A WHITE PAPER

Travel has a training problem.

Why the architecture of trade engagement is failing the industry that depends on it, and what an honest redesign looks like.

BY RUSSELL MEARA

Founder, Hablo

Published 2026

EXECUTIVE SUMMARY

The case in one page.

The global travel industry depends on the travel trade. The trade depends on training. And the training infrastructure the industry has built to support it is failing more often than it is serving, in ways that almost nobody is naming honestly.

This paper argues that the failure is architectural rather than executional. It is not a content problem, not a budget problem, not a motivation problem on the part of advisors. It is a structural mismatch between the platforms the industry deploys, the realities of an advisor's working life, and the biology of how human memory actually works. The result is a multi-billion-pound investment in trade engagement that produces certifications, completions, and dashboards, but does not produce confident advisor selling at the point of sale, which is the only outcome that should ever have mattered.

The argument proceeds in six parts. First, the operational reality: advisors have roughly twenty-four minutes a week for formal learning, and the industry has built training as if they had hours. Second, the biology: the Forgetting Curve, settled science for over a century, quietly destroys around 90% of unreinforced training within a week. Third, the engagement architecture that would defend against that decay: habit-forming product design and the visibility-based motivation structures that consumer technology has spent fifteen years perfecting. Fourth, the structural shift from the Learning Management System to the Learning Experience Platform that makes everything else possible. Fifth, the philosophical fork between AI built to replace the trade and AI built to amplify it. Sixth, the only metric that ultimately decides whether any of this is working: Return on Education, measured as confident advisor selling correlated to training, at the level of the individual advisor.

The paper closes with a single question the industry can no longer avoid: by the end of this decade, what is your trade engagement strategy actually being measured against?

ABOUT THIS PAPER

Who this is for, and what it argues.

This paper is written for the people who decide how the travel industry trains and engages the trade. Destination marketing leaders. Supplier engagement directors. Hotel-group commercial teams. Agency owners and senior advisors. Trade-press observers and analysts. Anyone holding the budget for trade engagement, or anyone evaluating where the next decade of investment should go.

It is the consolidated form of an eight-issue LinkedIn newsletter series I published in 2026 under the title Travel has a Training Problem. The series ran weekly across two months and developed a single connected argument across eight short instalments. This paper distils that argument into a single document that can be read in one sitting and referenced afterwards. Some material has been compressed. Some has been added. The core position is unchanged.

I am the founder of Hablo, a next generation trade engagement and training platform for the travel trade – aligned with the principles of a Learning Experience Platform. The platform is referenced briefly toward the end of this paper, as the working realisation of several of the arguments laid out. But the paper is not a Hablo pitch and is not structured as one. It is an industry-level argument about how trade engagement needs to be redesigned. Whether you ever use Hablo or not, the case being made here is one I am genuinely convinced any operator serious about the trade needs to engage with.

SECTION ONE

The architecture that fails the people it depends on.

The global travel industry is in a strange position. Consumer demand has held up far better than many forecast, even against a backdrop of geopolitical instability and regional volatility. Destinations are more complex and faster-moving than ever. Visa rules shift overnight, border policies change without warning, geopolitical events demand crisis communications the industry has no infrastructure to deliver in real time. Travellers are arriving better informed, more demanding, and less tolerant of generic itineraries than at any point in the last twenty years.

And yet the primary mechanism that destination marketing organisations, airlines, and hotel groups use to educate the trade is rooted in the digital architecture of the early 2000s. The Learning Management System. The one-and-done certification. The 45-minute webinar. The static destination specialist programme. The course catalogue that the advisor enrols in, completes, and receives a badge for.

> *This is not a content problem. The content is often excellent. The model the content lives inside is the problem.*

The operational reality of an advisor's day

The first thing to understand is how training fits into an advisor's actual working life. The answer is: barely.

