting towards treatment: useful, but relevant only to the minority of any workforce who are experiencing diagnosable illness at a given moment. A workplace that uses the dual-continuum model as part of its conceptual foundation can additionally ask a different, and arguably more immediately actionable, question: regardless of illness status, what proportion of our workforce is languishing rather than flourishing, and what features of working life are contributing to that languishing? This question connects naturally to the Job Demands-Resources model introduced in the previous lesson, because the resources that model identifies, autonomy, support, meaning, feedback, map closely onto the specific components Keyes includes within flourishing. A workplace rich in job resources is, almost by definition, a workplace more likely to produce flourishing employees, independent of whether any of those employees also happen to meet criteria for a diagnosable condition.

This distinction also has a direct bearing on the stigma and disclosure

problem that this pathway addresses elsewhere. If workplace wellbeing is framed exclusively in illness terms, then any conversation about an employee's wellbeing risks being heard, by that employee, as a conversation about whether they are mentally ill, with all the stigma, fear of career consequence and self-concealment that framing can provoke. A dual-continuum framing allows a manager to ask about flourishing and languishing in terms that do not presuppose illness at all: a conversation about whether someone feels they have a sense of purpose in their current role, whether they feel genuinely supported, or whether they have been finding work meaningful lately, can be had with any employee, at any time, without implying that the manager suspects the presence of a diagnosable condition. This is not a trivial rhetorical trick. It changes what the conversation is about, and in doing so it removes one of the principal barriers, namely fear of being labelled unwell, that prevents employees from engaging honestly with wellbeing conversations in the first place.

It is worth being precise about what the dual-continuum model does and does not establish, in keeping with the evidence-literacy habits this pathway is built around. Keyes' framework is well supported by factor-analytic research showing that illness symptoms and wellbeing indicators do indeed load onto statistically separable dimensions rather than a single bipolar scale, and the finding has been replicated across different populations and measurement approaches. It is less a predictive model, in the sense that the Job Demands-Resources model attempts to predict burnout and engagement from specific working conditions, and more a descriptive and conceptual reframing: it tells a practitioner that illness and wellbeing need to be assessed and discussed separately, without on its own prescribing which specific organisational interventions will shift someone from languishing towards flourishing. For that practical question, the Job Demands-Resources model and the HSE Management Standards, covered in the preceding lesson, remain the more directly actionable tools. The dual-continuum model's contribution is conceptual rather than operational, but the concept matters a great deal, because it determines what a workplace wellbeing programme is actually trying to achieve and, critically, how it talks about that aim to a workforce that may be wary of anything that sounds like a mental illness screening exercise.

Consider a worked example that brings these threads together. A regional retail chain commissions a foundational wellbeing briefing for its store

managers, initially framed by the client as “mental health awareness.” The practitioner delivering the briefing, informed by the dual-continuum model, reframes the opening content rather than simply delivering a symptom-recognition session. Rather than beginning with a list of mental health conditions and their warning signs, the session opens by asking managers to consider their own teams through the lens of flourishing and languishing: who on their team seems energised, connected and purposeful at the moment, and who seems to be going through the motions, present but disengaged, without necessarily showing any sign of diagnosable distress. This framing produces a noticeably different quality of discussion than a symptom-focused opening typically does. Several managers immediately identify long-serving staff who are not struggling in any clinical sense but who have become visibly disengaged since a recent restructuring removed much of their autonomy over scheduling, a clear job-resource loss in Job Demands-Resources terms. The practitioner is then able to connect this observation back to the organisational and framework content covered elsewhere in this pathway, helping managers see that addressing languishing in this part of the team is a legitimate and valuable wellbeing intervention in its own right, entirely separate from and prior to any question of whether a team member might also be experiencing a diagnosable condition requiring clinical support.

