



OPINION PIECE

The friction problem: why learners don't read learning outcomes and what design can realistically fix

Mark Frohnsdorff 

King Faisal Air Academy, Saudi Arabia

ABSTRACT

Learning outcomes are everywhere in higher education. Many learners do not engage with them. This piece argues that friction in design contributes to this gap: the cognitive effort created by sentence length, syntactic complexity, and abstract performance criteria. Psycholinguistic evidence and the limited empirical literature on learner engagement support this claim. When friction exceeds perceived benefit, learners choose more efficient paths. They scan assignment briefs, consult peers, or draw on prior experience. This is one layer of a multifactorial problem; learner maturity, motivation, and context also matter. Seven evidence-informed design principles are proposed to reduce friction and increase the probability of engagement. Better design does not guarantee changed behaviour. Poor design creates an unnecessary barrier. This piece is a provocation: are our well-intentioned instruments actually usable by the learners they are meant to serve?

KEYWORDS: learning outcomes, cognitive load, learner engagement, instructional design.

ARTICLE HISTORY: Received 15 April 2026. Accepted 29 July 2026.

Introduction: the gap between presence and use

Learning outcomes are everywhere in higher education. They populate course documents, assignment briefs, module handbooks, and virtual learning environments. Their presence is near-universal; their use is not.

This is not a new observation, but it is an underexamined one. Brooks et al. (2014, p. 725) conducted one of the few studies to investigate learner engagement with outcomes directly. They surveyed 918 students across three disciplines at the University of Leicester. While 81% agreed outcomes were useful learning aids, 49% reported that learning outcomes could only be fully understood at the end of a module, suggesting limited engagement during the learning process itself. The researchers noted that students engaged more with outcomes as the course progressed, implying a learning curve in decoding institutional language rather than immediate usability. Recent work on syllabus design reaches a similar conclusion: the information students value is not always the information that course documents foreground (Perrault, 2026).

Maier and Klotz (2025) report a related pattern in a different setting. Analysing over 20,000 formative assessment cases from a digital learning system used by secondary school students, they found that learners with low prior knowledge were strongly associated with reduced feedback seeking and comprehension. The setting differs from higher education, but the mechanism travels: instructional guidance is often least processed by those who need it most, whatever its quality. In higher education, learners are expected to be autonomous, yet they encounter complex academic prose with little teacher mediation. Lecturers rarely walk learners through the wording of an outcome; the text is expected to speak for itself. These conditions make disengagement particularly acute.

These findings challenge a core assumption in learning outcome design: that stating criteria enables their use. I propose that this assumption overlooks *friction*, the cognitive effort required to process complex academic prose. The term is borrowed from cognitive load theory and from user-experience design, where friction names any feature of a task that raises the effort needed to complete it (Sweller et al., 2011). When friction exceeds perceived benefit, learners rationally disengage. They turn instead to assignment instructions, peer advice, exemplars, or prior experience: paths that are more immediately legible.

For Learning Development practitioners, this matters directly. Much of our work involves helping learners decode the implicit conventions of academic assessment. If the instruments we direct learners toward are poorly designed, we are, in effect, directing them toward a barrier. This piece argues that linguistic complexity is the primary modifiable source of friction in learning outcomes, and proposes evidence-informed principles for reducing it. Abstract



criteria and unclear standards raise effort too; but linguistic complexity is the source most open to redesign. One distinction matters. Not all friction is unproductive: some effort is the work of learning itself. The target here is unproductive friction, the effort imposed by how an outcome is written rather than by what it asks the learner to think. Reducing it is necessary but not sufficient for learner self-regulation (Zimmerman, 2002).

The friction hypothesis: a design explanation

I propose a specific, bounded idea: learners may not use learning outcomes because of friction created by sentence length, syntactic complexity, and the cognitive effort required to process them. Consider the following, which is representative of the module-level learning outcomes typically found in higher education handbooks and learner-facing course documentation: 'On successful completion of this module, students will be able to critically evaluate theoretical perspectives on organisational change, drawing on empirical evidence to construct sustained arguments that acknowledge competing interpretations and demonstrate appropriate academic practice'.

This outcome contains 35 words, three embedded clauses, and abstract performance criteria. The learner must hold multiple elements in working memory, decode nested structures, and infer what success looks like. The cognitive cost is substantial. Faced with this effort, learners may instead scan assignment instructions, consult exemplars, ask peers, or rely on prior experience: paths that are often more immediately legible. Friction creates rational disengagement. Therefore, careful consideration of the language used in learning outcomes is necessary (Kennedy, 2007).

