THE FUTURE-READY CFO MASTERCLASS

MODULE 17 OF 20 | PILLAR V: FUTURE MARKETS, OPERATING EXCELLENCE AND REINVENTION

Strategic Foresight and Scenario Planning

MODULE SUMMARY

This module replaces the single point forecast with disciplined thinking about several possible futures. It covers probabilistic reasoning, with a Python Monte Carlo simulation of revenue and margin, calibration with a Brier score, tested leading indicators, and stress tests including a reverse stress test. It then teaches scenario planning on a two by two grid, a minimax regret decision table, long range planning, regime shifts in artificial intelligence, tariffs, rates and demand, and habits that keep foresight alive, linked to Modules 5, 12, 18 and 20. A fictional capacity decision, a one page scenario sheet and a self assessment follow.

MODULE AT A GLANCE

Core question	How does a CFO plan, invest and commit capital when no single forecast can be trusted, so that the company survives the futures it did not predict and profits from some of them?
Why it matters to the CFO	Boards and lenders still receive one number, yet the surprises of recent years came from tariffs, rates, demand and technology. A CFO who shows the range, the signposts and the prepared response protects both the balance sheet and credibility.
Key sources	McKinsey Quarterly (2000) on four levels of uncertainty; Shell scenario practice (Davis, 2002); Mellers and colleagues (2017) on forecasting skill; Federal Reserve 2026 stress test scenarios; Deloitte Finance Trends 2026 (2025); EY (2025); Deloitte on tariffs (2026).
Executive artifact	A scenario planning one pager with signposts and trigger actions, and a stress test template, both completed for a fictional company.

Important notice. This material is provided for educational and informational purposes only and does not constitute accounting, tax, legal, investment, or other professional advice. Readers should consult qualified professionals in accounting, tax, legal, and other relevant services before acting on any information presented. Fictional companies and figures are used for teaching and are labeled as such. The maturity ratings are an educational framework and have not been validated against a broad population of finance executives.

1. Executive overview and learning objectives

A note on Pillar V. The first four pillars of this program treated technology, capital, governance and leadership, which are the capabilities that a CFO brings to the company as it exists today. Pillar V looks outward and forward. Its four modules deal with the futures that the company must meet and with the operating discipline that lets it meet them: strategic foresight and scenario planning in this module, global business, tax and geoeconomic strategy in Module 18, operational excellence and performance architecture in Module 19, and the reinvention of the CFO role itself in Module 20. This module comes first because the other three depend on it. A decision to enter a new country, to redesign a performance system or to reinvent a career is a decision about a future that nobody can see, and the habits taught here are the habits that make such decisions sound.

The core question. How does a chief financial officer plan, invest and commit capital when no single forecast can be trusted, so that the company survives the futures that it did not predict and profits from some of them? The question is not an invitation to abandon forecasting. A company must still budget, report to lenders and tell investors what it expects. The question asks instead that the single number be treated as one point inside a range of outcomes, that the range be described honestly, and that the company decide in advance what it will do as events reveal which part of the range it has entered.

Why this matters over the next five years. The evidence in this section is advisory and survey based, and it describes expectation and not proof. Deloitte's Finance Trends 2026 report, published in October 2025 and based on a survey of 1,326 respondents, ranked economic uncertainty first among the risks named by finance leaders, at 26 percent, with geopolitical tension and supply chain disruption each at 21 percent. The same report found that 30 percent of the finance leaders surveyed plan to strengthen advanced scenario planning. An EY roundtable report of 2025 recorded that 55 percent of the Fortune 250 CFOs present said they had become more bearish about the rest of the year and about 2026, after only 25 percent had been bearish following the election, and it recorded an EY geostrategy leader urging CFOs to plan around four trade scenarios. The roundtable sample was small and the page did not state the number of respondents, so the figures should be read as a mood and not as a measurement. The practical lesson is that the next five years are likely to bring further shifts in trade policy, interest rates, customer demand and technology, and that a CFO who plans on one path will be corrected by events at the cost of the budget, the covenants and the credibility of finance.

What this module is and is not. It teaches a set of disciplines that fit together. Probabilistic reasoning describes the range of outcomes. Calibration tests whether the forecaster's confidence deserves belief. Leading indicators give early evidence of which future is arriving. Stress tests show where the company would break. Scenario planning builds a small number of consistent stories about different futures. A decision table then compares strategies across those futures, and long range planning and organizational habits carry the whole discipline into the budget and the board cycle. Nothing here predicts the future, and nothing here is a substitute for the judgment of the company's own management, its auditors, its lenders and its legal and tax advisers.

Learning objectives. After completing this module the reader should be able to do six things. **First**, explain why a single point forecast misleads, and describe a forecast as a distribution with a base rate, an expected value and a tail. **Second**, keep a forecast log, compute a Brier score and read a calibration table. **Third**, choose, test and monitor leading indicators, and distinguish a signal that earns its place from one that produces false alarms. **Fourth**, build a stress test, single factor and combined, and run a reverse stress test that finds the point of failure. **Fifth**, carry out the scenario planning method from the focal question to the signposts, and compare strategies across scenarios with a regret table. **Sixth**, link scenarios to capital allocation (Module 5), crisis response (Module 12), global strategy (Module 18) and the future of the role

(Module 20), and install the habits that keep foresight alive in the board cycle.

2. Conceptual foundations and frameworks

Four tools that are often confused. A **forecast** is the company's best estimate of what will happen. A **projection** or **plan** states what will happen if stated assumptions hold. A **stress test** asks what happens to the company if a specified bad event occurs, without any claim that the event is likely. A **scenario** is a plausible and internally consistent story of how the world might develop, built to explore what the company would need to do and not to predict. Ged Davis of Shell, in a paper of July 2002, wrote that scenarios are not projections, predictions or preferences, and described them as consistent and credible stories about different paths to alternative futures. Exhibit 1 sets the four tools side by side, with the question each answers and the usual misuse.

Exhibit 1 Forecasts, Probabilistic Forecasts, Stress Tests, Scenarios and Decision Tables Compared

Tool	The question it answers	Typical output	Strength	Usual misuse
Forecast	What do we expect to happen?	One number or a short range for each period	Simple and needed for budgets, covenants and guidance	Treated as a promise, or as the only future worth planning for
Probabilistic forecast	How likely is each outcome, and how wide is the range?	A distribution, percentiles and the chance of crossing a threshold	Shows the tail and the cost of being wrong	Invented precision, with probabilities that no one has tested
Stress test	Where would we break, and by how much?	Covenant, liquidity and coverage readings under a set shock	Protects the balance sheet and requires no prediction	Shocks chosen to be mild, or never reviewed by the board
Scenario	What would we do in each of several different worlds?	Three or four narratives with financial effects and signposts	Challenges assumptions that everyone shares	Used to predict, or reduced to a best case, base case and worst case
Decision table	Which strategy holds up across the scenarios?	Payoffs, expected value and regret by strategy	Makes the trade between return and robustness visible	Probabilities and payoffs presented as facts

Source: author's synthesis, stated from the author's own knowledge. The description of scenarios follows Ged Davis of Shell Global Business Environment, Scenarios as a Tool for the 21st Century (July 12, 2002), listed in Section 7.

Four levels of uncertainty. Not every uncertainty calls for the same method. Hugh Courtney, Jane Kirkland and Patrick Viguerie, in an article of the McKinsey Quarterly dated June 1, 2000 and adapted from a 1997 Harvard Business Review piece, argued that traditional strategy assumes that executives can predict the future well enough to build a single discounted cash flow model, and that this often understates uncertainty and produces strategies that miss both threats and opportunities. They proposed matching the method to the level of residual uncertainty that remains after the best analysis has been done. At the first level the future is clear enough, the remaining uncertainty does not change the decision, and one forecast will serve. At the second level the outcome is one of a few discrete futures, such as the passage or failure of a pending regulation, and the analysis cannot say which. At the third level the key variables define a continuous range with no natural scenarios, such as the penetration of a new market. At the fourth level, true ambiguity, even the relevant variables are hard to identify, and the authors note that this level is rare and tends to move toward the second or third over time. Exhibit 2 matches each level to a method that this module teaches, and adds a

column on the posture that the authors describe, which is to shape the industry, to adapt to it or to reserve the right to play by making small investments that preserve options.

Exhibit 2 Four Levels of Residual Uncertainty and the Method That Fits Each

Level	What remains unknown	Method of this module	Strategic posture	Illustration (fictional)
1. Clear enough future	Nothing that changes the decision	Single forecast with a sensitivity table	Adapt	Renewal of a multiyear supply contract at a fixed price
2. Alternative futures	Which one of a few discrete outcomes occurs	Scenarios with probabilities, a decision table and trigger events	Shape, adapt or reserve the right to play	A trade rule that will either be imposed or withdrawn
3. Range of futures	Where a continuous variable lands within a range	Monte Carlo simulation, and three or four scenarios that span the range	Adapt, or reserve the right to play	Demand for a product in a market that the company has not served
4. True ambiguity	Even the relevant variables	Leading indicators, analogies and a list of what must be believed	Reserve the right to play	A technology that may or may not change how customers buy

Source: the four levels and the three postures are from Courtney, Kirkland and Viguerie, Strategy under uncertainty, McKinsey Quarterly (June 1, 2000), as returned by the page that was read. The method column and the illustrations are the author's own, and the illustrations are fictional.

The vocabulary of probability, in plain words. A forecast is better described by a **distribution**, which is the set of possible outcomes with the chance of each, than by a single figure. The **mean** or **expected value** is the average outcome weighted by probability. The **median** is the middle outcome, with half of the chances above and half below. A **percentile** states the value below which a given share of outcomes falls, so that the fifth percentile is a bad outcome that occurs one time in twenty. The **base rate** is how often something happens among comparable cases, and it is the right starting point before the specifics of this case are considered. **Tail risk** is the chance of an outcome far from the middle, and its size matters as much as its chance. The **flaw of averages** is the error of planning on the average input and assuming that the result is the average output, which is false whenever the relationship between input and result is not a straight line, as it is not when fixed costs, covenants or tax losses are involved. Exhibit 3 gives each term with its meaning for a CFO.