It is important to name a specific way in which the flourishing-languishing framing can be misapplied, because the same reframing that makes it such a useful, non-stigmatising conversational tool also creates a genuine risk if it is handled carelessly. An organisation that adopts the language of flourishing without also internalising the underlying commitment to honest assessment can slide into treating flourishing as a performance standard that employees are implicitly expected to display, rather than as a description that a wellbeing programme is trying to understand and support. This produces what might be called wellbeing-washing: a workplace that talks fluently about flourishing, perhaps even running surveys that ask employees to rate their own sense of purpose and connection, while doing nothing to address the actual job demands and resource deficits identified elsewhere in this module, and treating any employee who reports languishing not as a signal that working conditions need attention but as an individual failure to adopt a sufficiently positive outlook. This is a serious misuse of the construct, and it inverts Keyes’ own purpose in developing it,

which was descriptive and diagnostic rather than aspirational or evaluative. A practitioner introducing the dual-continuum model to a client should be explicit that identifying languishing in a workforce is informationally useful precisely because it points towards specific, addressable organisational conditions, in the same way a low score on a Management Standards survey area points towards a specific working condition to examine, and should actively discourage any organisational tendency to treat flourishing scores as a target for individual employees to improve through their own effort alone, independent of whether the underlying working conditions that produce languishing have been addressed at all. Where a client's enthusiasm for flourishing language outpaces their willingness to examine working conditions honestly, the practitioner's task is to slow that enthusiasm down rather than to supply more flourishing-themed content that risks becoming a substitute for genuine organisational change.

The dual-continuum model also has a specific and underappreciated relevance to neurodivergent employees, including autistic employees and those with ADHD, precisely because neurodivergence itself is not a mental illness and sits awkwardly within any framework built exclusively around the illness continuum. An autistic employee may have no diagnosable mental health condition whatsoever and yet experience sustained languishing in a workplace whose sensory environment, communication norms and expectations around eye contact, small talk or open-plan collaboration are poorly matched to their own needs, producing chronic low-level strain that an illness-focused model has no vocabulary to capture and that a manager looking only for symptoms of depression or anxiety will simply miss. Equally, the same employee working in a role and environment that has been thoughtfully adapted, perhaps with clearer written communication, predictable routines and reduced sensory load, may report genuine flourishing on exactly the dimensions Keyes' model describes, a sense of purpose, mastery and connection, despite having a lifelong neurodevelopmental difference that an illness-only framework might otherwise treat as inherently incompatible with wellbeing. This matters for foundational practice because it illustrates with particular clarity why flourishing and languishing need to be assessed independently of any illness or diagnostic category at all: the relevant question for this employee is not whether they have a condition, which is a fixed and in this case irrelevant fact, but whether their current working environment is one in which a person

with their particular needs can flourish, which is a variable, addressable, and entirely legitimate question for a workplace wellbeing programme to ask regardless of diagnostic status.

The dual-continuum model, properly understood, does not replace the clinical and signposting content that generic awareness training provides, and this pathway does not ask practitioners to abandon the vocabulary of mental illness, which remains clinically necessary and genuinely important. What it provides is an additional, complementary lens that captures a much larger proportion of everyday workplace experience than an illness-only framework can reach, and it offers a way of opening wellbeing conversations that sidesteps much of the stigma and defensiveness that illness-framed conversations often provoke. A foundational, evidence-led workplace wellbeing offer benefits from being able to move fluently between both lenses: recognising and responding appropriately when illness is present, while also recognising that the much larger task of building a genuinely well workforce requires attention to flourishing and languishing across the entire workforce, illness or no illness.