The hypothesis connects to established mechanisms in cognitive psychology. Norman et al. (1992) demonstrated that syntactic complexity impairs comprehension under working memory constraints, with effects present in both timed and untimed conditions. Bartek et al. (2011) confirmed this using self-paced reading and eye tracking experiments. They found that placing more words between syntactically related elements slows comprehension consistently. Longer sentences impose this cost directly: more intervening material means more to hold in memory before the sentence can be resolved. The plain language scholar Iva Cheung (2017, cited in Henville, 2019), names the mechanism precisely. When a subject and verb are separated, the reader holds the subject in working memory, navigates all intervening



words, and only then processes the sentence. This creates extraneous cognitive load: effort imposed by poor design, not by content difficulty.

Criteria abstraction presents a further problem. León et al. (2023), in a meta-analysis of self-assessment accuracy, found that learners' ability to judge their own work depends critically on their understanding of the criteria against which they are assessed. Vague standards such as 'appropriate academic tone' and 'critical engagement' require interpretive resources that compete with comprehension. Carr (2017) establishes that modal verbs such as *must*, *should*, and *will* are context-sensitive and interpreted against inferred norms: 'students must demonstrate critical thinking' may signal strong obligation, weak recommendation, or an aspirational goal, depending on the reader's familiarity with institutional discourse. This ambiguity is not incidental. As Hussey and Smith (2002) argued in their critique of the learning outcomes movement, outcomes are frequently written to satisfy institutional and managerial requirements rather than to communicate with learners, and the language reflects that purpose.

A common response to the abstraction problem is the use of the action verbs associated with the revised Bloom's taxonomy (Anderson & Krathwohl, 2001): 'analyse,' 'evaluate,' 'synthesise.' These appear precise because they name a type of cognitive process. However, Hess et al. (2009) show that Bloom's verbs specify the type of thinking required but not the depth. The same verb can describe a surface task or a demanding one. 'Analyse the causes' can mean recall a list from notes or construct an original argument from competing evidence. Both appear to satisfy the same criterion. When depth of thinking remains implicit, learners can decode the verb but still cannot determine whether their work meets the intended standard. This is a design problem that a better vocabulary alone does not solve.

Quantified criteria do not escape the problem, they relocate it. A specification that measures amount, such as 80% accuracy or five factors listed, tells learners how much to produce but leaves the cognitive quality required just as unclear. A learner can list five factors without grasping why those factors matter. Precision about quantity can therefore create an illusion of clarity while the demand that determines success stays hidden. This is not a hypothetical concern. A Quality Assurance Agency for Higher Education (QAA)-funded project across four institutions investigated what it called 'mysterious quotients' of common assessment terms, and found that students report low confidence in understanding the very words these



outcomes depend on, including 'synthesise,' 'critical,' and 'analyse' (QAA, 2022). Learners already tell us which words they cannot reliably decode. These mechanisms, working memory load, modal ambiguity, and hidden cognitive depth, converge on a single implication: linguistic complexity in learning outcomes is not a trivial matter of style but a barrier that affects whether learners can act on what they read.

One layer among many: positioning the contribution

I acknowledge explicitly that other factors shape whether learners engage with learning outcomes. Self-directed learning capacity develops over time, and novice learners may lack the metacognitive resources to use outcomes effectively regardless of how they are written. Learners unfamiliar with criterion-referenced assessment may not recognise the utility of outcomes at all. Motivation, perceived relevance, assessment culture, and teacher guidance all operate independently of design quality.

My contribution addresses the design layer only: removing obvious friction so that higher-order capabilities (self-assessment, self-regulation, autonomous learning) can develop. Usable outcomes are a necessary but insufficient condition for these capabilities. A learner cannot self-assess against criteria they have not read or cannot process; but readable outcomes alone do not guarantee self-assessment will occur.

This positioning matters. I am not claiming that better-written outcomes will transform learner behaviour. I am claiming that poorly written outcomes create an unnecessary barrier that design can remove, thereby increasing the probability of engagement without ensuring it. This is recognisably the work of Learning Development. Johnson (2018) characterises the field as occupying a niche for mediation and demystification, helping learners interpret the conventions that institutions leave implicit. Practitioners who work at the boundary between institutional expectations and learner readiness are well placed to identify and remove such barriers. It is precisely our kind of problem.

Design principles: reducing friction

Based on the mechanisms reviewed, I propose seven evidence-informed principles. These address comprehension and accessibility, the prerequisite layer for any pedagogical function of learning outcomes.



Table 1. Evidence-informed design principles.