Exhibit 3 The Vocabulary of Probabilistic Reasoning for a CFO

Term	Plain meaning	What it tells a CFO	Common error
Distribution	All the outcomes that could occur, with the chance of each	How wide the range is and how it leans	Describing only the middle
Expected value	The probability weighted average of the outcomes	A fair single measure for repeated decisions and for ranking options	Believing that the expected value is the likeliest outcome, or that it will occur
Median and percentile	The middle outcome, and the value below which a stated share of outcomes falls	The range that covers most outcomes, such as the fifth to the ninety fifth percentile	Quoting one percentile as though it were the range

Term	Plain meaning	What it tells a CFO	Common error
Base rate	How often an event occurs among comparable cases	The anchor against which a plan should be judged before its special features are admired	Neglecting the base rate because this case seems different
Tail risk	The chance and size of an outcome far from the middle	What would threaten the covenants, the liquidity or the survival of the company	Treating a rare event as an impossible one
Correlation	The tendency of two variables to move together	Why shocks combine, so that revenue, margin and credit conditions fall together	Adding independent risks when they are linked
Flaw of averages	Planning on average inputs and expecting the average result	Why a plan built on average assumptions is usually worse than its average suggests	Computing the plan once, on the average case

Source: author's synthesis of standard definitions, stated from the author's own knowledge and not from a source read for this module.

Why forecasters err. Research on judgment under uncertainty finds that forecasters err in predictable ways, and a discipline that names the errors can reduce them. The study of Barbara Mellers, Joshua Baker, Eva Chen, David Mandel and Philip Tetlock, published in Judgment and Decision Making in July 2017, reports on the so called superforecasters, the top two percent of participants in government sponsored forecasting tournaments, who became more accurate in elite teams and resisted regression to the mean, which the work attributes at least in part to skill and not luck. The page that was read lists the factors that are thought to drive their success: actively open minded thinking, motivation shown by the number of questions attempted and the frequency of updating, scope sensitivity, nuanced probability judgments, and team environments with information sharing, discussion and post mortems. It also states that expert predictions are better when the setting gives regular, clear feedback and repeated chances to learn. The reading for a CFO is practical. Keep a record, update often, state probabilities in finer steps than likely and unlikely, and review the record as a team. Exhibit 4 lists the common errors with a counter for each.

Exhibit 4 Common Forecasting Errors, Their Appearance in Finance and a Counter for Each

Error	How it appears in finance	Counter
Anchoring	The budget starts from last year or from the first number said in the meeting, and later estimates stay near it	Build the range from the base rate and from outside views before looking at the plan
Overconfidence	Ranges that are too narrow, and forecasts that miss the range far more often than the stated chance	Keep a forecast log, compute a Brier score and widen ranges until the misses match the stated chances
Planning fallacy	Projects finish late and over budget because the plan assumes that nothing goes wrong	Reference class forecasting: compare with the record of similar projects
Base rate neglect	A new product is forecast to succeed because of its features, without regard to how often such products succeed	State the base rate first and justify each step away from it
Recency and narrative	The last quarter and a persuasive story dominate, and a single vivid event reshapes the forecast	Use several scenarios and require evidence for each story
Groupthink	Everyone in the room shares the same view of the world, so the same blind spot	Appoint a red team and ask each person to write a view before the discussion

Source: author's synthesis. The list of factors in superforecaster performance is from Mellers and colleagues (2017) as summarized on the page read. The mapping to finance and the counters are the author's own and are not drawn from that study.

Leading indicators and how to choose them. A **lagging indicator**, such as reported revenue, tells the company what has already happened. A **leading indicator** moves before the result that the company cares about, and so gives time to act. Good indicators share five properties. The indicator must have a plausible causal link to the result and not merely a past correlation. It must lead by enough time to allow a response. It must be measured reliably and often. It must have a threshold that triggers a named action. It must have an owner. Exhibit 5 turns these into a test that each candidate must pass. Section 3 shows a back test of four candidates for a fictional company, in which the broadest indicator signals most often and is wrong most often.

Exhibit 5 A Five Part Test for a Leading Indicator

Test	Question to ask	Evidence required	Reject if
Causal link	Is there a mechanism by which the indicator moves the result?	A short written explanation, supported by data	The link is only a coincidence in past data
Lead time	How many weeks or months does it move before the result?	Back test over several periods	The lead is shorter than the time needed to respond
Hit rate and false alarms	How many past misses did it flag, and how often did it flag a miss that did not occur?	A table of signals against outcomes	It signals nearly always, or nearly never
Measurability	Can it be measured on a fixed definition, at least monthly?	A data source and a definition	The definition changes, or the data arrive after the decision date
Action link	What will we do when the threshold is crossed, and who decides?	A trigger action and an owner	No one can say what the signal would change

Source: author's synthesis, stated from the author's own knowledge. The case for tracking leading indicators where scenarios cannot be formed follows Courtney, Kirkland and Viguerie (2000), as returned by the page read.

The scenario planning method. The method that Shell has long practiced, as Davis described it in 2002, begins by identifying the frames of reference and dominant mental models in the organization that matter for the decisions at hand, surfaces the assumptions behind current plans, identifies the key drivers of change over the relevant horizon, and separates predetermined elements, which are nearly certain, from the fundamental uncertainties. The planner then maps the directions that the uncertainties could take, selects those most useful for shaping present decisions, and builds scenarios as alternative contexts for testing strategies. Davis states that the process matters as much as the stories, since it widens perspectives and creates a common language among the people who must decide, and that scenarios help to stress test strategy and challenge planning assumptions. Courtney and his co-authors suggest keeping to no more than four or five scenarios at the third level of uncertainty. Exhibit 6 sets out a seven step method in the form that this module uses, with an output for each step.

Exhibit 6 A Seven Step Scenario Planning Method

Step	Task	Output	Common error
1. Focal question	State the decision that the scenarios must inform, and its horizon	One sentence, such as whether to add capacity within three years	A question so broad that every scenario seems relevant
2. Forces and drivers	List the external forces that bear on the question, and mark each as predetermined or uncertain	A short list of drivers, with the near certain ones set aside	Mixing company decisions with outside forces

Step	Task	Output	Common error
3. Critical uncertainties	Pick the two uncertainties that matter most to the question and are least predictable	Two axes, each with two opposite end states	Choosing axes that are correlated, so that one is redundant
4. Scenario grid	Combine the end states into a two by two grid of four futures, and name each	Four named scenarios	Naming scenarios as good, bad, and average
5. Narratives and numbers	Write the story of each, in a page, and translate it into revenue, margin, cash and covenant effects	A short narrative and a table of financial effects for each	Stories without numbers, or numbers without stories
6. Implications and options	Test strategies against each scenario, and find the moves that work in most of them	A decision table with expected value and regret	Optimizing for the likeliest scenario alone
7. Signposts and triggers	Choose observable signs that show which scenario is arriving, and agree the action for each	A list of signposts with thresholds, owners and trigger actions	Signposts that no one watches

Source: the sequence of identifying frames, drivers, predetermined elements and critical uncertainties follows Davis (2002) as returned by the page read. The seven step form, the outputs and the common errors are the author's own adaptation.

Options, hedges and no regret moves. A strategy under uncertainty need not be a bet on one scenario. A **no regret move** pays off in every scenario, such as lengthening the maturity of debt when the cost is small or building the data needed to monitor a signpost. A **hedge** gives up some return in the likeliest scenario in return for protection in a bad one, such as a swap that fixes the rate on floating debt. An **option** is a small commitment now that preserves the right, though not the duty, to commit more later, as when a company buys land or pays a deposit on long lead equipment before deciding whether to build. A **big bet** commits the company to one future and pays off greatly if it arrives. Courtney and his co-authors name three postures, which are shaping, adapting and reserving the right to play, and the last of the three is the posture of options. Module 5 develops the valuation of real options, and Exhibit 7 shows how the four kinds of move compare.

Exhibit 7 No Regret Moves, Hedges, Options and Big Bets

Kind of move	What it does	Cost	When it fits	Illustration (fictional)
No regret move	Improves the position in every scenario	Low, and recovered in all futures	Always, and first	Extend the maturity of the revolving facility; build the data for signposts
Hedge	Reduces loss in a bad scenario at a price in the others	A premium or forgone gain	When a loss would threaten covenants or liquidity	Swap part of the floating rate debt to a fixed rate
Option	Preserves the right to commit more later	A small outlay, which may be lost	When waiting will reveal which scenario is arriving	Pay for permits and long lead equipment deposits for a plant, without building it
Big bet	Commits the company to one future	The full investment, with the loss if the future fails to arrive	When the company can shape the outcome, or the future is clear enough	Build the full plant now

Source: author's synthesis. The three postures of shaping, adapting and reserving the right to play are from Courtney, Kirkland and Viguerie (2000). The four kinds of move and the illustrations are the author's own and are fictional.

3. Practical applications and tools

A running example. The exhibits of this section use one fictional company, Aldermoor Industrial Supply, a distributor and light manufacturer of industrial components with revenue of $500 million in the year just ended, a plan for the coming year of $530 million of revenue at a 15.0 percent EBITDA margin, which is $79.5 million of earnings before interest, taxes, depreciation and amortization, a $250 million term loan, a $100 million undrawn revolving credit facility and $40 million of cash. All figures are in millions of dollars, all are fictional, and every sum, probability and count was computed in Python and checked by assertion. The assumptions are the author's and are chosen to teach, and a real company would calibrate them to its own history.

Step one: replace the point forecast with a distribution. The plan of Aldermoor rests on revenue growth of 6.0 percent and a margin of 15.0 percent. A **Monte Carlo simulation** is a method that draws many random combinations of the uncertain inputs from stated distributions and records the result of each, so that the outcome is described by a distribution and not a figure. For Aldermoor, 20,000 trials were run with a fixed random seed of 17, so that the result can be reproduced exactly. In 85 percent of the trials the company is in a normal year, in which growth is drawn from a normal distribution with an average of 7.5 percent and a standard deviation of 4.0 percent, and the margin moves with growth, rising by 0.25 point for every point of growth above the average, because fixed costs are spread over more revenue, with a further random variation of 1.2 points. In 15 percent of the trials the company is in a demand shock, in which growth averages minus 6.0 percent with a standard deviation of 3.0 percent and the margin averages 11.5 percent with a standard deviation of 1.2 points. The 15 percent chance of a shock is an assumption of the author and is the weakest input in the model, which is itself a lesson: the answer is only as good as the stated beliefs, and the beliefs should be argued over in the open. Exhibit 8 gives the results.

Exhibit 8 Monte Carlo Result for Aldermoor: Revenue, Margin and EBITDA for the Coming Year (Fictional, Millions of Dollars)

Measure	Revenue	EBITDA margin (percent)	EBITDA
Point forecast (the plan)	530.0	15.0	79.5
Mean of 20,000 trials	527.4	14.5	76.8
Fifth percentile	463.3	10.9	51.2
25th percentile	514.1	13.2	68.4
Median	533.2	14.6	78.2
75th percentile	548.3	15.9	86.5
95th percentile	568.9	17.5	97.9

Source: author's calculation in Python, 20,000 trials with a random seed of 17. Inputs are the author's assumptions and are fictional.