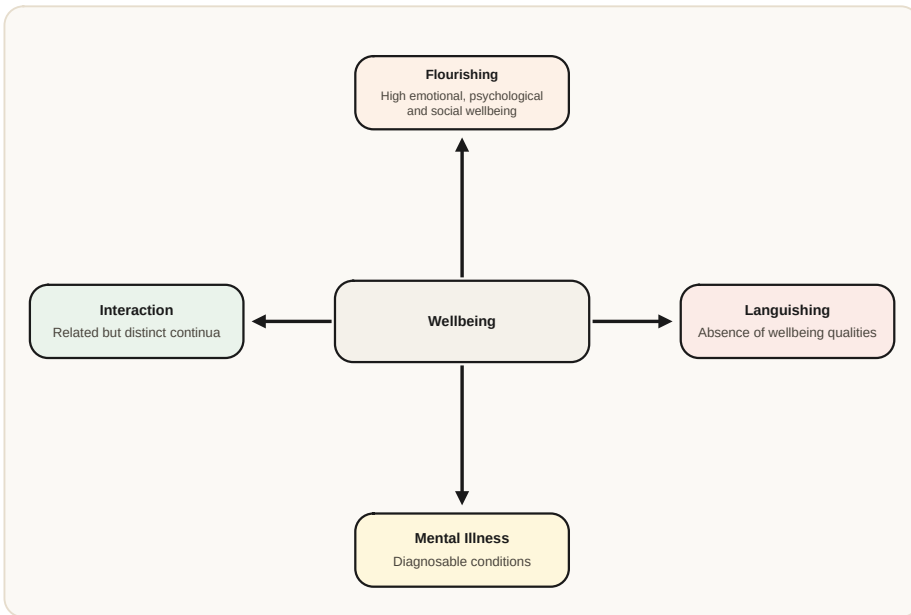


FIGURE 5.1 Dual-Continuum Model

The dual-continuum model treats mental illness and mental health as related but separate continua. It describes flourishing and languishing as states of wellbeing, independent of diagnosable illness.

Points worth keeping

- Corey Keyes' dual-continuum model treats mental illness and mental health/wellbeing as related but statistically separable continua, so a person can be free of diagnosable illness while languishing, or managing a diagnosed condition while flourishing.
- Flourishing describes high emotional, psychological and social wellbeing (positive emotion, purpose, mastery, connection); languishing describes the absence of these qualities without necessarily involving diagnosable illness.
- Framing wellbeing conversations around flourishing and languishing, rather than illness symptoms, lets managers engage any employee without implying suspicion of a diagnosable condition, reducing one major barrier to honest disclosure.
- The model is conceptual and descriptive rather than predictive; for actionable guidance on which specific working conditions to change, it is best used alongside the Job Demands-Resources model and the HSE Management Standards covered earlier in this module.

PRACTICE TASK

Think of three people you currently work with or have worked with recently. For each, place them, as best you can judge from their outward presentation alone, somewhere on an illness continuum (no apparent difficulty through to significant difficulty) and separately somewhere on a flourishing-languishing continuum. Write one sentence per person explaining your reasoning, and note whether any of the three occupy a position — such as apparently well but languishing, or managing a known difficulty while still seeming to flourish — that an illness-only model would have struggled to capture.

CHAPTER 6

The economic and business case: presenteeism, the Thriving at Work review, and what the numbers can and cannot tell you

Practitioners entering the workplace wellbeing field quickly discover that evidence of clinical effectiveness, however carefully established, is often not sufficient on its own to secure organisational investment. Commissioners, finance directors and boards frequently want to understand the business case: what does poor mental health actually cost the organisation, and what return can reasonably be expected from investing in addressing it. This lesson equips practitioners to engage with the economic case for workplace mental health investment honestly, which means understanding both the genuine substance behind the business case and the considerable caution that evidence-literate practice requires when numbers are presented with more confidence than the underlying data can support.

The starting point for the economic case is a distinction that is frequently blurred in promotional material but that matters enormously for honest practice: the distinction between absenteeism and presenteeism. Absenteeism is the more familiar concept, referring to time away from work due to sickness, and it is relatively straightforward to measure because organisations already record it for payroll and HR purposes. Presenteeism is a less visible but, according to a substantial body of occupational health research, often more costly phenomenon: it refers to attending work while unwell, whether physically or psychologically, and functioning at reduced capacity as a result. An employee experiencing presenteeism is recorded as present, generates no absence statistic, and may show no obvious outward sign of difficulty, while nonetheless working at a fraction of their usual productivity, making more errors, struggling to concentrate, or withdrawing from the collaborative and creative aspects of their role. Because presenteeism leaves no administrative trace in the way absence does, it is inherently harder to quantify, and any figure purporting to measure its cost should be treated with proportionately more caution than an absence figure

drawn directly from payroll records.