Principle	Application	Rationale
Constrain sentence length	Maximum 20 words; 12–15 for multilingual contexts	Longer sentences increase distance between dependent elements, degrading memory processes in comprehension (Bartek et al., 2011; Norman et al., 1992)
Minimise syntactic complexity	One clause per outcome; subject and verb adjacent	Subject-verb separation creates extraneous cognitive load (Cheung, 2017, cited in Henville, 2019)
Prefer performative verbs	'Analyse,' 'write,' 'demonstrate' rather than modal constructions	Modals introduce context-sensitive ambiguity (Carr, 2017)
Quantify abstract criteria	'Fewer than three errors per 100 words' rather than 'appropriate accuracy'	Concrete standards give learners a check they can apply without interpreting vague terms, supporting more accurate self-assessment (León et al., 2023)
Sequence by cognitive demand	Isolated criteria before combined criteria	Introducing criteria one at a time before combining them manages intrinsic load through progressive disclosure, reducing what the learner must hold at once (Sweller et al., 2011)
Mandate exemplar pairing	Worked examples at threshold, merit, and distinction levels	Exemplars make abstract criteria concrete (Andrade & Du, 2007)
Apply heightened simplicity for multilingual contexts	Reduce length and complexity further; consider first language (L1) glosses	Learners of second or additional languages (L2) face amplified extraneous cognitive load; plain language is an ethical imperative for those managing additional cognitive burdens (Cheung, 2017, cited in Henville, 2019)

These principles are testable and refinable, not prescriptive. Whether the recommended sentence length of 20 words is optimal, whether exemplar pairing consistently increases engagement, and whether the principles hold equally across disciplines and institutional contexts are open empirical questions. The aim is to give the Learning Development community a set of specific, falsifiable propositions to examine rather than a set of rules to adopt.

One design tension is worth naming here. The principle of quantifying abstract criteria — ‘fewer than three errors per 100 words’ rather than ‘appropriate accuracy’ — increases accessibility and enables learner self-checking. However, greater quantification tends toward lower-order demonstrations. A criterion that asks learners to list five factors is self-checkable; a criterion that asks them to argue the relative significance of those factors is not, at least not without further specification. Performative verbs reach their structural limit at higher cognitive demands: ‘analyse’, ‘evaluate’, and ‘synthesise’ are clear in form but opaque in depth. Addressing this tension is a design challenge that extends beyond sentence-level friction, and one that warrants its own treatment (Hess et al., 2009; León et al., 2023).

Conclusion: from presence to usability

The learning outcomes movement has achieved near-universal presence in higher education. What remains contested is whether learners actually use them. I have proposed that friction in design, specifically sentence length, syntactic complexity, and cognitive load, contributes to this gap. The evidence is suggestive rather than conclusive: psycholinguistic research establishes the mechanisms, but direct empirical testing with learning outcomes remains limited. Brooks et al. (2014) found that nearly half of students understood outcomes fully only in retrospect and Maier and Klotz (2025) found that low-prior-knowledge students actively avoid instructional guidance. These findings align with the friction hypothesis without proving it.

The argument is bounded. I address one layer of a multifactorial problem, the design layer, while acknowledging that learner maturity, motivation, and context also shape engagement. Poor design creates unnecessary barriers; good design removes them, no more. These principles apply to learner-facing outcome statements, not to academic writing in general, where specialist readers bring different resources and motivated attention to the task.



The Learning Development community is well placed to test, refine, or refute this hypothesis, and to act on it. The role is not only one of mediation but of consultancy. Learning developers can review the language of outcomes, briefs, and handbooks, and feed back to course teams on where unproductive friction has crept in. They can run benchmarking workshops in which students apply outcomes to exemplars and surface the points where the wording fails them. Recent sector work points the same way. The QAA (2022) toolkit addresses unclear assessment language through staff and student dialogue. The principles proposed here address the layer beneath that conversation: the toolkit asks staff to explain the language; these principles ask them to write less that needs explaining. The two are complementary, and learning developers are placed to carry both. Beyond testing the principles, we can ask whether they hold across disciplines, year groups, and linguistic backgrounds, and whether exemplar pairing changes what learners actually do with criteria. The emergence of generative artificial intelligence (AI) adds a further dimension worth attending to. If learners increasingly outsource outcome decoding to AI tools, the design problem may be displaced rather than solved and this warrants its own empirical attention from our community.

The aim is not to prescribe a new orthodoxy but to provoke critical examination of whether our instruments serve their intended audiences. If learning outcomes are not usable, they will not be used, regardless of their pedagogical intent. That gap, between an outcome that is present and an outcome that is usable, is the friction problem, and it belongs to all of us in Learning Development.