What the distribution shows. The plan of $79.5 million of EBITDA sits at the 54th percentile of the simulation, which means that the plan is slightly more likely to be missed than met, and the mean of the trials is $76.8 million, which is $2.7 million or 3.4 percent below the plan. The gap arises because the bad outcomes are further from the plan than the good outcomes are, and because a shock lowers growth and margin together. The point forecast is also nearly always wrong in the narrow sense, since only 9.9 percent of trials fall within 2 percent of the plan and only 24.9 percent within 5 percent. The fifth percentile is $51.2 million, and the average of the worst 5 percent of trials, a measure often called conditional value at risk, is $47.4 million, which is 40 percent below the plan. Exhibit 9 puts the same results as questions a board might ask.

Exhibit 9 **Probabilities That a Board Can Use: Aldermoor EBITDA, Revenue and Margin (Fictional)**

Question	Trials (of 20,000)	Share of trials
EBITDA below the plan of 79.5	10,762	53.8
EBITDA within 2 percent of the plan	1,988	9.9
EBITDA within 5 percent of the plan	4,970	24.9
EBITDA below 66.0, an internal alert level that is assumed	4,255	21.3
EBITDA below 60.0, an assumed escalation level for the board	2,893	14.5
Revenue below the plan of 530.0	8,973	44.9
Margin below 14.0 percent	7,411	37.1
Revenue and margin both at or above the plan	7,205	36.0
Average EBITDA in the worst 5 percent of trials	1,000	47.4

Source: author's calculation in Python, same trials as the preceding exhibit. The last row shows the average EBITDA in millions of dollars and not a share. The alert levels are assumptions chosen for teaching.

The lesson of the simulation. The plan is not wrong. It is incomplete, because it shows no tail and no probability, and the incompleteness has a cost. Only 36.0 percent of trials meet both the revenue plan and the margin plan, so a CFO who promises both has promised something that this model gives one chance in three. A fairer statement to the board would be that the plan is a central case, that about one year in five produces EBITDA below $66 million, and that the company has agreed in advance what it will do if the early months fall in that range. The simulation can be rebuilt in a spreadsheet, and the choice of distributions should be reviewed by the people who know the business best. Section 7 explains the limit of this method: a simulation describes the uncertainty that the modeler has thought of and not the uncertainty that remains.

Step two: test the forecaster with a Brier score. A forecast of the form "the chance of this event is 80 percent" can be scored once the event has occurred or not. The **Brier score** is the average of the squared differences between the stated probability and the outcome, in which the outcome is 1 if the event occurred and 0 if it did not. A score of 0 is perfect, and a forecaster who always says 50 percent scores 0.25. The score punishes confident misses heavily, since a forecast of 90 percent for an event that does not occur costs 0.81. Aldermoor's finance team kept a fictional log of 20 forecasts of events such as whether a customer would renew, whether a shipment would clear by a date or whether a cost line would stay within budget. The Brier score of the log is 0.219, which is better than the 0.25 of a coin flip by 12.4 percent, and an unremarkable record. Exhibit 10 groups the forecasts by the confidence stated, and the pattern is the usual one: the team is well calibrated in the middle and overconfident at the top, since the three forecasts at 90 percent came true two times in three.

Exhibit 10 **Calibration Table for a Fictional Forecast Log of 20 Forecasts**

Stated confidence	Forecasts	Average stated chance (percent)	Actual frequency (percent)	Brier score in group
10 to 30 percent	6	20.0	16.7	0.180
40 to 60 percent	4	50.0	50.0	0.260
70 to 80 percent	7	75.7	71.4	0.204
90 percent	3	90.0	66.7	0.277

Stated confidence	Forecasts	Average stated chance (percent)	Actual frequency (percent)	Brier score in group
All 20 forecasts	20	56.0	50.0	0.219

Source: author's calculation in Python on a fictional log. The groups are 10, 20 and 30 percent; 40 and 60 percent; 70 and 80 percent; and 90 percent. A score of 0.250 is that of a forecaster who always states 50 percent, and of one who always states the base rate of 50 percent in this log.

How to keep a forecast log. The practice is simple and rarely followed. Write each forecast down before the event, with a date, a clear yes or no resolution rule and a probability in steps of five points. Score the log each quarter. Compare the Brier score with the 0.25 of the coin and with the score of a naive rule such as the base rate. Read the calibration table and widen the confidence statements where they are too high. The research summarized in Section 2 suggests that this feedback is the condition under which expert judgment improves, and it adds that teams that share information and hold post mortems do better. A CFO who keeps such a log, and shows it to the audit committee, makes the company's forecasts more credible and also forces honesty about how much is known. The log is also a defense against hindsight, which makes every outcome seem to have been foreseeable.

Step three: choose, test and monitor leading indicators. Aldermoor considered four candidate indicators of a revenue miss, defined as revenue more than 3 percent below plan in a quarter. Over twelve quarters there were three such misses, in the fourth, eighth and eleventh quarters. A signal given in a quarter is counted as a hit if a miss follows in the next quarter, and as a false alarm if it does not. Eight of the eleven possible signal quarters were followed by no miss. The back test in Exhibit 11 shows that the broad customer sentiment index flagged all three misses but also gave six false alarms in eight quiet quarters, which makes it nearly useless as a trigger. Quote volume gave the cleanest signal, with two of three misses flagged and one false alarm, while the backlog to billings ratio caught all three misses at the cost of two false alarms.

Exhibit 11 Back Test of Four Candidate Leading Indicators Over Twelve Quarters (Fictional)

Candidate indicator	Signals	Misses flagged (of 3)	False alarms (of 8)	Hit rate (percent)	False alarm rate (percent)	Share of signals followed by a miss (percent)
Backlog to billings ratio below 1.05	5	3	2	100	25.0	60
Distributor inventory above 10 weeks	5	2	3	67	37.5	40
Quote volume down 5 percent or more year on year	3	2	1	67	12.5	67
Broad customer sentiment index below its average	9	3	6	100	75.0	33

Source: author's calculation in Python on fictional data. Misses occurred in quarters 4, 8 and 11, so that eight of the eleven signal quarters preceded no miss. With only three misses the statistics are fragile, and a real back test needs more history and an out of sample check.

A dashboard that a board can read. Once the indicators are chosen, they are monitored on a one page dashboard with an amber line and a red line for each, an owner and a trigger action. Exhibit 12 shows the dashboard of Aldermoor at the end of the year. The status of each indicator was computed by rule in Python, with amber when the reading crosses the amber line and red when it crosses the red line. Three indicators are amber and one is red, which tells the board that no single alarm has sounded but that the backlog, the volume of quotes and the cost of imported inputs are each weakening and that the collection of receivables has

already crossed the red line. The trigger actions are stated in advance so that the discussion in the board meeting concerns the action and not the reading.

Exhibit 12 Leading Indicator Dashboard for Aldermoor at Year End (Fictional)

Indicator	Reading	Amber line	Red line	Status	Trigger action
Backlog to billings ratio	1.04	below 1.05	below 1.00	Amber	Review the order book weekly; hold phase two spending
Quote volume, change year on year (percent)	-6.0	below -5.0	below -15.0	Amber	Ask sales for a customer by customer review of lost quotes
Distributor inventory (weeks of sales)	9.4	above 10.0	above 13.0	Green	None; continue to monitor monthly
Effective duty on imported inputs (percent)	11.0	above 10.0	above 15.0	Amber	Open second source and domestic quotes for the three largest inputs
All in floating borrowing rate (percent)	7.0	above 7.5	above 8.5	Green	None; review the swap if the rate crosses 7.5 percent
Days sales outstanding	62	above 55	above 60	Red	Escalate the ten largest overdue accounts to the CFO; review credit terms

Source: author's construction, with status computed in Python by comparing each reading with its lines. The figures and thresholds are fictional. A negative quote volume reading is below the amber line of minus 5.0 percent and above the red line of minus 15.0 percent.

Step four: stress test the balance sheet. A stress test applies a stated shock to the plan and reads the result in the measures that matter to lenders and to survival. For Aldermoor the plan year has revenue of $530.0 million, EBITDA of $79.5 million, interest of $16.5 million, cash taxes of $10.8 million, capital expenditure of $25.0 million and free cash flow of $27.3 million. Opening net debt is $210.0 million, so that net debt at the end of the plan year would be $182.8 million and net leverage 2.30 times EBITDA, against a covenant maximum of 3.75 times. The model assumes that each dollar of lost revenue loses 40 cents of EBITDA, because the variable margin is high, that working capital is 20 percent of revenue and is released as revenue falls, that $150 million of the debt floats at 7.0 percent and the rest is fixed at 6.0 percent, and that tax is 25 percent of earnings after depreciation of $20 million and interest. The first four shocks are single factors. A fall of 10 percent in revenue lowers EBITDA to $58.3 million and raises leverage to 3.23 times. A tariff cost equal to 2.0 percent of revenue lowers EBITDA to $68.9 million. A rise of 200 basis points in floating rates costs $3.0 million of interest, and receivables that slow by ten days absorb $14.5 million of cash. None of the four breaches the covenant. The combined shock, in which all four occur together, lowers EBITDA to $48.8 million and lifts leverage to 4.32 times, which breaches the covenant. Exhibit 13 gives the full table.

Exhibit 13 Stress Test of the Aldermoor Plan Year: Single Factor and Combined Shocks (Fictional, Millions of Dollars)

Case	EBITDA	Free cash flow	Net debt	Net leverage (times)	Interest cover (times)	Liquidity with revolver	Cash only	Covenant
Plan	79.5	27.3	182.8	2.30	4.82	167.3	67.3	Met
Revenue down 10 percent	58.3	22.0	188.1	3.23	3.53	162.0	62.0	Met
Tariff cost of 2.0 percent of revenue	68.9	19.3	190.7	2.77	4.18	159.3	59.3	Met

Case	EBITDA	Free cash flow	Net debt	Net leverage (times)	Interest cover (times)	Liquidity with revolver	Cash only	Covenant
Floating rates up 200 basis points	79.5	25.0	185.0	2.33	4.08	165.0	65.0	Met
Receivables slow by 10 days	79.5	12.7	197.3	2.48	4.82	152.7	52.7	Met
Combined shock	48.8	(0.5)	210.5	4.32	2.50	139.5	39.5	Breach
Severe combined shock	24.4	(17.8)	227.8	9.35	1.16	122.2	22.2	Breach

Source: author's model, computed in Python from the assumptions in the text. Combined shock: revenue down 10 percent, tariff cost of 2.0 percent of revenue, floating rates up 200 basis points and receivables slower by 10 days. Severe: revenue down 20 percent, tariff cost of 3.0 percent, rates up 300 basis points and receivables slower by 15 days. Figures in parentheses are negative. Covenant is tested against a maximum of 3.75 times. Cash only is cash plus free cash flow, without the revolver.

