


Decision Intelligence QUARTERLY

The board-grade strategic briefing that turns a monitored topic into decisions, with its evidence, its counter-arguments, its options and its own report card attached.




1 What it is

Decision Intelligence is the flagship strategic briefing of a Shaping Tomorrow continuous-intelligence programme, the deliverable that opens and closes each monitoring cycle on your topic. Where the weekly [Signal Scanner](#) detects and the monthly [Change Trackers](#) assess, Decision Intelligence decides: it converts a cycle's accumulated evidence into weighted themes, decision postures, scenarios set in the horizon year, options and board-ready questions, each traceable to a tiered, dated source. Every briefing now ships with two companion pages: a [System Map](#) that shows why its themes move together, and a [Scenario Stress Test](#) that grades its options in every future and tracks which future is gaining momentum.

The Shaping Tomorrow intelligence cycle

What changed this week, what is moving this month, what to decide this quarter, and one page that always shows where things stand. This deliverable's lane is highlighted.



Each briefing reads every scan and tracker in the quarter, grades the milestones the previous briefing set, and feeds the Action Triggers that close the cycle and the Live Summary that keeps the position current.

A recent edition, “Where Your Money Will Live in 2030”, the third in the Financial Services and Future of Money series, runs on 48 verified sources, five weighted themes each with an Outlook to 2030, four owner-assigned implications, a scenario matrix set in 2030, twelve options across four decision areas, three explicit exclusions, a claim-fidelity appendix and an abbreviations appendix. Before it makes any new calls, it grades the milestones its previous edition set.

THE PROBLEM IT SOLVES

Strategic reports fail in two directions: they are either evidence-dumps that leave the decision to the reader, or opinion pieces that hide their sourcing and never revisit their own calls. Decision Intelligence is engineered against both. Every recommendation carries a posture, an owner and a dated action, and every decision area comes with a range of options; every claim carries a tiered source; every argument travels with its counter-argument; and every edition grades the previous edition's homework.

NEW IN THE SEPTEMBER AND OCTOBER 2026 FORMAT

The briefing was restructured in September 2026 around what boards asked for: less to read before the verdict, a clearer view forward and a choice of responses rather than a single recommendation. The Executive Synthesis is now about two minutes long. Each theme is weighted for how much it shapes your decisions, with who decides and how much room you have to act, and each carries an Outlook to the horizon year built on sourced forecasts. The Scenario Matrix is set in that year, on two uncertainties that produce different kinds of world. A new Options section sets out three choices in each of three decision areas, from least change to biggest shift, with the trade-off of each and the scenarios it fits. October 2026 added a “Shift from today” line to every scenario, a Potential Surprises to 2040 section with signposts the public Track record grades, and a fixed closing question that turns the scenarios on the reader’s own current activities. The whole briefing is written in plain English, with every abbreviation explained in an appendix and a link to how it was made.

From 21 October 2026 every briefing ships with two companion pages, each on its own link and published with the briefing by one button. The [System Map](#), also drawn inside the briefing after the last theme and listed in its Contents, shows the feedback loops that connect the themes, which of them are running this quarter, and the leverage points where the four Implications act. The [Scenario Stress Test](#) grades every option in each scenario and against a wild card, scores twelve leading indicators on the matrix’s two axes at every reading to show which future is gaining momentum, and turns the reading into budget positions and signposts.

2 Why clients choose it

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
It grades its own homework	Every edition grades the milestones the previous edition set, with a fresh source for each verdict, before making new calls.	Accountability compounds into trust. Renewal conversations start from a track record, not a promise.
It argues against itself	Counter-arguments, stress tests, exclusions and a claim-fidelity appendix are structural, not optional.	Your board gets analysis that has already survived adversarial review, and can see exactly where to push back.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
It offers choices, not verdicts	Postures, owners and dated actions, then three options in each decision area with their trade-offs and the scenarios they fit.	The distance from reading the briefing to running the meeting is close to zero, and the meeting debates real alternatives.
It looks ahead to a stated year	Every theme carries an Outlook to the horizon year; the scenarios are set in that year; the options name the scenarios they fit.	You plan against a named year and named forecasts, rather than against a mood.
It is written for everyone at the table	A two-minute synthesis, plain English throughout and every abbreviation explained.	Board members outside the specialism can follow and challenge it, so the briefing works in the meeting, not only in the strategy team.
It shows the system, and watches the scenarios	A System Map of the loops behind the themes, and a Scenario Stress Test that grades every option in every future and scores which future is gaining momentum, month by month.	You see why the themes move together and where effort counts, and the scenarios become a monitoring instrument that sets budgets, not a workshop exercise that is filed.
It compounds	Series continuity, a continuity ledger, and weak signals promoted or retired each cycle.	Each edition is worth more than the last, because it carries the whole programme's memory. That is what a subscription buys that a one-off report cannot.