The average employee in any sector has roughly twenty-four minutes a week available for formal learning during the working week, around 1% of a full-time job. The figure comes from research by Bersin & Associates and has been cited continuously by learning and development research firms for the past decade. It is a corporate average, not a travel-trade specific number, but every advisor I have spoken to has confirmed it sounds about right, and many would argue it is generous.

Twenty-four minutes is what is left over once the cognitive load of an advisor's actual job has been accounted for. Quoting a complex multi-centre itinerary. Holding for forty minutes on a supplier line to rebook a disrupted flight. Managing a client whose visa has changed status overnight. Writing tomorrow's proposals. Following up on yesterday's deposits. The 24 minutes lives in the cracks: the moment between calls, the gap before a meeting, the train ride home. And the suppliers asking advisors to spend those minutes on a 45-minute webinar are not competing with idleness. They are competing with the moment the advisor could be using to actually breathe.

There is a deeper issue underneath the maths. Cognitive psychology has a concept called cognitive load, which says that human working memory has limited capacity. When that capacity is full, which it is for most advisors for most of their day, new information does not get encoded. It bounces off. A 45-minute webinar that tries to teach an entire destination strategy in one session is not failing because the content is bad. It is failing because the brain it is being delivered to is already at capacity. The training does not stick because biologically, in that moment, it cannot. This is not a motivation problem. The advisors are not lazy and they are not disengaged. They are full. And the response to a full vessel is not a bigger jug.

The other forty hours

The twenty-four-minute number is a working-week statistic. It tells us how much formal learning capacity exists during the contracted hours of an advisor's job. It does not tell us anything about the rest of their life.

Take eight hours for sleep, eight for work, and there are still roughly forty waking hours every week that sit outside the office. The commute. The evening on the sofa. The half-hour at the gym. The Sunday morning scroll. Those forty hours are where the more interesting opportunity lives, because almost none of those minutes belong to anyone in the travel industry. They belong to TikTok, Instagram, the BBC, Spotify, WhatsApp, the football scores. They belong to the platforms that have spent fifteen years and billions of dollars perfecting the art of fitting useful, interesting, addictive content into the spaces between everything else.

The supplier-side conversation about trade engagement is almost entirely focused on the twenty-four-minute working-week window. It should also be focused on the other forty hours, because that is where habit forms, where loyalty is built, and where the next generation of trade engagement will actually be won.

SECTION TWO

The biology nobody is accounting for.

Even short content, delivered in the right moments, is not enough on its own. There is a deeper biological problem underneath every certification campaign the industry runs: one that has been quietly destroying the return on training investment for decades, that is not a matter of opinion or interpretation, and that almost nobody on the supplier side is talking about honestly.

It is called the Forgetting Curve.

The science nobody disputes

In the late 1880s, a German psychologist called Hermann Ebbinghaus did one of the most important experiments in the history of cognitive science. He memorised lists of nonsense syllables, then tested himself at intervals to see how much he could recall. What he found became known as the Forgetting Curve. Without active reinforcement, memory decays exponentially. By the end of the first 24 hours, retention has dropped to roughly 30 to 50% of the original. By the end of the first week, around 90% of the new information is gone.

Ebbinghaus published this in 1885. The research has been replicated continuously for over 140 years, including by modern neuroscience using techniques he could not have imagined. The curve has held up under every test. It is not contested in the scientific literature. It is one of the most settled findings in the entire field of memory research.

> *By the end of the first week, around 90% of the new information is gone. This is not opinion. It is biology.*

What this means in money

A traditional DMO destination specialist programme costs anywhere from a few thousand pounds for a small market to several hundred thousand for a flagship multi-region campaign. Whatever the budget, the content is usually excellent, the production values where the budget allows are high, and the campaign messaging is polished. Agents are invited, incentivised, sometimes recognised publicly at industry events.