Read the table for what binds. The surprise of the table is that liquidity is not the binding constraint. Even in the severe case Aldermoor holds $122.2 million of liquidity counting the revolver, well above the minimum policy of $60.0 million. The covenant is the constraint. In the combined shock the covenant is breached, and a breach can make the revolver unavailable, so that the usable liquidity is the cash alone, $39.5 million, which is below the policy minimum. A CFO who tested liquidity without the covenant would report comfort that the lenders could withdraw. The correct report shows both, and Module 12 treats the lender conversation and the playbook for such a quarter. Module 16 treats the negotiation of an amendment before a breach and not after it.

Step five: the reverse stress test. A stress test asks what a chosen shock would do. A **reverse stress test** starts from the failure, here a breach of the covenant, and asks what combination of events would cause it, which often reveals risks that nobody had chosen to test. Exhibit 14 reports four results for Aldermoor, each found by solving the model in Python. Revenue alone must fall 13.6 percent, to $458 million, with EBITDA at $50.7 million, before the covenant is breached, and the margin of safety on that single factor is wide. A tariff cost alone must reach 4.8 percent of revenue. But if the tariff cost of 2.0 percent, the rate rise of 200 basis points and the slowing of receivables by ten days have already occurred, a fall of only 6.5 percent in revenue is enough. The reverse test therefore shows the true message, which is that combinations of moderate events, which are far more common than one extreme event, bring the company to the covenant at about half of the single factor distance.

Exhibit 14 Reverse Stress Test for Aldermoor: What Would Have to Be True for the Covenant to Be Breached (Fictional)

Question asked in reverse	Break point	What it means for the CFO
How far can revenue fall alone before leverage passes 3.75 times?	13.6 percent (EBITDA of 50.7)	The company is robust to a pure demand shock of the size of a mild recession
How large a tariff cost, as a share of revenue, produces the same breach?	4.8 percent of revenue	A cost shock is more dangerous than a demand shock of equal size, since no cash is released from working capital
If a 2.0 percent tariff cost, a 200 basis point rate rise and slower receivables have already occurred, how far can revenue fall?	6.5 percent	Combined moderate shocks reach the covenant at less than half of the single factor distance, so lenders should be approached early

Question asked in reverse	Break point	What it means for the CFO
What is the lowest EBITDA that keeps leverage at or below 3.75 times at plan net debt of 182.8?	48.7	A floor of 38.7 below the plan, but net debt itself rises as EBITDA falls, so the true floor is higher

Source: author's calculation in Python, solving the model of the stress test by repeated halving of the interval until the leverage equals 3.75 times. Millions of dollars unless stated. The last row divides plan net debt by the covenant and does not allow for the cash that a fall in earnings would remove.

Where severity comes from. A common difficulty is the choice of how severe a shock should be. One public reference is the Federal Reserve's severely adverse scenario for its 2026 supervisory stress test of large banks, published in February 2026, in which the unemployment rate rises 5.5 percentage points from 4.5 percent to a peak of 10 percent, real gross domestic product falls 4.6 percent, house prices fall about 30 percent and equity prices fall about 58 percent. The Federal Reserve states that the scenario is a hypothetical for testing the capital of banks and that it is not a forecast. It was designed for banks and not for a distributor of industrial components, so a company should not copy it, but it shows how a severe case is built, as a connected path of several linked variables over several quarters and not as a single number. A company may set its own severe case by asking what its worst quarter and worst year on record looked like, what its customers' worst year looked like, and how each of its main variables would move in a recession such as that scenario describes.

Step six: build scenarios on a two by two grid. Scenario planning begins with a focal question. For Aldermoor the question is whether to commit $120 million to new manufacturing capacity within three years, and on what terms. The drivers that bear on the question include the trade regime, the strength of customer demand, interest rates, labor cost and technology. Two of the drivers are predetermined: the installed base of equipment that Aldermoor's customers own is aging, and labor cost is rising, so neither is treated as uncertain. The two **critical uncertainties** are the trade regime, which may be open or fragmented, and the strength of customer demand, which may be strong, driven in part by customers' investment in automation and artificial intelligence, or weak. These two were chosen because they matter most to the capacity decision and are the least predictable, and because they are not strongly related, so that all four combinations are plausible. Exhibit 15 shows the grid, and Exhibit 16 gives the financial effect of each scenario in the third year, with the probabilities that the management team assigned, which are judgments and not facts.

Exhibit 15 Scenario Grid for the Aldermoor Capacity Question: Trade Regime and Customer Demand (Fictional)

	Open trade	Fragmented trade
Strong demand	**Broad Tailwind** (20 percent). Customers invest heavily, inputs are available at low duty, competition is global	**Costly Boom** (30 percent). Customers invest heavily, but duties and supply disruption lift costs and favor domestic capacity
Weak demand	**Long Plateau** (35 percent). Customers defer investment, supply is ample, prices are under pressure	**Hard Squeeze** (15 percent). Customers cut spending while duties and borrowing costs press on margins

Source: author's construction for a fictional company, following the method of identifying two critical uncertainties and combining their end states described in Section 2. Probabilities are management judgments for teaching.

Exhibit 16 Financial Effect of the Four Scenarios in Year Three (Fictional, Millions of Dollars)

Scenario	Trade	Demand	Probability	Year three revenue	EBITDA margin (percent)	EBITDA
Broad Tailwind	Open	Strong	20 percent	620.0	17.0	105.4
Costly Boom	Fragmented	Strong	30 percent	585.0	13.0	76.1
Long Plateau	Open	Weak	35 percent	520.0	14.0	72.8
Hard Squeeze	Fragmented	Weak	15 percent	430.0	9.5	40.9
Probability weighted			100 percent	546.0	13.8	75.5
The plan of the budget cycle				590.0	15.0	88.5

Source: author's model, computed in Python. Year three revenue and margins are assumptions chosen for teaching. The probability weighted margin is the weighted EBITDA divided by the weighted revenue. The probability that EBITDA falls short of the plan is 80 percent, the combined probability of the three scenarios other than Broad Tailwind.

Reading the scenarios. The probability weighted EBITDA is $75.5 million, which is $13.0 million, or 14.7 percent, below the $88.5 million that the plan of the budget cycle assumes for the third year. Only in Broad Tailwind does the company reach that plan, and the management team gave that scenario a 20 percent chance, so the plan is a good outcome described as a central case. This is the most common discovery of scenario work. The variation among scenarios is also valuable in itself. EBITDA in Hard Squeeze, at $40.9 million, is less than half of EBITDA in Broad Tailwind, at $105.4 million, so that the capacity decision looks very different depending on the world that arrives. Exhibit 17 adds a short story, the implications and the signposts for each scenario. A good narrative explains why events unfold as they do, in a page and not in a paragraph, so that managers can recognize the scenario when its early events appear.

Exhibit 17 Narratives, Implications and Signposts for the Four Scenarios (Fictional)

Scenario	Story in brief	Implications for Aldermoor	Signposts to watch
Broad Tailwind	Customer plants add automation and software, order books fill and imported inputs stay cheap	Capacity is short within eighteen months, and lead times lose orders to rivals	Backlog to billings above 1.10 for two quarters; quote volume up more than 10 percent
Costly Boom	Demand is strong, but duties of 15 percent or more on key inputs and slow ports raise costs, and customers prefer a domestic supplier	Domestic capacity earns a premium, while margin falls unless prices rise	Effective duty on inputs above 15 percent for two quarters; domestic share of customer tenders rising
Long Plateau	Customers hold spending, prices are flat and the market is crowded	Surplus capacity hurts returns, and cash discipline and cost position matter most	Quote volume down year on year for two quarters; distributor inventory above 10 weeks
Hard Squeeze	Customers cut spending sharply while duties and borrowing costs remain high	The covenant becomes the binding constraint, and survival depends on liquidity and speed	Days sales outstanding above 60; floating rate above 8.5 percent; covenant headroom below 15 percent

Source: author's construction for a fictional company. The signposts connect to the indicators of the dashboard earlier in this section, so that the monitoring system and the scenarios are one system.

Step seven: compare strategies across the scenarios. The team then lists the strategies open to it and estimates the value that each would create in each scenario, expressed as net present value in millions of dollars. Four strategies were considered: to expand capacity now, to hold the present footprint, to expand in phases with stage gates, and to acquire a regional competitor. Exhibit 18 shows the payoffs, the probability weighted expected value and the worst case for each. The expected value criterion selects strategy A, expanding now, at $64.0 million, narrowly ahead of the phased strategy at $62.3 million. The maximin criterion, which chooses the strategy with the best worst case, selects B, holding the footprint, whose worst case is a gain of $10 million. The two criteria disagree, as they often do, because one weighs the average and the other fears the worst.

Exhibit 18 Payoffs by Strategy and Scenario: Net Present Value (Fictional, Millions of Dollars)

Strategy	Broad Tailwind (20 percent)	Costly Boom (30 percent)	Long Plateau (35 percent)	Hard Squeeze (15 percent)	Expected value	Worst case
A. Expand capacity now	150	120	20	(60)	64.00	(60)
B. Hold the present footprint	20	40	30	10	28.00	10
C. Phased expansion with stage gates	105	85	45	0	62.25	0
D. Acquire a regional competitor	90	60	15	(30)	36.75	(30)

Source: author's construction for a fictional company, with expected values computed in Python. Payoffs are assumptions chosen for teaching and are not forecasts. Figures in parentheses are losses.

Minimax regret. A third criterion asks a different question. The **regret** of a strategy in a scenario is the amount by which its payoff falls short of the best payoff available in that scenario, which is the sum that the company would regret not having earned. The **minimax regret** rule chooses the strategy whose largest regret is smallest. In Broad Tailwind the best payoff is $150 million, so that holding the footprint carries a regret of $130 million. Exhibit 19 shows the full table. Expanding now has the largest single regret of $70 million, in Hard Squeeze, holding has a maximum of $130 million, and acquiring has $60 million, while the phased strategy has a maximum regret of only $45 million. The phased strategy therefore wins on minimax regret, and its expected regret of $21.0 million is close to the $19.25 million of expanding now.

Exhibit 19 Regret by Strategy and Scenario, and the Minimax Regret Choice (Fictional, Millions of Dollars)

Strategy	Broad Tailwind	Costly Boom	Long Plateau	Hard Squeeze	Largest regret	Expected regret
A. Expand capacity now	0	0	25	70	70	19.25
B. Hold the present footprint	130	80	15	0	130	55.25
C. Phased expansion with stage gates	45	35	0	10	45	21.00
D. Acquire a regional competitor	60	60	30	40	60	46.50
Best payoff in the scenario	150	120	45	10		

Source: author's calculation in Python from the payoffs of the preceding exhibit. Regret is the best payoff in the scenario less the payoff of the strategy. Minimax regret selects strategy C, with a largest regret of 45.