FROM A RECENT EDITION, "WHERE YOUR MONEY WILL LIVE IN 2030" (SEPTEMBER 2026)

In a two-minute synthesis and a 25-minute full read, the briefing: graded the five milestones its July edition set (three **on track**, two **unresolved**); weighted five themes, three of them **high**, and gave each an Outlook to 2030; issued four implications, one to **decide**, two to **prepare** and one to **monitor**; built four scenarios for 2030 on who issues money and what protects its users; set out twelve options across four decision areas, each naming the scenarios it fits; excluded three futures with reinstatement triggers; and put 48 verified sources, a claim-fidelity appendix and an abbreviations appendix behind all of it.

3 One document, two reading speeds

The briefing declares its own read estimate on the masthead, a 2-minute Executive Synthesis and a full read of around 25 minutes, and the architecture honours both. A director can read the plain-English “Why this matters now” abstract and the three-paragraph synthesis, see where the analysis could be wrong and which decisions cannot be deferred, and leave with the cycle’s verdict. A strategy team can descend through weighted themes, outlooks, scenarios, options and the source register and audit every claim to its origin. Nothing in the fast path contradicts the deep path, because both are generated from the same evidence base.

4 The design system: why the format itself is a benefit

Like every Shaping Tomorrow deliverable, the briefing is engineered to be skimmed, read or audited, and its design carries the method:

- **A consistent posture, tier and weight grammar:** Decide / Prepare / Monitor pills, Immediate / Near-Term / Longer-Range tiers and high / medium / low weights use the same colours here as in the Signal Scanner, Action Triggers and Live Summary, so the whole programme reads as one system.
- **Interactive where it earns it:** the Futures Wheel expands on click and shows its ratings on hover, and the contents cross-reference in both directions (theme → implication → theme), while the whole document remains print-clean for board packs.
- **Colour-coded epistemics:** counter-arguments, weak signals, outlooks, stress tests, options and decisions each carry a distinct colour, so you can see the confidence status of any block before reading a word of it.
- **Evidence at the point of claim:** links, dates and tier ratings sit inline and on the figures themselves, never banished to a bibliography the reader will not visit.
- **Plain language by design:** short sentences, institutions named in full, every abbreviation explained, so the briefing works for everyone at the table, not only the specialists.

5 Anatomy of the deliverable: element by element

The background detail for readers who want to see exactly what is in the briefing and why, section by section.

Click a section to expand it. Green tags mark what was added or rebuilt in September and October 2026, including the companion pages of 21 October.

The front door: masthead, continuity and the Futures Wheel

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Editorial title	A named thesis, “Where Your Money Will Live in 2030”, with its edition number and the titles of the previous briefings in the series.	Signals a maintained body of work, not a one-off PDF. The series-continuity line tells you your topic will be followed, not sampled.
Scope meta-line	Cycle date, named audience and an explicit horizon year, for example “Horizon: to 2030”.	Sets expectations honestly. You know the year every outlook, scenario and option is calibrated to, a discipline most foresight work skips.
Read estimate	“2 min Executive Synthesis · 25 min full read”, stated up front.	Respects executive time and doubles as a design promise: the document commits to being useful at both speeds.
How this was made	A line under the title, and a card at the foot of the briefing, linking to Shaping Tomorrow's six-stage method: Scan, Filter, Stress-test, Judge, Trigger and Actionable Intelligence.	A reader who meets the briefing on its own can see how it was produced before deciding how far to rely on it.
Interactive Futures Wheel	A clickable radial map of the first-, second- and third-order consequences of the lead theme, tinted by risk and opportunity, with line weight showing impact × likelihood, each node's rating on hover and a short note on how to read the figure.	The whole causal landscape in one visual. Executives explore “what follows from this?” themselves, a workshop tool embedded in a report, and a striking screen-share moment in any board meeting.
What you'll learn / Key takeaways	Twin cards of three bullets each: what the reader walks away knowing, and three dated, sourced findings the body of the briefing defends.	The contract with the reader, stated before the argument begins, and proof on page one that the answers are specific, not directional hand-waving.