Within seven days of completion, by the science we have known for over a century, roughly 90% of the specific selling points that programme was designed to embed are no longer accessible to the advisors who completed it. The certificate has been awarded. The badge appears on a profile. The compliance form has been ticked. But the knowledge, the actual product knowledge the campaign existed to install, is gone.

Put more bluntly: some campaigns are spending six figures, others a few thousand, and the biological maths is the same regardless. Whatever the budget, the industry is paying full price to embed information that has a half-life shorter than the time it took to design the campaign. And nobody is measuring the decay, because the metrics the industry uses, including the number of certified agents, the completion rate, and the course satisfaction score, are all captured at the moment of completion, before the curve has done its work.

This is the model failing, not the advisors

It is tempting, for the suppliers who hold these budgets, to interpret this data as a problem with advisors. They were not engaged enough. They did not study hard enough. They are not committed to learning. This interpretation is wrong, and it has cost the industry millions.

Ebbinghaus's research applies to every human brain. It applies to doctors memorising drug interactions, to pilots learning new airframes, to lawyers preparing for case hearings. The Forgetting Curve is not a measure of advisor commitment. It is a measure of how human memory works. If the design of a training campaign produces 90% knowledge loss within a week, the failure is in the campaign design, not in the people who completed it.

The science also points at the antidote. Ebbinghaus found that if the same material is re-encountered at calculated intervals, the principle that later became known as spaced repetition, the curve flattens dramatically. Each reinforcement reduces the slope of subsequent decay. After several well-timed reinforcements, the material moves from working memory into long-term memory, where it stays.

This is the only intervention the research consistently supports. It is also, structurally, almost impossible to deliver through a traditional Learning Management System. An LMS is built around the certification event: the assessment, the badge, the completion record. It has no native capability to track an individual learner's decay curve, predict when reinforcement is needed, and proactively surface the right piece of content at the right moment. The architecture was not designed for that job. It was designed for compliance reporting.

SECTION THREE

The engagement architecture that biology demands.

Knowing that spaced repetition is the answer the science supports raises an immediate practical question that the research itself does not address. How do you actually get a busy travel advisor to engage with reinforcement content frequently enough for it to work? You cannot legislate it. You cannot mandate it. You cannot incentivise your way through it. The training has to become a habit, or it does not happen at all.

There is one body of work the industry can borrow from. Consumer technology has spent fifteen years and billions of dollars learning how to build daily habits. The playbook exists. It is not a secret. It is described in detail in Nir Eyal's 2014 book Hooked, which set out a four-phase loop that turns a behaviour into an automatic response. The framework has been used by almost every major consumer technology company in the world over the last decade. It is also, almost entirely, absent from B2B travel training.

The Hook Model, applied to the trade

The four phases are simple. A trigger prompts the behaviour. An action is taken in anticipation of a reward. A variable reward arrives: not always the same, not always predictable. And finally, a small investment is made that loads the next trigger. Repeat the loop enough times and the trigger no longer needs to come from outside. The user provides their own.

Translating that loop into B2B trade engagement requires four shifts from how the industry currently operates.

First, the shift from external to internal triggers. Most trade engagement today relies on email reminders, sales rep follow-ups, and ‘have you completed your training yet?’ messages. These work for acquisition. They are unsustainable for retention. What builds habit is the cultivation of internal triggers: the mental association between a professional moment and the platform. ‘I need to feel confident before this client call.’ ‘I want to see if anything has changed about this destination before I draft the proposal.’ ‘I want to see if I have moved up the leaderboard.’ When the platform becomes the automatic response to those professional needs, a habit has formed.

Second, the relentless reduction of friction in the action itself. Most legacy training platforms ask too much: find the right URL, remember the password, navigate the catalogue, watch a video, take a quiz, complete a certification. Habit-forming platforms strip the action to a single tap. Open. Scroll. Done. Trade engagement platforms, almost without exception, have not made this shift.