How much to trust the answer. Three checks protect a decision table from false precision. The first is sensitivity. Expanding now and phasing have equal expected value when the probability of Hard Squeeze is 17.4 percent, with the other probabilities scaled in proportion, against the 15 percent that was assumed, so a small change in a judgment reverses the ranking by expected value and the phased strategy deserves more weight than the 1.75 million of expected value that it gives up. The second is the value of perfect information, which is the gain from knowing the scenario in advance. Choosing the best strategy in each scenario would give an expected $83.25 million, which is $19.25 million above the best expected value of $64.0 million. That figure is the most that the company should pay for information and flexibility, and it explains why the phased strategy, which buys some of the flexibility, is worth its cost. The third is the pattern across criteria. When the expected value, regret and worst case criteria point to three different strategies, as they do here, the decision turns on the company's capacity to bear loss and not on arithmetic, and the board must decide how much loss the company can bear. Module 5 provides the valuation of staged investment as a real option.

Long range planning and the budget. Scenarios must be tied to the plan and the budget, or they remain an exercise. Four horizons are useful, and Exhibit 20 sets them out. The annual budget, over twelve months, remains a commitment and an instrument of control. A **rolling forecast** adds a new period as each period closes, so that the company always looks the same distance ahead, and it is the main device for updating expectations between budgets. The strategic plan, over three to five years, is where scenarios and capital allocation meet. The long view, over about ten years, states the themes that the company cannot ignore and is the subject of Module 20. Deloitte's Finance Trends 2026 report quotes the corporate controller of a large retailer as saying that the company used to run scenarios monthly and now runs models and analysis almost daily, which shows the direction of practice, though the speed that suits a retailer may not suit every company.

Exhibit 20 Planning Horizons and Their Link to Capital Allocation and the Budget

Horizon	Purpose	Method	Cadence	Link to capital and the budget
Operating budget, 12 months	Commit resources and control spending	A base budget plus contingent budgets with pre approved triggers	Annual, with quarterly reviews	Base budget funds the central case. Contingent budgets are released only when a signpost trips
Rolling forecast, 18 months	Update expectations as facts arrive	Add a quarter as each quarter closes, with a range and not a point	Monthly or quarterly	Updates the covenant and liquidity outlook, and tells lenders early
Strategic plan, 3 to 5 years	Choose where to invest and how much to commit	Three or four scenarios, a decision table and stage gates	Refresh each year, and after any major event	Each large project carries scenario NPVs and a gate for the next tranche (Module 5)
Long view, about 10 years	Name the themes that could change the business	Foresight on technology, regulation, demographics and geopolitics	Every two or three years	Shapes research, partnerships and options, and informs Module 20 and Module 18

Source: author's synthesis, stated from the author's own knowledge. The remark on the frequency of scenario analysis is from Deloitte Insights, Finance Trends 2026 (October 6, 2025), as returned by the page read.

Linking scenarios to capital allocation. Module 5 treats the evaluation of investment. Scenarios add three practical rules to that evaluation. The first is to present each large project with its net present value in each scenario, not in one case, so that the board sees the spread. The second is to divide a large commitment into

tranches, each released by a stage gate that is tied to a signpost, so that the company pays for flexibility in small amounts. The third is to hold unallocated capital and debt capacity, which is an option on the scenarios that cannot yet be seen. A hurdle rate is not raised to cover uncertainty, since a higher discount rate penalizes distant cash flows without regard to which scenario they belong to, and the more honest course is to show the scenarios and let the board see the range. The budget carries the same logic. The base budget funds the central case, and a short list of contingent budgets, such as a cost reduction plan or a deferral of capital spending, is approved in advance with the trigger that releases each, so that the response to a signpost takes days and not months.

Planning for regime shifts. Four shifts deserve a place in almost every set of scenarios, though their form depends on the company, and Exhibit 21 shows how each enters the model. The first is the pace of adoption of artificial intelligence, which may change both a company's cost structure and its customers' demand, and which Modules 1 and 2 treat in detail. The second is the course of trade policy. The EY report of 2025 cited above set out four trade scenarios, namely hardening US and China blocs, a US that stands alone, proliferating protectionism and a grand bargain that lowers tariffs, and reported that, among the CFOs in its roundtable, 67 percent planned to pass higher costs on to customers, 50 percent to explore domestic sourcing and 42 percent to shift sourcing to non-tariffed regions. A Deloitte page on tariffs, which carried no date, reported that 55 percent of CFOs cited limited data access and 54 percent a fragmented strategy as key blockers, but did not name the survey. The third is the level of interest rates and the cost of refinancing, which belongs in every stress test of a company with floating debt, and the fourth is the strength of demand, which the recession case of the Federal Reserve illustrates. Module 18 treats tariff, currency and geopolitical exposure for a company that operates across borders.

Exhibit 21 Four Regime Shifts and How Each Enters the Scenario Model

Regime shift	What changes	Variable in the model	Early signposts	Finance response
Artificial intelligence	Cost to serve, speed of work, customers' capital plans, new competitors	Productivity gain, adoption pace, price pressure	Customer budgets for automation; competitors' pricing and service levels	Pilot with stage gates; reinvest only measured gains (Module 1)
Tariff and trade policy	Input cost, lead times, location of demand, retaliation risk	Effective duty rate, share of inputs affected, ability to pass on cost	Announced duties; port delays; supplier quotes	Second sources, pricing clauses, contingent sourcing budget (Module 18)
Interest rates	Cost of floating debt, refinancing risk, valuation	Floating rate, share of debt that floats, maturity ladder	Central bank path; credit spreads; lender appetite	Hedge part of exposure, extend maturities, keep covenant headroom (Module 12)
Demand regime	Volume, mix, payment behavior of customers	Revenue change, decremental margin, days sales outstanding	Backlog, quote volume, order cancellations	Contingent cost plan, working capital drive, early lender contact (Module 12)

Source: author's synthesis, stated from the author's own knowledge. The EY trade scenarios and the survey responses are from the EY report How CFOs are striving to get past uncertainty (2025, undated page), and the Deloitte tariff page was undated; both are described in Section 7.

Habits that keep foresight alive. Scenarios decay when they are written once and shelved. Five habits keep them in use. A **pre mortem** asks the team, before a decision, to imagine that the project has failed a year later and to write down the reasons, which draws out doubts that would otherwise stay unspoken. A **red team** is a

group or a person charged to argue against the plan, with access to the facts and no stake in the answer. The **signpost review** puts the dashboard and the status of each scenario on the board agenda every quarter, with the question whether the probabilities should change. The **forecast log** described above gives the finance team a record of its own accuracy. **Kill criteria** state before a project begins the facts that would stop it, so that stopping does not require anyone to admit error. Exhibit 22 sets out the cadence, and the practices are the author's recommendations from general experience and not findings of the sources of this module.

Exhibit 22 Organizational Habits That Keep Foresight Alive

Habit	How it works	Cadence	Output	Owner
Pre mortem	Imagine the decision failed, and write reasons on paper before any discussion; rank the causes	Before each major commitment	A list of failure causes, each with a mitigation or a signpost	Project sponsor, with the CFO
Red team	A small group with access to the facts argues the opposite case in writing and in person	Annual plan, and each major investment	A short memo of the strongest case against, and management's reply	Board or CFO appointed
Signpost review	The dashboard and the scenario status are reviewed against thresholds, and probabilities are revisited	Every board meeting, with a full refresh each year	Updated scenario probabilities and any trigger actions taken	CFO, with the audit or finance committee
Forecast log	Record forecasts with probabilities and score them	Quarterly scoring	Brier score, calibration table and lessons	Head of FP&A
Kill criteria and stage gates	State in advance the facts that stop or release each tranche of a project	At approval, and at each gate	Written gate conditions in the board paper	Project sponsor and CFO

Source: author's recommendations, stated from the author's own knowledge. The pre mortem and the red team are established practices, but no source on them was read for this module.

A scorecard for the foresight practice. A company can assess its own practice against five levels, which correspond to the scale of the self assessment of Section 8. Exhibit 23 describes the first, third and fifth levels for eight dimensions. A CFO should rate each dimension honestly, by documents and not by impression, and choose the two lowest as the program of work for the coming year.

Exhibit 23 Scorecard for the Foresight Practice: Levels 1, 3 and 5 for Eight Dimensions

Dimension	Level 1: Unprepared	Level 3: Competent	Level 5: Future-ready
Form of forecast	One number for each period	A range, with a sensitivity table	A distribution, with tail measures, reviewed by the board
Calibration	No record of past forecasts	Forecasts compared with results each quarter	A scored log with calibration, shared with the audit committee
Leading indicators	Reported results only	Several indicators, with thresholds	Back tested indicators, owners and trigger actions on a board dashboard
Stress testing	None, or a one case sensitivity	Single factor and combined shocks on covenants and liquidity	Reverse stress tests, refreshed twice a year and after trigger events

Dimension	Level 1: Unprepared	Level 3: Competent	Level 5: Future-ready
Scenarios	Best, base and worst case	Three or four scenarios with financial effects	Scenarios with signposts, probabilities and strategies tested against each
Strategic decisions	Chosen on the central case	Scenario NPVs shown for large projects	Regret and staged commitments used; unallocated capacity held
Link to budget and capital	Annual budget only	Rolling forecast and strategic plan with scenarios	Contingent budgets with pre approved triggers; gates tied to signposts
Habits	Pre mortem and red team unknown	Pre mortems for major decisions	Standing red team, quarterly signpost review, kill criteria in each approval

Source: author's proposal. The levels follow the five point scale of the proposal for this program, 1 Unprepared, 2 Developing, 3 Competent, 4 Advanced and 5 Future-ready, and the descriptions are not empirically validated.

4. Executive case study and decisions

The setting (fictional). Aldermoor Industrial Supply is a fictional company, and every name and figure in this case is invented for teaching. The chief executive has proposed to commit $120 million to a new plant that would add manufacturing capacity, citing a strong order book and the cost of lost orders in the past two years. The board will vote in six weeks. The CFO has been asked for an opinion, and the lenders of the $250 million term loan and the $100 million revolving facility will need to be told of any large commitment. Exhibit 24 gives the position of the company before the decision.

Exhibit 24 Position of Aldermoor Before the Capacity Decision (Fictional, Millions of Dollars)

Item	Value	Comment
Revenue, year just ended	$500.0 million	Plan for the coming year: $530.0 million at a 15.0 percent EBITDA margin
Plan EBITDA, coming year	$79.5 million	Free cash flow of $27.3 million after interest, tax and capital expenditure of $25.0 million
Net debt at the start of the year	$210.0 million	Term loan of $250.0 million less cash of $40.0 million; the revolver of $100.0 million is undrawn
Net leverage at plan year end	2.30 times	Covenant maximum of 3.75 times; headroom of 38.7 percent at plan
Liquidity policy	$60.0 million minimum	Policy is measured on cash alone in a stress, because a breach could close the revolver
Year three EBITDA: plan against scenarios	$88.5 million against $75.5 million	Only the Broad Tailwind scenario reaches the plan, and management assigned that scenario a 20 percent chance
The proposal	$120.0 million for new capacity	Expected to produce no new EBITDA in the first year, because the plant would be commissioned in the second

Source: author's construction for a fictional company, with figures computed in Python and consistent with the exhibits of Section 3.