The Executive Synthesis, the two-minute path REBUILT

The synthesis is not a summary bolted on top; it is a short, self-contained decision document with its own discipline, written so that a reader outside the sector can follow it.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
“Why this matters now”	A plain-English abstract of three short paragraphs, with no statistics and no jargon, ending with the one thing to take into the next meeting.	The one-minute on-ramp, added after readers found the full synthesis too dense as a first encounter. Anyone on the board can repeat the cycle’s central problem from memory.
The question on your desk	A single question about the decisions in front of the reader this week, answered in one line by the evidence.	Exposes the assumption the briefing is about to challenge before the argument begins.
Three paragraphs	The cycle’s central tension with the facts that prove it; what is settled and what is still open; and a third paragraph that begins “Looking ahead to 2030”, gives sourced forecasts and points to the options.	The interpretive lens in about two minutes, with at least five citations, so the verdict and its receipts travel together. There is no “trust us” paragraph anywhere in the synthesis.
“Where this analysis could be wrong”	A stress-test box naming the single most consequential assumption and the evidence that would force a revision, stated as a falsifiable claim.	The feature sceptical buyers remember. A briefing that names its own kill conditions earns the right to be believed on everything else.
The decisions that cannot be deferred	Three or four actions, each opening with a bold imperative and one sentence of justification, developed later in the implications and options.	A director who reads nothing else leaves with the actions that cannot wait for the next quarter.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Continuity ledger UPDATED	Item 3 in the Contents, straight after the synthesis: every milestone the previous edition set, each graded against a source dated after it (CONFIRMED, ON TRACK, UNRESOLVED, BROKEN or REOPENED). Each grade is copied to the public Track record and never edited afterwards.	Cycle-over-cycle accountability made visible, and now scoreable: nothing raised in a previous cycle is silently dropped, the old calls are graded before new ones are made, and the misses stay on the public record.

Audience Snapshots: four lenses, one evidence base

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Role cards	Four colour-coded cards, one per reader group (for example “If you run payments”), each stating in one cited sentence what is changing for that group.	One briefing serves four executive audiences without four documents. Each stakeholder finds their own stake in under a minute.
Question for your next meeting	One plain, decision-shaped question per card: “When an AI agent initiates a payment through us, what exactly has the customer authorised?”	Converts intelligence into agenda items. Your internal meetings are pre-drafted, the fastest route from reading to acting.
What it is for NEW	A line saying what answering the question gives the team: “Gives legal and product teams one test to apply to every payment an agent starts.”	Each team knows not only what to discuss but what the discussion is meant to settle.

Themes, the argument layer **WEIGHTS AND OUTLOOKS**

The cycle’s signals are organised into five or six themes (four when a client asks for four), ranked by when they bite and weighted by how much they shape your decisions to the horizon year. Each theme is carried by the same scaffold:

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Weight table NEW	One row per theme: its weight (high, medium or low), who decides and how much room you have to act, repeated as a line under each theme heading.	Shows where your effort counts. A heavy theme you cannot influence calls for preparation; one you can shape calls for a decision. The weights are stated as the analysis's judgement, so you can challenge them.
Impact-tier ranking	A colour-coded tier pill on every theme, defined in decision terms: Immediate, Near-Term or Longer-Range.	Triage built in. You know which themes bear on this year's decisions and which are multi-year watches.
Framing paragraph	The consequence for you first, then what the theme forces you to confront, in three to five plain sentences.	Analysis that holds its nuance in plain words is what separates board-grade from newsletter-grade. You can judge the quality from any single paragraph.
Data figures	At least two purpose-built charts across the briefing, of different kinds, each with its sources on the figure.	The evidence visualised for lifting straight into your own decks, sources attached.
Sourced signals	Three or four evidence items per theme, each leading with the finding, dated and linked.	Due-diligence-proof. An analyst can verify any item in one click; nothing rests on unattributed assertion.
Weak signals to watch	Early, low-confidence signals, each flagged with the specific development that would give it more weight.	Shows the detection net running ahead of consensus, and defines in advance exactly what evidence would promote a hunch to a finding.
Outlook to the horizon year NEW	Where the theme is heading by the horizon year, carried by sourced forecasts or dated commitments and reported as forecasts, followed by a "What it means" line.	Every theme looks forward as well as back, and the forecasts are named so you can check them yourself.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Counter-argument box	The most credible challenge to the theme's reading, argued properly and sourced.	The briefing argues against itself so your team doesn't have to. Boards trust analysis that has already survived its own opposition.
On the System Map NEW	A line at the end of every theme naming the loop it sits on, the loop's status this cycle and the leverage point on it, linked to the System Map after the last theme.	Each theme is read as part of the system, not as a separate pressure.
Decision link	A cross-reference from every theme to the implication it feeds.	No orphaned analysis. Every theme exists because a decision depends on it, and the link proves it.