This is not a reason to dismiss presenteeism as a concept; the underlying phenomenon, that psychological distress reduces functional capacity even when a person remains at their desk, is well supported by research on the cognitive and attentional effects of conditions such as depression and anxiety, which reliably show impairments in concentration, working memory and decision-making speed even in presentations that would not necessarily lead to a period of sickness absence. The caution is specifically about the economic figures that are frequently attached to presenteeism in promotional and policy literature, which typically rely on self-report productivity scales, applied to estimated salary costs, aggregated across an entire workforce, to produce a total cost figure. Each step in that chain introduces uncertainty: self-report productivity measures are approximate at best, the proportion of a workforce assumed to be experiencing presenteeism at any given time varies enormously between sources, and the resulting aggregate figures, while individually plausible, are considerably less precise than the single, confident-sounding number with which they are often presented. A practitioner should be able to use the concept of presenteeism to help a commissioner understand why absence figures alone understate the true cost of poor workplace mental health, without presenting a specific aggregate cost figure as though it carried the same evidential weight as a directly measured absence statistic.

The most significant UK policy document in this area, and one that practitioners in this field should be familiar with regardless of which sector they work in, is the independent review commissioned by the UK government and published in 2017, widely known by the title of its report, Thriving at Work, and associated with its two lead authors, Dennis Stevenson and Paul Farmer. The review was commissioned specifically to examine how employers can better support the mental health of people currently in employment, and how individuals with mental health problems can be helped to enter and stay in work. Its significance for this pathway lies less in any specific statistic it produced, since cited figures from reviews of this kind have a tendency to circulate long after their original context and methodology have been forgotten, and more in the framework it proposed for what responsible employer action actually looks like. The review introduced the idea of a set of core mental health standards that all employers, regardless of size or sector, should be expected to meet,

alongside an additional set of enhanced standards for larger employers with greater capacity to invest. These core standards included producing, implementing and communicating a mental health at work plan; developing mental health awareness among employees; encouraging open conversations about mental health; providing employees with good working conditions; promoting effective people management; and routinely monitoring employee mental health and wellbeing.

Readers who have worked through the earlier lessons in this module will recognise the family resemblance between these core standards and both the Job Demands-Resources model and the HSE Management Standards: good working conditions and effective people management sit squarely within the demands-and-resources framework, while routine monitoring echoes the assessment-before-action principle this pathway has already emphasised repeatedly. The Thriving at Work review did not invent a new theoretical model; it translated existing occupational psychology evidence into a set of practical commitments that a government-commissioned review could credibly ask all UK employers to adopt, and in doing so it gave practitioners in this field a widely recognised reference point that long pre-dates NICE NG212 and that many organisational buyers, particularly those in larger or more policy-literate organisations, will already have some familiarity with.

When a practitioner is building the business case for a prospective client, the Thriving at Work framework offers a genuinely useful structure, because it allows the practitioner to frame recommendations not as novel, unproven ideas but as alignment with an established, government-commissioned standard that the client's own board or senior leadership may already feel some obligation to meet, particularly if the organisation operates in a sector where corporate governance or environmental, social and governance reporting expectations are becoming more attentive to employee wellbeing. This is a legitimate and honest use of the framework. What is not legitimate, and what this lesson asks practitioners to avoid, is presenting the review's existence as though it constituted direct evidence that implementing its standards produces a specific, quantified financial return. The review itself was a policy document synthesising expert judgement and available evidence to recommend standards; it was not a randomised controlled trial of those standards' economic effect, and treating it as though it had quantified a guaranteed return would repeat exactly the kind of evidence

overreach that the earlier lesson on reading evidence claims in this module was designed to help practitioners recognise and avoid.