The decision point. The board must decide whether to approve the full $120 million now, as the chief executive proposes, or to choose another course. The finance team has prepared the scenario table, the decision table and the stress test, and the CFO must say which strategy she recommends and why. The chief executive's case is straightforward: the strategy of expanding now has the highest expected value of the four,

$64.0 million, and in the best scenarios it has by far the highest payoff, $150 million in Broad Tailwind and $120 million in Costly Boom. The risks are on the other side of the table. In Hard Squeeze the same strategy loses $60 million, and, as Exhibit 25 shows, it breaches the leverage covenant even at plan, since net debt would rise to $302.8 million, or 3.81 times EBITDA, before the new plant earns anything.

Exhibit 25 Four Options Under the Covenant: Leverage at Plan and in a Demand Shock (Fictional, Millions of Dollars)

Option	Year one outlay	Net debt, plan year end	Leverage at plan (times)	Leverage if revenue falls 10 percent (times)	Headroom at plan (percent)	Expected value	Largest regret
A. Expand now	120.0	302.8	3.81	5.28	(1.6)	64.00	70
B. Hold	0.0	182.8	2.30	3.23	38.7	28.00	130
C. Phased: phase one	25.0	207.8	2.61	3.65	30.3	62.25	45
D. Acquire	90.0	263.8	2.88	3.83	23.1	36.75	60

Source: author's calculation in Python with the stress model of Section 3. Option A spends $120.0 million in the year, option C spends $25.0 million in phase one and $95.0 million later, and option D spends $90.0 million and adds $12.0 million of EBITDA. Headroom is the covenant of 3.75 times less leverage, divided by 3.75, and a negative figure is a breach. Figures in parentheses are negative.

The options. Option A, to expand now, has the highest expected value and the largest tail loss, and it breaches the covenant at plan. Option B, to hold the footprint, is safe, with leverage of 2.30 times, and has the worst regret, $130 million, because the company would watch the best scenarios arrive without capacity. Option D, to acquire a regional competitor, spreads the risk across a larger company but takes leverage to 2.88 times at plan and to 3.83 times in a demand shock, which is a breach. Option C, a phased expansion, commits $25 million now for permits, design and deposits on long lead equipment, and holds the remaining $95 million back until stated conditions are met. It leaves leverage at 2.61 times at plan and 3.65 times if revenue falls 10 percent, inside the covenant in both cases, and it has the smallest largest regret, $45 million, though its expected value is $1.75 million lower than that of option A.

The recommended path. The CFO recommends option C, with three conditions attached to the second phase. The first is that the signposts of Exhibit 26 must have been met, which means that the demand that the chief executive sees is confirmed in the backlog and not only in the sales team's view. The second is that the company must have either an amendment to the leverage covenant, negotiated before it is needed, or pro forma headroom of at least 15 percent after the spending, because the stress test shows that a breach would close the revolver. The third is that the effective duty on the main inputs, as measured by the dashboard, does not exceed 15 percent for two quarters, or that if it does, the plant is justified in the Costly Boom scenario, in which domestic capacity earns a premium. The CFO also recommends that the company take two no regret moves in the meantime: to extend the maturity of the revolver, and to open second sources for the three largest inputs.

The reasoning. The recommendation does not rest on a belief that Broad Tailwind is unlikely. It rests on three observations. The first is that the phased option gives up only $1.75 million of expected value, which is 2.7 percent of the $64.0 million of option A, in exchange for a reduction of the worst case from a loss of $60 million to zero and of the largest regret from $70 million to $45 million. The second is that the expected value ranking is fragile, since a rise in the probability of Hard Squeeze from 15 percent to 17.4 percent is enough to remove the advantage of option A. The third is that the value of knowing the scenario is as much as $19.25 million, and option C buys a large part of that value for a small premium by deferring the largest commitment

until the backlog and the duties have shown which world is arriving.

Exhibit 26 Stage Gates for the Phased Expansion: Conditions and Actions (Fictional, Millions of Dollars)

Gate	Spending released	Conditions to be met	If the conditions fail
Gate 1: approval	$25.0 million, phase one	Board approval; deposits are placed with terms that allow the order to be reduced or cancelled; the revolver maturity extension is under way	Stop at phase one, and keep the permits and the design as an option
Gate 2: after two quarters	$45.0 million, phase two	Backlog to billings above 1.10 for two consecutive quarters; quote volume not down; covenant amended or pro forma headroom of at least 15 percent; cash only liquidity above the policy minimum of $60.0 million	Hold, review the dashboard monthly, and shift the budget to the cost plan if two or more indicators turn red
Gate 3: after four quarters	$50.0 million, phase three	Gate 2 conditions still met; effective duty on inputs not above 15 percent for two quarters, or the Costly Boom case confirmed by the domestic share of tenders	Slow or stop, and consider a smaller final tranche or a partnership for the last stage
Standing rule	Any tranche	Any red status on days sales outstanding or on the floating rate, together with covenant headroom below 15 percent, pauses all further spending until the board has reviewed the stress test	Pause and review

Source: author's construction for a fictional company. The tranches sum to $120.0 million: 25.0 plus 45.0 plus 50.0. The thresholds are examples and a real company would set them from its own history and agreements.

What a weaker answer misses. A weaker answer would take one of three forms. The first would present the expected value of $64.0 million as the answer and recommend expansion, without noticing that the strategy breaches the covenant at plan, which no spreadsheet that ignores the debt agreement would reveal. The second would see the breach and recommend holding the footprint on the ground of safety, and so would accept the largest regret of $130 million and leave the company without capacity in the scenarios in which capacity is most valuable. The third would adopt the phased option but without signposts, so that the second tranche would be released by habit or by the enthusiasm of the sponsor and not by evidence. The stronger answer keeps the strategy tied to a monitoring system, to a financing plan and to a lender conversation that takes place before it is needed.

Questions to consider. What would the board need to see to move the probability of Hard Squeeze above 17.4 percent, and who would present that evidence? Which of the signposts could be gamed or delayed, and who owns each? What would the lenders ask for in an amendment, and what could Aldermoor offer in exchange (Module 16)? How would the answer change if the probabilities were assigned by the sales team and not by the finance team? Which of the four strategies has the least value if the model of the demand shock is wrong in the direction of a deeper fall?

5. Boardroom discussion and key takeaways

The board should expect to see ranges, scenarios and signposts, and should ask about them. Module 9 describes the oversight role of directors, and the questions in Exhibit 27 may be put by any director to the CFO. A CFO who has done the work of this module can answer each of them with evidence, and a CFO who cannot is giving the board a number and not a view of the future.

Exhibit 27 Ten Questions Directors Should Ask About Foresight and Scenario Planning

Question a director should ask	What a ready CFO can say
What is the probability that we miss the plan, and how far, and where did that number come from?	The simulation or equivalent, with stated inputs and the percentiles (Exhibit 8, Exhibit 9)
How good have our own forecasts been, and how do we know?	The forecast log, the Brier score and the calibration table (Exhibit 10)
Which indicators warn us first, and have they ever been wrong?	The back test with hits and false alarms (Exhibit 11, Exhibit 12)
Which single event or combination would breach our covenants, and how close are we?	The stress test and the reverse stress test (Exhibit 13, Exhibit 14)
Is our liquidity measured with or without a revolver that a breach could close?	Both figures, and the policy minimum measured on cash alone (Exhibit 13)
What are the two or three uncertainties that matter most to this decision, and what are the scenarios?	The grid, the narratives and the financial effects (Exhibit 15, Exhibit 16, Exhibit 17)
How does this investment perform in each scenario, and what is its largest regret?	The decision and regret tables, and the sensitivity of the ranking (Exhibit 18, Exhibit 19)
What would make us stop, and who will tell us?	Gates, kill criteria and the signpost review (Exhibit 26, Exhibit 22)
How does the budget respond if a signpost trips, and who may release contingent funds?	Contingent budgets with pre approved triggers (Exhibit 20)
When did we last change our view because of a signpost, and what did it cost or save?	A dated example from the signpost review log (Exhibit 28)

Source: author's synthesis. Cross references are to modules in this program: Module 5 for capital allocation, Module 9 for governance, Module 12 for crisis management and financial resilience, Module 18 for global strategy and Module 20 for the future of the role.

Advising the chief executive and the board. The CFO should make three arguments. The first is that the company should be shown a range and not a point, since a point conceals the tail that the covenants and the liquidity policy must survive. The second is that large commitments should be staged and tied to signposts, because the company does not know which future it is entering and the cost of learning in small steps is low. The third is that the habits of foresight, the log, the dashboard, the red team and the signpost review, are cheap and are the only part of this module that compounds over time. EY's 2025 report urged CFOs to plan around scenarios and to build agility in supply chains and financing, and Deloitte reported that 30 percent of the finance leaders in its survey planned to strengthen scenario planning, which supports the view that the practice is becoming a normal expectation.

Common errors. Seven errors recur. The first is to treat the budget as the forecast and the forecast as a promise. The second is to build scenarios that are the same story at three volumes, a good case, a middle case and a bad case, without any difference in the underlying cause. The third is to assign probabilities with false precision, or to refuse to assign any, so that the decision table cannot be used. The fourth is to stress test the plan with shocks that are mild or that were chosen after the fact. The fifth is to test liquidity without testing covenants. The sixth is to choose indicators for their convenience and not for their record. The seventh is to write scenarios once and file them, so that the board does not see them again until the next plan.

Key takeaways. First, a single point forecast is a statement about the middle of a distribution, and the simulation of Aldermoor shows that the plan was met or exceeded in barely more than 46 percent of trials and that its mean was below the plan. Second, a forecast log and a Brier score turn confidence into evidence, and

most teams are overconfident at the top of the scale. Third, a leading indicator earns its place by a back test, which counts hits and false alarms, and by a named action. Fourth, a stress test should read the covenant and the liquidity together, and a reverse stress test shows that moderate shocks combine to reach a covenant at about half of the single factor distance. Fifth, scenarios are built from a focal question and two critical uncertainties, and they are stories with numbers and signposts. Sixth, a decision table that shows expected value, regret and the worst case, together with sensitivity, tells the board how much loss the company can bear. Seventh, large commitments should be staged and tied to signposts, so that flexibility is bought in small amounts. Eighth, the habits of the pre mortem, the red team and the quarterly signpost review keep the discipline alive. Ninth, the cross references matter: Module 5 values the options, Module 12 supplies the playbook when a signpost trips, Module 18 treats the trade and currency scenarios and Module 20 looks beyond the planning horizon.