Third, the embrace of variable rewards. This is the part of the Hook Model that goes hardest against B2B engagement instinct. Predictable rewards, such as five points for completing a module or ten for tracking a booking, do not generate the dopamine response that builds craving. What works is unpredictability in the right form. Not unfair. Not arbitrary. Surprising. A badge the advisor did not know existed. A public acknowledgement from a DMO rep that lands without warning. A new module on a topic the advisor did not know they wanted to read about until they saw it.

Fourth, the investment loop. Small acts of personalisation, such as listing specialisms, following partners, tracking bookings, and writing reviews, make the platform more valuable to the user over time. Every interaction trains the platform to predict the user better. Walking away means walking away from years of compounded value.

Why visibility is the actual currency

There is a related insight that the Hook Model does not quite capture but that becomes obvious once it is named. Almost every supplier-side loyalty programme in travel has been built around the wrong currency.

The default architecture issues points. Earn the points, redeem the points, the points are the motivation. From conversations with people who have been close to these programmes, the pattern that comes back is consistent. The first few months produce a burst of activity. Around month six, engagement starts to drop. By twelve months, the programme is quietly winding down. The team running it is having difficult conversations about whether to relaunch, rebrand, or retire it altogether.

The points were never the motivation. They were a proxy for the motivation. What advisors actually want is the recognition that comes with earning the points: the public, named, specific acknowledgement of doing the work well, in front of their peers. And in programmes that delivered the proxy but never delivered the recognition, advisors lost interest the moment the novelty wore off.

> *Points are a proxy. Visibility is the currency.*

What motivates sales professionals consistently, not just in travel but in every industry where the work is complex, relationship-led, and visible, is three structures of recognition. Status: being publicly named as expert in something specific. Mastery: being publicly named as having advanced. Belonging: being publicly named as part of an elite community. All three are visibility mechanics. None of them is a points mechanic.

The diagnostic question every supplier should be asking is the simple one. Is the structure of this programme designed to make advisors more visible, or to make their behaviour more measurable? Because most programmes are doing the second when they should be doing the first.

SECTION FOUR

From compliance machine to engagement ecosystem.

Everything argued so far, including short-form content, habit-forming engagement, visibility-based motivation, and spaced repetition delivered automatically, has a single underlying requirement that the industry has been avoiding for two decades. None of it is deliverable through the Learning Management System architecture the industry has built around. The shift the industry has to make is from the LMS to the Learning Experience Platform. It is the architectural change that makes everything else possible.

The question every LMS was built to answer

Every Learning Management System in the travel industry was designed to answer one question with high precision: did the agent complete the course? The certification record. The completion percentage. The compliance dashboard. The exportable spreadsheet of who has finished what. The architecture of an LMS, including every database table, every administrative screen, and every reporting view, is built around producing accurate, auditable answers to this one question.

Which would be useful if the question mattered.

The question that actually matters

Consider what a DMO, a supplier, or a hotel group is really paying for when they invest in trade training. It is not certificates. It is not completion percentages. It is not the satisfaction score from a post-course survey. What they are paying for is conversion. When an advisor is on a call tomorrow with a client deciding between two destinations, the question is not whether that advisor was certified last quarter. The question is whether they can confidently recall the differentiators that turn the conversation into a booking. The training existed to make that recall possible. The certification is just a paperwork artefact of the moment the training happened.

So the question every training platform should actually be optimised to answer is this: can the advisor sell the destination confidently tomorrow? The LMS has no way of answering it. Not because the data does not exist, but because the architecture was never built to look for it. The system stops measuring at the moment of completion, which is the moment before the Forgetting Curve does its work.

What an LXP actually is

It is tempting to think the LMS could be fixed with the right add-ons. Bolt on a quiz module for retention testing. Add a recommendation engine for personalisation. Layer some gamification on top. Plug in an AI assistant. This is the same instinct that produces what software engineers call ‘spaghetti code’: systems patched together over decades, each new layer working around the limitations of the layer beneath. The travel industry's training infrastructure is, almost without exception, in this state.