Strategic Implications, the decision layer

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Four implications	Each titled as an action, for example "Implication 1: Decide how you will keep deposits: with rewards, with tokens, or both", with the reasoning in plain prose: who should do what, by when, and why.	The recommendation reads as a decision to take, not a finding to interpret.
Posture pill	Every implication is graded Decide, Prepare or Monitor.	Separates "act now" from "get ready" from "watch", a shared triage vocabulary across the whole programme.
Dated action and owner	One concrete action line with a deadline and an accountable role.	The recommendation arrives pre-routed and calendar-ready. Your cost of acting approaches zero.
Draws on	Links back to the themes that carry the evidence.	Every recommendation can be traced to its evidence in one click.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
On the System Map NEW	The leverage point each implication acts on, and what it does to which loop.	The actions are where they are because that is where the system can be moved, and the map shows why.

Scenario Matrix, set in the horizon year REBUILT

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Named axes	A 2×2 built on two uncertainties whose answers produce different kinds of world by the horizon year (the matrix is headed with it: “Scenario Matrix 2030”), for example “Money issued by: banks → non-banks” and “Protection: firm terms → public rules”.	Scenarios that differ in kind test a strategy; scenarios that differ only in speed or degree test nothing but its timing. Trends that advance in every future are shown inside each scenario rather than spent on an axis.
Four named worlds	Memorable environments written as the world in the horizon year: The upgraded bank, Regulated new money, Walled gardens, Buyer beware.	Names that stick become your internal shorthand. Six months later the team is still asking “are we drifting toward Walled gardens?”
Shift from today NEW	Each scenario opens with one sentence, in the cell’s own colour, naming what would have moved from today’s position for that world to be true, anchored on a figure the reader can check now: “Europe’s average wait for new medicines, 607 days today, lengthens rather than shortens”.	Every scenario becomes a measurable distance from now rather than a picture floating in 2030, so the next edition can say which world you are drifting toward.
Scenario Stress Test NEW	A line under the matrix linking the companion stress test, which scores twelve leading indicators on the two axes each month and says which future, if any, now leads.	Ongoing scenario monitoring with the scoring in the open: you can see which future is gaining momentum, and challenge exactly how it was scored.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Early indicators	Every cell lists two or three observable events that would show the world moving into it.	Converts scenario planning from a workshop exercise into a monitoring instrument the programme then tracks cycle by cycle.

Options, the choice layer NEW

Options sit after the scenarios because every option names the scenarios it fits. They are the analysis's proposals for discussion, and the claim-fidelity appendix says so.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Three decision areas	The three places in the evidence where readers hold the biggest levers, each titled with the area and who holds it, for example "Options: AI agents that pay (any firm that sells online)".	You can see at a glance which choices are yours to make and which belong to someone else.
Three options per area	Option A · Least change, Option B · Middle course and Option C · Biggest shift, each stated in one plain sentence.	Turns a recommendation into a choice. The board sees the range from staying put to changing course, not a single answer to accept or reject.
Trade-off	What each option costs now and what it leaves exposed later.	The price of every option is on the card, so the discussion starts from trade-offs rather than preferences.
Graded in every future NEW	A line opening the Options linking the stress test, where each option is graded in all four futures and a wild card and marked Robust, Bet or Fragile, with a budget position and its release trigger.	The board knows which choices are safe in any future and which should wait for a signpost, so the scenarios inform investment priorities directly.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Fits scenarios	The scenarios in which each option holds up, including “All four” where it is robust.	Links each choice to the futures it is meant for: a robust option can be taken now; one that fits a single scenario can wait for that scenario’s early indicators.

Potential Surprises to 2040, the far-horizon layer NEW

Everything above is calibrated to the horizon year. This section, added in October 2026, looks past it: two or three developments to 2040 that few in the sector yet expect, each already visible in an early, sourced sign. Widely discussed futures do not qualify; the test is whether the card makes a well-read executive stop.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Two or three surprise cards	Each labelled Weak signal or Emerging, with a dated window ending by 2040 and a probability band carrying its percentage, such as “Low-Medium (20 to 40%)”.	The far horizon arrives as scorable claims, not scenario wallpaper. A band with a percentage can later be judged right or wrong.
“Why it would surprise”	One line naming the assumption the development would break: “Medicines are licensed to treat named diseases, and ageing itself is not one of them.”	Shows exactly which of your planning assumptions is at stake, so the card can be routed to whoever owns that assumption.
Signposts within three years	Two or three dated, observable events that would show the surprise forming, each recorded as a call in the public Track record .	You never have to watch 2040; you watch the next three years, and we grade our own signposts in public as they fall due.
“What it touches today”	One line naming the decisions already being made on the assumption the surprise breaks: “Drug pipelines, pension funds and life insurers set their assumptions now on the basis that ageing cannot be treated.”	Connects the far horizon to this year’s budget lines, which is what earns the section its place in a long briefing.