6. The executive artifact: a scenario planning one pager and a stress test template

The artifact has two parts, each one page. The first is the **scenario planning one pager**, which holds the focal question, the scenarios, the chosen strategy, the signposts and the trigger action for each signpost, and which the board should see at every meeting. The second is the **stress test template**, which records the shocks, the results on the covenant and liquidity and the action to be taken if a shock is approached. Both are shown completed for the fictional Aldermoor, with the figures of Sections 3 and 4. The reader should complete them for the company's own most important decision and review them at each board meeting.

Exhibit 28 Scenario Planning One Pager, Completed for the Fictional Aldermoor Capacity Decision

Element	Entry for Aldermoor (fictional)
Focal question and horizon	Should the company commit $120.0 million to new capacity within three years, and on what terms? Horizon: three years. Reviewed each board meeting; full refresh each year
Critical uncertainties	Trade regime (open or fragmented) and customer demand (strong or weak)
Scenarios and probabilities	Broad Tailwind 20 percent; Costly Boom 30 percent; Long Plateau 35 percent; Hard Squeeze 15 percent. Year three EBITDA of 105.4, 76.1, 72.8 and 40.9, with a weighted value of 75.5 against a plan of 88.5
Strategy chosen	C: phased expansion with stage gates. Largest regret of 45, expected value of 62.25, worst case of zero
No regret moves now	Extend the maturity of the revolver; open second sources for the three largest inputs; start a scored forecast log
Signpost 1: demand	Backlog to billings ratio. Reading 1.04 (amber). Release phase two if above 1.10 for two quarters; hold and review if below 1.00. Owner: head of sales operations
Signpost 2: trade	Effective duty on imported inputs. Reading 11.0 percent (amber). If above 15.0 percent for two quarters, activate second sources and test the Costly Boom case. Owner: head of procurement
Signpost 3: credit	Days sales outstanding and floating rate. Readings 62 days (red) and 7.0 percent (green). Escalate overdue accounts now; review the swap if the rate crosses 7.5 percent. Owner: treasurer
Signpost 4: covenant	Pro forma leverage headroom. Plan 38.7 percent. Pause all uncertain spending if below 15 percent. Owner: CFO
Trigger actions	Contingent cost plan approved in advance; pause rule for capital spending; lender call within five business days of any red status on two indicators

Element	Entry for Aldermoor (fictional)
Review	Signpost review at every board meeting (about ten minutes); probabilities revisited each quarter; red team memo before gate 2

Source: author's template, completed with the figures of Sections 3 and 4. All names and figures are fictional. The thresholds and owners are examples that each company must set for itself.

Exhibit 29 Stress Test Template, Completed for the Fictional Aldermoor Plan Year (Millions of Dollars)

Shock	Parameter	Net leverage (times)	Interest cover (times)	Cash only liquidity	Action if reached
Plan	As budgeted	2.30	4.82	67.3	None; the plan case is the reference
Revenue down 10 percent	Revenue minus 10 percent	3.23	3.53	62.0	Contingent cost plan; defer uncommitted capital spending
Tariff cost of 2.0 percent of revenue	Tariff cost of 2.0 percent of revenue	2.77	4.18	59.3	Pricing action and second sources; review the covenant forecast
Floating rates up 200 basis points	Floating rates plus 200 basis points	2.33	4.08	65.0	Review hedges; check the interest cover limit of any lender
Receivables slow by 10 days	Receivables slower by 10 days	2.48	4.82	52.7	Collections drive; review credit terms and factoring options
Combined shock	All four shocks together	4.32	2.50	39.5	Lender call before the breach; pause capital; negotiate the amendment
Severe combined shock	Revenue minus 20 percent, tariff 3.0 percent, rates plus 300, receivables plus 15 days	9.35	1.16	22.2	Liquidity plan of Module 12 and advisers; consider asset sales or new capital

Source: author's template, completed with the results of the stress test of Section 3. Covenant maximum 3.75 times; liquidity policy minimum $60.0 million on cash alone. The company should add the shocks that are specific to its own business and agreements.

How to use the artifact. Complete the one pager first and keep it to one page, since a document that cannot be read in ten minutes will not be read. Show it to the board with the same layout each quarter, so that changes of reading and probability stand out. Complete the stress template next, using the actual terms of the company's debt agreements, which a qualified adviser should read for the exact definitions of the covenants. Neither document is a prediction, and neither is advice on law, tax or accounting. Each is a way of agreeing in advance, in calm conditions, what the company will do when events begin to show which future is arriving.

7. Research and Perspectives from the Leading Firms

The publications below were found by search and each page was opened before it was described. Where a page returned only part of its content, or did not state a date or an author, the description says so, and only what the page returned is stated. The models of Sections 3 and 4, namely the simulation, the Brier score, the back test, the dashboard, the stress test, the scenarios, the decision table and the case, are the author's own and are not taken from these pages. Practice changes quickly in this field, and the reader should confirm the current text before relying on any item.

Exhibit 30 Publications Used in This Module

Source	Title and URL	What it says
McKinsey and Company	Hugh Courtney, Jane Kirkland and Patrick Viguerie, Strategy under uncertainty, McKinsey Quarterly (June 1, 2000), adapted from a 1997 Harvard Business Review article. https://www.mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/strategy-under-uncertainty	Argues that traditional strategy assumes a single predictable future, and proposes four levels of residual uncertainty, from a clear enough future to true ambiguity, with a method and a posture for each, namely to shape, to adapt or to reserve the right to play. Suggests no more than four or five scenarios at the third level. The article is from 2000 and the text was read as returned by the page
Shell (Ged Davis)	Ged Davis, Scenarios as a Tool for the 21st Century, presented July 12, 2002 at the Probing the Future Conference, Strathclyde University. https://www.pik-potsdam.de/ateam/avec/peyresq2005/talks/0921/leemans/literature/davis_how_does_shell_do_scenarios.pdf	A Shell vice president describes how the company builds scenarios, from the frames of reference to the drivers, the predetermined elements and the critical uncertainties, and states that scenarios are not projections, predictions or preferences but consistent and credible stories. The document is hosted by a third party, and the paper is more than twenty years old
Shell (corporate site)	What are Shell scenarios, Shell (undated page). https://www.shell.com/news-and-insights/scenarios/what-are-shell-scenarios.html	Only the page description was returned, which says that Shell scenarios go beyond traditional energy outlooks and explore economics, geopolitics, social change and environmental stresses. No body text, date or method was returned, and nothing further is claimed
Deloitte	Steve Gallucci, Ed Hardy, Justin Silber, Dave Turk and Timothy Murphy, Finance Trends 2026: Navigating the expanded scope of finance, Deloitte Insights (October 6, 2025). https://www.deloitte.com/us/en/insights/topics/leadership/finance-trends-leadership/2026.html	Reports a survey of 1,326 respondents in which economic uncertainty was the top ranked risk at 26 percent and 30 percent of finance leaders planned to strengthen advanced scenario planning, and quotes a retailer's controller on running scenarios almost daily. The survey sample and method beyond the count were not described on the page
Deloitte	Impact of tariffs: Overcoming volatility, Deloitte US, Tariff Impact on Business: CFO Strategies (undated, with a 2026 notice). https://www.deloitte.com/us/en/what-we-do/capabilities/finance-transformation/articles/tariff-impact-business-cfo-strategies.html	Advises finance leaders to refresh models quickly, to gain visibility into performance drivers and to align tariff consequences with supply chain choices, and cites survey figures of 55 percent of CFOs for limited data access and 54 percent for a fragmented strategy, without naming the survey. The page did not mention scenario planning or stress testing
EY	EY Center for Executive Leadership, How CFOs are striving to get past uncertainty (2025, no date on the page). https://www.ey.com/content/dam/ey-unified-site/ey-com/en-us/campaigns/ey-center-for-executive-leadership/documents/ey-how-cfos-are-striving-to-get-past-uncertainty.pdf	Reports a roundtable of Fortune 250 CFOs in which 55 percent were more bearish about the rest of 2025 and 2026, lists four trade scenarios from an EY geostrategy leader and gives CFO responses to tariffs. The number of respondents was not stated, so the statistics describe a small group and not a population
Federal Reserve	Board of Governors of the Federal Reserve System, 2026 Stress Test Scenarios, final supervisory scenarios (February 2026). https://www.federalreserve.gov/publications/files/2026-final-supervisory-stress-test-scenarios-20260204.pdf	Sets out the hypothetical severely adverse scenario for large banks, with unemployment rising from 4.5 percent to a peak of 10 percent, real GDP down 4.6 percent, house prices down about 30 percent and equity prices down about 58 percent, and states that it is not a forecast. It is designed for banks and not for other companies

Source	Title and URL	What it says
Academic (Mellers and colleagues)	Barbara Mellers, Joshua Baker, Eva Chen, David Mandel and Philip Tetlock, How generalizable is good judgment? A multi-task, multi-benchmark study, Judgment and Decision Making, vol. 12, no. 4 (July 2017), pp. 369 to 381. https://dlab.sauder.ubc.ca/sjdm/journal/17/17408/jdm17408.html	Describes superforecasters as the top two percent of participants in government sponsored tournaments, scored by Brier scores, lists factors thought to explain their accuracy, and notes that expert predictions improve where feedback is regular and clear. The page did not report new Brier score data, and the book Superforecasting itself was not read

Each page was opened and the descriptions follow what the tool returned. No publication of PwC, KPMG, RSM US, BDO USA, Grant Thornton, Baker Tilly, Crowe, Forvis Mazars, CBIZ, Moss Adams, CohnReznick, CLA or EisnerAmper was read for this module, and none is cited.

What the research agrees on. The sources differ in kind, from a regulator's scenario to a consulting firm's survey, a company's account of its own method and a study of forecasting skill, and they should not be read as one body of evidence. They agree on a few practical points. A single forecast understates uncertainty, and the method should fit the level of uncertainty. Scenarios are stories for testing strategy and not predictions. Severe cases are built as linked paths of several variables. Forecasting skill improves with feedback and with open minded updating. CFOs report that uncertainty is high and that scenario planning is a growing priority. None of the sources shows, by independent evidence, that the use of scenarios or a stress test improves the financial results of a company, and the tools of this module are offered as sound practice and not as proven remedies.

Limits of the evidence. Four cautions apply. First, several sources are surveys or roundtables of executives, whose samples are small or undescribed, and several come from firms that also sell services. Second, a simulation and a decision table are only as good as their assumptions, and the probabilities in this module are judgments chosen for teaching, so a real company should treat them as the start of an argument. Third, a model built from past data describes the uncertainty that the modeler has already imagined, and the surprises that matter most are often outside it. Fourth, the Federal Reserve scenario is a regulatory tool for banks, and the practice of Shell is that of a large energy company with a long planning horizon, so neither transfers to a mid sized company without adaptation.