Disclosure statement

The authors used the following generative AI tools in the preparation of this manuscript: Claude (Anthropic), via the claude.ai platform. The tasks performed by Claude include: critical reader, to act as a peer reviewer and identify potential grounds for rejection relating to academic rigour, reasoning gaps, citation standards, and APA style compliance. For example, Claude flagged a source on cognitive load and sentence length as not peer-reviewed and suggested that a peer-reviewed alternative be located to support the same point. The author addressed feedback and completed revisions to the manuscript. Claude did not generate ideas, text, or structure; it functioned as a critical reviewer only. All content and revisions are the author's own work. The authors have complied with the journal's principles of AI use.



Funding

The author received no funding for this research and declares no competing interests.

References

- Anderson, L. W., & Krathwohl, D. R. (Eds.). (2001). *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives*. Longman.
- Andrade, H., & Du, Y. (2007). Student responses to criteria-referenced self-assessment. *Assessment & Evaluation in Higher Education*, 32(2), 159–181.
<https://doi.org/10.1080/02602930600801928>
- Bartek, B., Lewis, R. L., Vasishth, S., & Smith, M. R. (2011). In search of on-line locality effects in sentence comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(5), 1178–1198. <https://doi.org/10.1037/a0024194>
- Brooks, S., Dobbins, K., Scott, J. J. A., Rawlinson, M., & Norman, R. I. (2014). Learning about learning outcomes: The student perspective. *Teaching in Higher Education*, 19(6), 721–733.
<https://doi.org/10.1080/13562517.2014.901964>
- Carr, J. (2017). Deontic modals. In T. McPherson, & D. Plunkett (Eds.), *The Routledge handbook of metaethics* (pp. 194–210). Routledge.
- Henville, L. (2019, November 8). *Reducing the weight of your words*. University Affairs.
<https://universityaffairs.ca/career-advice/reducing-the-weight-of-your-words/>
- Hess, K. K., Jones, B. S., Carlock, D., & Walkup, J. R. (2009). *Cognitive rigor: Blending the strengths of Bloom's taxonomy and Webb's depth of knowledge to enhance classroom-level processes*. ERIC. <https://eric.ed.gov/?id=ED517804>
- Hussey, T., & Smith, P. (2002). The trouble with learning outcomes. *Active Learning in Higher Education*, 3(3), 220–233. <https://doi.org/10.1177/1469787402003003003>
- Johnson, I. (2018). Driving learning development professionalism forward from within. *Journal of Learning Development in Higher Education*, (13). <https://doi.org/10.47408/jldhe.v0i0.470>
- Kennedy, D. (2007). *Writing and using learning outcomes: A practical guide*. University College Cork. <http://hdl.handle.net/10468/1613>
- León, S. P., Panadero, E., & García-Martínez, I. (2023). How accurate are our students? A meta-analytic systematic review on self-assessment scoring accuracy. *Educational Psychology Review*, 35, Article 106. <https://doi.org/10.1007/s10648-023-09819-0>



- Maier, U., & Klotz, C. (2025). Students ignore their mistakes: Elaborated error feedback processing in a digital learning system. *Contemporary Educational Psychology*, 82, Article 102395. <https://doi.org/10.1016/j.cedpsych.2025.102395>
- Norman, S., Kemper, S., & Kynette, D. (1992). Adults' reading comprehension: Effects of syntactic complexity and working memory. *Journal of Gerontology*, 47(4), 258–265. <https://doi.org/10.1093/geronj/47.4.P258>
- Perrault, E. K. (2026). Syllabus information and the relative importance of content they can contain. *College Teaching*, 1–9. <https://doi.org/10.1080/87567555.2026.2612968>
- Quality Assurance Agency for Higher Education. (2022). *Toolkit: Making the language of assessment inclusive*. <https://www.qaa.ac.uk/membership/collaborative-enhancement-projects/assessment/making-the-language-of-assessment-inclusive>
- Sweller, J., Ayres, P., & Kalyuga, S. (2011). *Cognitive load theory*. Springer. <https://doi.org/10.1007/978-1-4419-8126-4>
- Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory Into Practice*, 41(2), 64–70. https://doi.org/10.1207/s15430421tip4102_2

Author details

Mark Frohnsdorff works in professional training contexts where language functions as a performance and safety tool. His interests span classroom practice, vocabulary development, and evidence-informed course and task design, with a focus on how learners perform under cognitive and communicative pressure. He draws on research, classroom observation, and cross-disciplinary thinking to prepare learners for real-world demands beyond the classroom.

ORCID

Mark Frohnsdorff  <https://orcid.org/0009-0000-6669-7208>