A worked example illustrates how a practitioner might navigate this terrain in a live commissioning conversation. A finance director at a manufacturing company, considering whether to commission wellbeing training and policy support, asks the practitioner directly what return on investment the organisation can expect. An inexperienced or overselling practitioner might reach for one of the widely circulated aggregate cost-of-presenteeism figures and present it as though it directly translated into savings the organisation could expect to realise. A more honest response acknowledges the genuine substance of the economic case while being transparent about its limits: presenteeism, meaning reduced functioning while at work due to psychological distress, is a well-established phenomenon in the occupational health literature, even though precise organisation-level costings are inherently difficult to produce with confidence; the Thriving at Work review, commissioned by government and grounded in expert synthesis of the available evidence, sets out a widely recognised set of standards that this proposal is designed to help the organisation meet; and the most defensible way to estimate the financial case for this specific organisation is to look at its own absence data, its own turnover costs in roles affected by stress-related attrition, and its own recruitment costs for replacing staff who leave under those conditions, rather than importing a generic industry-wide figure that may bear little relation to this organisation's actual circumstances. This response gives the finance director something more useful than a borrowed statistic: a method for building an organisation-specific business case from data the organisation already holds, which is both more defensible and, in most cases, considerably more persuasive to a financially literate audience than an imported aggregate figure would be.

Beyond absence data, two further organisation-specific figures are worth a practitioner's attention when building a defensible, locally grounded business case, precisely because both are typically already captured somewhere within the organisation's own systems rather than requiring any new, uncertain measurement. The first is turnover attributable to roles that have experienced documented stress-related difficulty, which most organisations can approximate by cross-referencing exit interview data, where stress or mental health is cited as a contributing factor, against the

organisation's own standard calculation of the cost of replacing an employee in that role, typically including recruitment costs, lost productivity during the vacancy and the ramp-up period for a new hire, and any temporary cover arrangements. This figure will rarely be precise, since exit interview disclosure is itself subject to the stigma dynamics discussed elsewhere in this module and will understate the true proportion of stress-related departures, but it has the considerable advantage of being built from the organisation's own, already-trusted costing methodology rather than an imported industry average, which makes it both more defensible under scrutiny and more persuasive to a financially literate audience who will recognise the inputs as their own. The second figure worth examining is the pattern of short-term, repeated absence episodes, sometimes called absence frequency as distinct from absence duration, in teams or roles that other evidence, such as a Management Standards survey or exit interview data, suggests are experiencing high demands relative to resources. A pattern of frequent short absences is a recognised indicator in occupational health practice of accumulating strain that has not yet escalated to a single long absence episode, and demonstrating to a client that their own existing absence data already shows this pattern in the specific teams the organisational assessment has flagged gives the economic case a level of local specificity that no borrowed national statistic could provide, while remaining entirely within what the organisation's own records can support without requiring the practitioner to make any claim that exceeds the evidence.

A final honest framing worth equipping practitioners with concerns the cost of inaction, which is a genuinely different argument from a projected return on investment and is considerably easier to defend evidentially. Rather than claiming that a specific programme will produce a specific saving, a practitioner can honestly point out that the organisation is already incurring costs, through the absence, turnover and recruitment figures discussed above, whether or not it invests in addressing them, and that the real choice facing the organisation is not between a cost and a saving but between an unmanaged, ongoing cost and a deliberate, evidence-informed attempt to understand and reduce it. This framing does not require the practitioner to quantify the size of any future saving, which the available evidence genuinely cannot support with precision, and it has the further advantage of being essentially unfalsifiable in the way an inflated return-on-investment claim is

not: it is simply true that unaddressed psychological strain in a workforce carries a cost, visible in the organisation's own data, and that doing nothing is not a neutral or cost-free option merely because its costs are already familiar and therefore less salient than the cost of a new training programme would be. Presenting the choice this way, as a decision about which already-present cost to engage with honestly rather than a speculative bet on an uncertain future saving, is both more defensible and, in the experience of practitioners who use it consistently, frequently more persuasive to a sceptical financial decision-maker than a confident-sounding but unsupportable return-on-investment figure would be.