Companion pages: the System Map and the Scenario Stress Test NEW

Every briefing from 21 October 2026 ships with two companion pages, built from the briefing's own evidence, each on its own link. One "Publish all" step puts the briefing, the map and the stress test online together, briefing first.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
System Map section	Inside the briefing, after the last theme and listed in the Contents: the map of the loops that connect the themes, a table of each loop and its status this cycle (Turning, Priming, Weakening or Stalled), a "What it means" line and the evidence for every link.	The connections between the themes are drawn where the argument needs them, and every arrow can be checked in place.
The System Map page	The full interactive map with a card for each loop, what the map shows that the briefing alone cannot, the leverage points with their postures, how the loops explain the four scenarios, and every link with its source and tier. Product sheet .	A shared picture for the leadership meeting, built to be argued over in a 90-minute challenge session and kept live by the monthly tracker.
The Scenario Stress Test page	The nine options graded in the four futures and a wild card to 2040; the trail of readings on the matrix with the twelve indicators behind it; budget positions with release triggers; signposts and the decisions they trigger. Product sheet .	Scenario development and ongoing scenario monitoring in one place: which moves hold in every future, which future is gaining momentum, and where the money should go.
Carried forward	Each new briefing carries every earlier reading forward with the score it was given at the time, and the Change Tracker's signals feed the indicators in between.	After a year you have a documented record of which future the evidence favoured, and when.

What We Are Not Planning For, the exclusions layer

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Three explicit exclusions	Futures deliberately held out of the plan, each with the reasoning stated.	Saying “no, and here is why” is rarer and more valuable than another risk. It protects you from resourcing against phantoms.
Reinstatement triggers	Each exclusion carries the named event that would put it back on the table.	The exclusions are falsifiable too. Nothing is dismissed forever; everything dismissed is monitored for re-entry.

Questions, Source Register, Claim Fidelity and Abbreviations

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Questions For The Next Strategy Meeting	Five questions, each forcing a trade-off, surfacing risk posture or challenging an assumption, at least one drawn from a weak signal or a non-consensus reading.	The last mile of intelligence: the agenda for your next strategy session, written for you. Chairs lift these questions verbatim.
The fixed fifth question NEW	Every edition ends its questions the same way: “Which of our current activities looks least wise in these futures?”	Scenarios are most useful as a test of today’s commitments. A leadership team that only adds options never takes anything off its list; this question asks for the subtraction, every quarter.
Source Confidence Register	All sources in per-theme tables (source, tier, date, key claim) under a four-tier taxonomy, with the source count and tier mix stated.	Radical sourcing transparency. A procurement or research team can audit the entire evidence base in minutes, and compare it to the unsourced decks it replaces.
Recency discipline	A declared recency window, older sources named as structural anchors, and undated pages recorded as undated rather than given an inferred date.	Method hygiene that survives scrutiny: you know exactly how fresh the evidence is and how older anchors are being used.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Sources in their own language	Sources in other languages read in the original and reported in plain English.	The evidence is not limited to what happens to be published in English, and nothing is lost in a reader's own translation.
Conflict notes	Genuine source disagreements named, with the side the briefing weighted, and an honest account of where the source set is thin.	Most reports hide their contradictions; this one curates them. Confessing the evidence base's own gaps is the deepest trust signal in the document.
Claim-Fidelity Appendix	Every analyst inference and editorial framing separated from what sources assert, including the theme weights, scenario axes and options, which are marked as proposals for discussion, and the forecasts behind each Outlook.	You always know whether you are standing on evidence or interpretation, and can disagree productively. Effectively unique in commercial foresight.
Abbreviations NEW	Every abbreviation and unit in the briefing, in alphabetical order, each with a plain explanation rather than only the expansion.	A reader outside the sector never has to stop and search, so the briefing can travel beyond the specialist team.

Decision Intelligence briefings open and close each monitoring cycle of a Single-Topic or Multi-Topic continuous-intelligence programme, working in concert with the weekly [Signal Scanner](#) and monthly [Change Trackers](#) that feed them and the [Action Triggers](#) synthesis that closes the cycle. Each briefing ships with its [System Map](#) and [Scenario Stress Test](#), published with it. Their findings also feed the programme's Live Summary, the one-page view that is refreshed after every deliverable. Each edition is delivered as an interactive, print-ready briefing tailored to your topics, horizon and executive owners.

[Discuss a Tailored Programme](#)

[Request a Sample](#)