It does not work because the underlying architecture is wrong. The LMS was built to be a filing cabinet for certification records. No amount of layering will turn a filing cabinet into a habit-forming experience. The categories of thing are different.

A Learning Experience Platform is not an LMS with better visuals. It is a different category of software entirely, built to answer a different question and optimised for a different outcome.

Where the LMS is administrator-centric, the LXP is user-centric. The first asks: what does the organisation need the user to learn? The second asks: what does the user want to learn to be successful? Where the LMS pushes mandatory courses at users on a schedule, the LXP pulls users into a dynamic ecosystem of content they want to spend time in. Where the LMS treats content as static and dated, for example a course written eighteen months ago, still in the catalogue, still being assigned, the LXP treats content as a living feed, updated continuously, surfaced based on relevance. Where the LMS sees learners as isolated individuals taking solitary courses, the LXP is designed to connect them, oriented toward community and shared learning. Where the LMS reports on completion, the LXP measures engagement, retention, and the correlation between training behaviour and commercial outcomes.

> *These are not the same platform with different settings.*
> *They are different species.*

Why user-centric architecture matters

An LMS treats expertise as something that flows downward. The head office of the tourism board creates the curriculum, the regional teams approve it, the agents consume it, the certifications are awarded. Every advisor sees the same catalogue, in the same order, regardless of who they are, where they sell, or what they specialise in. The agent is a passive recipient of a one-size-fits-all programme.

An LXP starts from the opposite premise. The most useful thing it can do is show each advisor the content that is relevant to them, specifically. In practice, this means tailoring what each advisor sees based on the parameters that actually matter to their work: their region, their role, and their specialisms. An advisor in Germany who sells luxury and cruise does not see the same feed as an advisor in Brazil who sells adventure and family travel. The platform uses those parameters to recommend the modules, updates, and destination content most relevant to each individual, and to nudge them toward what they are most likely to need, while always leaving the final decision about what to engage with in the advisor's own hands.

It guides. It does not dictate. The advisor remains the protagonist of their own professional development, not the subject of it.

SECTION FIVE

The philosophical fork that determines the future.

Two philosophies of AI are being deployed across the travel industry right now. They produce different products, serve different customers, and lead to different futures for the trade. Most operators have not stopped to notice which one they are buying.

Replace or amplify

Walk through any travel-tech conference in 2026 and you will see two fundamentally different visions of what AI is for, often presented on neighbouring stands by people using almost identical vocabulary. The vocabulary is the same. The products are not.

The first philosophy says the value of AI in travel is to replace the trade. AI travel planners marketed directly to consumers. Booking assistants that bypass the advisor entirely. Aggregators racing to build the conversational interface that turns travel research into a chatbot conversation and ends in a direct booking. The end-state of this philosophy is a world in which the advisor is the friction the technology is built to remove.

The second philosophy says the value of AI in travel is to give the trade superpowers. AI that sits beside the advisor, not in front of them. AI that makes a good advisor great, and a great advisor extraordinary. AI that hands the advisor knowledge they would never have had time to acquire, anticipates the questions their client is about to ask, and turns thirty years of supplier complexity into a single conversational layer the advisor can rely on at the point of sale. The end-state of this philosophy is a world in which the advisor is the asset the technology is built to amplify.

> *These are not marketing distinctions. They are architectural ones.*

They produce different products, sold to different customers, with different commercial models, optimised for different outcomes. The choice between them is one of the most important decisions a DMO, a supplier, or a hotel group will make about AI in the next five years.

Why amplification needs a walled garden

Once the amplification philosophy is chosen, several downstream design decisions follow that go against the grain of how most of the AI industry has been building consumer-facing tools.

The most important is what is sometimes called a walled garden. An amplification-philosophy platform has to be built around content that has been verified: the DMO's official destination information, the supplier's actual product data, the operator's confirmed routes and schedules. It does not draw from the open internet. It does not improvise. It does not fill in gaps with plausible-sounding inventions.