Practitioners should leave this lesson with two things held together rather than in tension: a genuine appreciation that the economic case for workplace mental health investment is real, well grounded conceptually in the distinction between absenteeism and presenteeism, and given practical shape by the Thriving at Work review's core standards; and an equally genuine commitment to resisting the temptation to borrow precise-sounding aggregate figures from secondary sources whose methodology the practitioner has not examined. The combination of genuine substance and honest caution is, as earlier lessons in this module have already argued, what distinguishes a credible, evidence-led practitioner from one who simply repeats whatever statistic was most recently circulating in a conference slide deck.

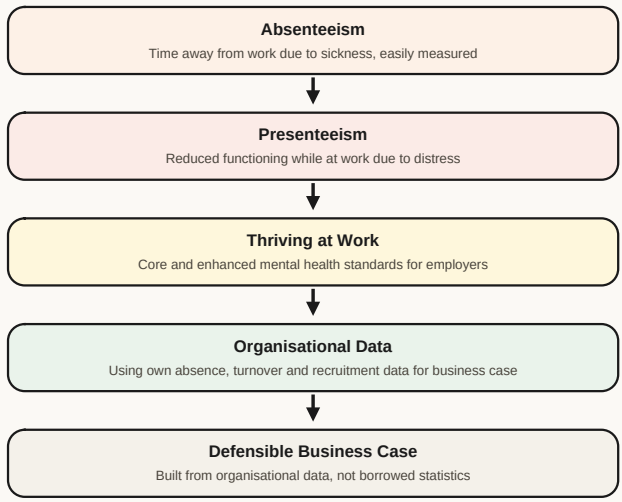


FIGURE 6.1 Economic Case for Workplace Mental Health

This figure outlines the key steps in understanding the economic case for workplace mental health investment, from distinguishing between absenteeism and presenteeism to building a defensible business case. It highlights the importance of using organisational data and avoiding over-reliance on aggregate statistics.

Points worth keeping

- Presenteeism (reduced functioning while present at work due to distress) is a genuine, well-supported phenomenon distinct from absenteeism, but aggregate cost-of-presenteeism figures rest on self-report productivity measures and should be used with more caution than directly measured absence data.
- The Stevenson/Farmer 'Thriving at Work' review (2017), commissioned by the UK government, set out core and enhanced mental health standards for employers (a communicated mental health plan, awareness-raising, open conversation, good working conditions, effective people management, and routine monitoring) that map closely onto the Job Demands-Resources model and HSE Management Standards covered earlier in this module.
- The review is a policy and standards document grounded in expert synthesis, not a trial that quantified a guaranteed financial return, and practitioners should not present it, or imported industry-wide cost figures, as though they carried that level of evidential weight.

- The most defensible organisation-specific business case is typically built from the client's own absence, turnover and recruitment data rather than from borrowed aggregate statistics.

PRACTICE TASK

Draft the two or three sentences you would say to a finance director who asks what return on investment they can expect from a workplace wellbeing programme, using the honest, standards-based and organisation-specific approach described in this lesson rather than quoting an imported aggregate statistic.

The end of the sample

This is a free sample of “Workplace Wellbeing: Foundations”, and it ends here.

Everything you have read here is from the opening of the book, set exactly as it is printed in it. The rest follows in the same way, chapter by chapter, with its worksheets, its examples and its reference list.

The full book is on sale in the Talking Therapies UK shop. A copy is yours to keep, as a PDF and an ebook, licensed to you by name: a short line saying so is printed at the foot of each page of the PDF and on the title page of the ebook.