How a CFO can use the research. The McKinsey article is the best guide to choosing a method, because it asks first how much uncertainty remains and then fits the tool to the answer. The Davis paper supplies the order of work for a scenario project and the warning that the process matters as much as the stories. The Deloitte and EY pages are most useful as evidence of what peers report and worry about, and as a source of questions for the board, with the caution that their figures come from surveys of unstated or small samples. The Federal Reserve document is a model of how a severe scenario is written down, with variables, a path over several quarters and an explicit statement that it is not a forecast. The study of Mellers and colleagues is a reminder that forecasting is a skill that can be measured and improved, and that the improvement comes from feedback, scoring and teamwork. The reader's own log of forecasts, stress tests and signposts, reviewed honestly each quarter, will teach more than any average of what other executives report.

What the sources do not settle. Three questions remain open and each deserves the reader's own judgment. The first is how many scenarios are enough. The McKinsey authors suggest no more than four or five, and the two by two grid gives four, but a company with several independent uncertainties may need a second grid for a second decision. The second is how to assign probabilities, since a scenario purist may refuse to do so on the ground that probabilities invite prediction, while a decision analyst needs them to compute an expected value. This module takes the middle course of assigning rough probabilities, showing the sensitivity of the answer to them and relying on regret and the worst case when they are in doubt. The third is how often to

refresh the work. A yearly refresh with a quarterly signpost review is a reasonable default, but a company in a fast moving market may need a monthly rhythm, as the retailer quoted by Deloitte suggests.

8. The ten question self assessment

Rate each statement from 1 to 5 using the scale below, and mark one box in each row. The ratings must rest on actual experience or documented evidence, such as a forecast log you kept, a stress test you ran, a scenario one pager you completed, a decision table you presented or a signpost review in a board minute, and not on confidence. A rating of 1 means Unprepared: limited understanding and no meaningful application. A rating of 2 means Developing: conceptual familiarity, relying heavily on others. A rating of 3 means Competent: able to apply the capability independently in conventional situations. A rating of 4 means Advanced: the capability is integrated into strategy and delivers measurable results. A rating of 5 means Future-ready: enterprise wide impact, anticipation of change and the development of others.

No.	Statement of demonstrated capability	1	2	3	4	5
1	I present forecasts to the board as ranges with probabilities and tail measures, and I can show the simulation or equivalent analysis behind the current plan.	☐	☐	☐	☐	☐
2	I keep a dated forecast log with probabilities, I compute a Brier score or an equivalent measure, and I have changed how I state confidence because of the calibration results.	☐	☐	☐	☐	☐
3	Each leading indicator that I monitor has been back tested for hits and false alarms, and has an owner, a threshold and a named trigger action.	☐	☐	☐	☐	☐
4	A one page dashboard of indicators with amber and red lines is reviewed by the board or its committee at each regular meeting.	☐	☐	☐	☐	☐
5	I run single factor and combined stress tests on revenue, margin, covenants and liquidity at least twice a year, and I measure liquidity both with and without facilities that a breach could close.	☐	☐	☐	☐	☐
6	I have run a reverse stress test, and I can state which combination of moderate events would breach our covenants or our liquidity policy.	☐	☐	☐	☐	☐
7	I have led scenario planning from a focal question through critical uncertainties to named scenarios with financial effects, probabilities, signposts and strategies tested in each.	☐	☐	☐	☐	☐
8	Large investment proposals show value in each scenario, use staged commitments with gates tied to signposts, and compare strategies by regret and worst case as well as expected value.	☐	☐	☐	☐	☐
9	Our planning cycle includes a rolling forecast, contingent budgets with pre approved triggers and an annual refresh of scenarios, and I can link scenarios to capital allocation.	☐	☐	☐	☐	☐
10	We use pre mortems and a red team for major decisions, we write kill criteria into approvals, and the board reviews signposts and scenario probabilities every cycle.	☐	☐	☐	☐	☐

Total score: _____ out of 50. Ten statements rated from 1 to 5 give a minimum of 10 and a maximum of 50 points. Exhibit 31 gives the proposed interpretation. The scale is that of the source proposal for this program, and the statements are specific to this module.

Status of this module. Module 17 is not one of the five critical modules named in the proposal, which are Modules 1, 5, 9, 10 and 13, so the rule that a score below 30 of 50 blocks the highest designation does not apply to it as proposed. A low score here, however, may still weaken the effect of strong scores elsewhere, since the capital allocation of Module 5 and the resilience of Module 12 both depend on the habits of foresight that this module describes.

Exhibit 31 Proposed Interpretation of the Module 17 Score

Score	Reading	Suggested use
10 to 19	Unprepared	Study the module in full, start a forecast log, and run a simple stress test of covenants and liquidity on the current plan before the next board meeting
20 to 29	Developing	Build the first dashboard with tested indicators, compute a Brier score on a quarter of forecasts, and write one scenario one pager for the largest open decision
30 to 39	Competent	Add a reverse stress test and a decision table with regret, introduce stage gates into the next large approval, and begin a quarterly signpost review
40 to 45	Advanced	Hold a red team review each year, refresh probabilities each quarter with documented reasons, and link contingent budgets to triggers
46 to 50	Future-ready	Develop successors in probabilistic judgment, share the practice with peers and boards, and test whether the company changed a decision because of a signpost

These bands are proposed program standards and have not been empirically validated against a broad population of finance executives. Module 17 is not designated a critical module in the proposal.

Development actions by pattern. A single total can hide the pattern behind it, so look at which statements received the lowest ratings. Exhibit 32 suggests three actions for each pattern.

Exhibit 32 Development Actions for Low Score Patterns

Low scores on	Pattern	Three suggested development actions
Statements 1 and 2	Probability and calibration gap: forecasts are points and confidence is untested	1. Rebuild the plan as a simulation like Exhibit 8. 2. Start the log and score it as in Exhibit 10. 3. Show the percentiles and the probabilities of Exhibit 9 to the board
Statements 3 and 4	Indicator gap: warnings are informal or untested	1. Back test candidates as in Exhibit 11. 2. Build the dashboard of Exhibit 12 with owners and triggers. 3. Put it on the board agenda
Statements 5 and 6	Stress gap: balance sheet shocks are mild, single factor or ignore the covenant	1. Complete the template of Exhibit 29. 2. Run the reverse test of Exhibit 14. 3. Measure liquidity with and without facilities that a breach could close
Statements 7 and 8	Scenario and decision gap: stories without numbers, or choices on one case	1. Complete the grid of Exhibit 15 and the narratives of Exhibit 17. 2. Build the payoff and regret tables of Exhibit 18 and Exhibit 19. 3. Stage the next large commitment as in Exhibit 26
Statements 9 and 10	Planning link and habit gap: foresight is an annual exercise	1. Add contingent budgets as in Exhibit 20. 2. Hold a pre mortem before the next major decision as in Exhibit 22. 3. Put the signpost review in the board calendar using Exhibit 28

Source: author's suggestions. The mapping of statements to patterns is the author's proposal, and ratings should rest on experience or documented evidence.

A twelve month plan for the reader. The module is most useful when turned into a plan with dates. In the first month the reader should start the forecast log, rebuild the plan of the coming year as a range with the main uncertainties named, and run the single factor stress tests on covenants and liquidity. In the second and third months the reader should complete the reverse stress test, choose four or five candidate indicators and begin to back test them with the history available. In the fourth to sixth months the reader should hold the first scenario workshop on the largest open decision, with a red team, and bring the one pager and the dashboard to the board. In the second half of the year the reader should stage the next large commitment with gates, add contingent budgets to the planning cycle and score the first two quarters of the forecast log. At the end of the year the reader should repeat the self assessment of this section and compare the result with the first, using documents and not impression.

Using the assessment with a team. The assessment may be completed by the CFO alone, but it is more useful when the head of FP&A, the treasurer and one director rate the same ten statements independently, and the group then discusses the statements on which the ratings differ most. A gap between the CFO's rating and a director's rating is itself a finding, since it often shows that the work is done but not shown to the board. Each rating should be supported by a document that a colleague could read, and a rating without a document should be lowered by one point until the document exists.

9. Currency check and sources

Currency check: this field changes, and surveys, scenarios and published guidance shown here may already have changed. The publications in Exhibit 30 were opened in October 2026. The following were confirmed on pages read: the four levels of uncertainty and the three postures of the McKinsey Quarterly article; the steps and the statements on scenarios in the Davis paper of 2002; the Deloitte survey figures of 26 and 30 percent and the sample of 1,326; the Deloitte tariff page figures of 55 and 54 percent; the EY roundtable figures and the four trade scenarios; the figures of the Federal Reserve severely adverse scenario for 2026; and the description of superforecasters and of the factors behind their accuracy in the paper of Mellers and colleagues. Only the page description of the Shell corporate page was returned, and the text of the Davis paper was returned in summary by the reading tool. The following come from the author's own knowledge and were not confirmed in a source read: the definitions of the Brier score, the Monte Carlo method, conditional value at risk, regret and minimax regret; the vocabulary of probability and the list of forecasting errors beyond those named in the Mellers page; the reference class forecasting and planning fallacy counters; the pre mortem, the red team, kill criteria and stage gates; the structure of rolling forecasts and contingent budgets; and the description of how a covenant breach can affect a revolving facility, which depends on the terms of each agreement. The company, its people and every figure in the exhibits are fictional. All sums, ratios, probabilities, percentiles and counts were computed in Python with a fixed random seed of 17 and checked by assertion, and the probabilities, payoffs and thresholds are assumptions chosen for teaching and not forecasts. This module is educational and is not accounting, tax, legal, securities, insurance or lending advice. Readers must consult qualified professionals in accounting, tax, legal and other relevant services before acting, and debt agreements, covenant definitions, disclosure of forward looking information, tariff and trade rules and any lender discussion in particular require the advice of qualified professionals. The templates, scorecards and scales shown here are proposed program standards and examples, not rules of any regulator, and each company must adapt them to its agreements, its sector and the law of the places in which it operates.

Sources: McKinsey Quarterly, Strategy under uncertainty (2000); Ged Davis of Shell, Scenarios as a Tool for the 21st Century (2002); Shell, What are Shell scenarios (undated); Deloitte Insights, Finance Trends 2026 (2025); Deloitte US, Impact of tariffs (undated); EY Center for Executive Leadership, How CFOs are striving to get past uncertainty (2025); Federal Reserve Board, 2026 stress test scenarios (2026); Mellers and colleagues, Judgment and Decision Making (2017). Titles, dates and URLs appear in Exhibit 30. Aldermoor Industrial Supply and every figure in its exhibits are fictional.