The reason is straightforward. The AI has to make the advisor more knowledgeable, more confident, more accurate at the point of sale than they would be without it. Which means the AI cannot, under any circumstances, hallucinate. An advisor who confidently recommends a hotel that does not exist, a visa requirement that was repealed last year, or a route that no carrier actually operates is not amplified. They are damaged. The AI has made them a less reliable professional than they were before they used it.

Open large language models hallucinate. They are statistically likely to do so on any sufficiently niche query, which describes most destination knowledge. The consumer using a general-purpose AI to plan a holiday can verify, hedge, or live with the occasional fabrication. The advisor on a call with a high-value client at 11.30am cannot.

An amplification AI also acknowledges, honestly, what AI cannot do. It cannot replicate what an experienced advisor does when things go wrong: when the flight cancels at midnight, when border policy changes three days before departure, when a client's grandfather dies in another country and they need to be flown home tomorrow. And underneath all of that there is something more practical that almost nobody designing AI travel tools has thought about. When the system breaks, the resolution does not happen online. It happens on the phone. The advisor sits on hold to a supplier reservation line for forty minutes, explains a complex multi-party situation to a duty manager in another timezone, and negotiates an outcome that no AI can negotiate because no AI can pick up the phone and speak human-to-human with a hotelier at 11pm. While airlines, hotels, and ground operators still resolve disruption through human-staffed phone lines, which they will for years to come, the advisor's phone is the resolution mechanism. The AI cannot replace that. Not because it is not clever enough. Because it physically cannot do the thing.

SECTION SIX

The only metric that ultimately matters.

Every argument in this paper has been chipping away at the same insight. Time poverty means most training does not even land. The Forgetting Curve means most of what lands does not stick. The legacy LMS architecture cannot deliver the reinforcement biology requires. AI without amplification philosophy hollows out the trade rather than empowering it. Points programmes plateau because the wrong currency is being earned.

Every one of those arguments has been a way of saying the same thing. The metrics the industry has been reporting on do not correspond to the outcomes the industry actually cares about. The programmes look successful by their own measurements while the real commercial question, namely whether the training made the advisor sell more confidently and more often, remains unanswered.

Solving the glocal paradox along the way

Before reaching the metric itself, one operational tension is worth dispatching briefly. Every DMO and major supplier operating across multiple markets eventually runs into the same problem. Head office wants brand consistency. Markets want local relevance. The two pull against each other, and almost every programme that tries to solve the problem ends up compromising on both.

It is not a trade-off. It is an architecture problem with a structural answer.

The reason the contradiction feels unresolvable is that the industry has historically thought of a training programme as a single piece of content distributed across many markets. A modern engagement platform does not distribute one piece of content. It surfaces different content to different advisors based on who they are and where they sell. The Italian feed is genuinely Italian: the language, the visa updates that matter to Italian outbound clients, the supplier campaigns most active in that market, the FAM trip opportunities relevant to Italian advisors. The Brazilian feed is genuinely Brazilian, and shares almost nothing with it. But both feeds sit inside the same platform, under the same brand, governed by the same editorial standards. The consistency is structural. The relevance is operational. The contradiction dissolves.

Return on Education

Which brings us to the metric that decides whether any of this is worth doing.

Most current programmes report against the wrong things. Completion rates. Certified-agent counts. Course-satisfaction scores. The number of badges issued in a quarter. Engagement minutes per advisor per month. These are all measurements of activity at the point of training. None of them is a measurement of commercial outcome at the point of sale. And the gap between those two things is exactly the gap this paper has been arguing about.

> *Return on Education is the deliberate counterpart to Return on Investment. For every pound spent educating advisors, how much commercial outcome did that education produce?*

Operationally, Return on Education measurement requires three things most current programmes do not collect.

First, the correlation between training engagement and booking activity at the level of the individual advisor. Did the agents who completed the destination specialist programme actually sell more of that destination in the following quarter than the agents who did not? Most programmes never track this. The data exists, with the certification record sitting in the LMS and the booking record sitting in the supplier system, but the two are almost never joined.

Second, the retention curve for product knowledge over time. Six months after an advisor completed a certification, do they still recall the specific differentiators the programme was designed to embed? Almost no programme in the industry currently measures it.

Third, the confidence signal. When an advisor is on a call with a client, are they recommending the destination, supplier, or product the training was about, or are they recommending an easier sell because the training did not give them the recall they needed at the moment of truth? This is the hardest of the three to measure but the most important. Because confident recommendation at the point of sale is what every trade engagement programme exists to produce, even if no programme currently measures it.

CLOSING

The challenge this paper closes with.

If you have read this paper to its close, the question worth carrying out of it is the simplest one I can ask.

> *By the end of this decade, what is your trade engagement strategy actually being measured against?*

If the answer is completions, certifications, badges, course satisfaction, or engagement minutes, then none of what this paper has argued has landed yet. Those metrics describe activity. They are not measurements of outcome. They tell you the programme is running. They do not tell you the programme is working.

If the answer is Return on Education, meaning confident advisor selling, correlated to training, measured against commercial outcome at the point of sale, then the rest of this paper has been preamble. It has been the diagnostic work of showing why most programmes cannot currently answer that question. The diagnostic work is done. The decision is about whether the operators reading this are willing to redesign their measurement architecture around it.

Some will. Most will not. The ones who do will spend the next decade quietly out-performing the ones who do not, because they will be optimising for the right thing while the rest of the industry continues optimising for the wrong one.

About Hablo

Hablo is a Learning Experience Platform built for the travel trade. It is the working realisation of the architectural argument laid out in this paper. The platform is mobile-first, content-led, AI-augmented, and built around the principle that advisors are the protagonists of their own professional development, not the subjects of it.

Hablo's personalisation is rule-based today, tailoring what each advisor sees to their region, role, and specialisms, with AI recommending and nudging while leaving the advisor in control. The platform's behavioural and predictive personalisation, its full visibility-based motivation architecture, and the joined-up Return on Education measurement that connects training engagement to supplier booking systems are the platform's direction of travel: the roadmap, not the current feature set. I would rather be precise about what ships today than oversell.

Hablo is currently being deployed by Visit California as part of their next-generation global trade engagement strategy. Conversations about deployments with other destinations and suppliers are welcome through hablo.travel or directly with the author.

A note on what comes next

This paper is the consolidated form of the inaugural Travel has a training problem series. But the conversation it has tried to start is not finished. There is a deeper question that has been quietly building underneath the Return on Education argument: about how expert trade knowledge actually reaches the consumer who needs it, and how that connection itself can be the next chapter of the industry's commercial story. That is the territory of the next paper, which will be published alongside the second series of newsletters in due course.

REFERENCES

Light references

The data on the average employee having roughly 24 minutes a week for formal learning comes from research originally conducted by Bersin & Associates, since incorporated into Deloitte's learning research practice, and has been cited continuously by L&D research firms across the past decade as a corporate average rather than a travel-trade-specific figure.

The Forgetting Curve is the foundational research of Hermann Ebbinghaus, published in his 1885 monograph Über das Gedächtnis (Memory). The curve has been replicated continuously across the field of cognitive psychology for over 140 years and remains one of the most settled findings in memory research.

The Hook Model framework is drawn from Hooked: How to Build Habit-Forming Products by Nir Eyal (Portfolio/Penguin, 2014). The four-phase loop, comprising trigger, action, variable reward, and investment, is the author's framework, applied here to travel trade engagement.

Behavioural design principles referenced in Section 3, including the role of ability, motivation, and triggers, draw on the work of BJ Fogg, particularly the Fogg Behaviour Model developed at Stanford's Behaviour Design Lab.