

# **THE FOUNDER FINANCIAL INFRASTRUCTURE (FFI) STANDARD**

## **Book 2: Performance Modeling and Forecasting** Beta v0.5

April 2026

<https://ffistandard.org/>

© The Oakworth Group, 2026.

Published under Creative Commons Attribution 4.0 International (CC BY 4.0).  
This Standard may be freely used, cited, adapted, and distributed with attribution.  
Full license terms: [creativecommons.org/licenses/by/4.0](https://creativecommons.org/licenses/by/4.0)

## **BOOK 2: PERFORMANCE MODELING AND FORECASTING**

Book 2 defines the standards for the financial models a company must maintain to represent its forward performance: the forecasting methodology, the unit economics framework, the scenario and sensitivity architecture, the growth modeling standard, and the cost structure standard. These five domains together constitute the analytical layer of financial infrastructure. The integrated three-statement model governed by Book 1 is the structural layer. Book 2 governs the logic that drives it.

### **SECTION 2.1: THE FORECASTING METHODOLOGY STANDARD**

#### **PURPOSE**

The forecasting methodology standard governs how a company constructs, documents, and maintains its financial forecasts. A forecast is not a financial statement. A financial statement records what occurred. A forecast represents what a company projects will occur, derived from documented assumptions about how its operations will develop. The reliability of a forecast is determined by the transparency and methodological soundness of its assumptions, not by the proximity of its outputs to actual results.

#### **DEFINITIONS**

**FORECAST:** A projection of a company's future financial performance, expressed through revenue, expense, and cash position estimates for one or more future periods, derived from documented operational assumptions.

**BOTTOM-UP FORECAST:** A forecast in which revenue and expense projections are derived from granular operational assumptions, including customer count, pricing, conversion rates, sales capacity, and headcount, that aggregate to company-level totals. A bottom-up forecast is the required methodology at Level 2 and above. Its defining characteristic is that every output can be traced to an operational driver.

**TOP-DOWN FORECAST:** A forecast in which revenue projections are derived from market size estimates and assumed capture rates, rather than from operational drivers. A top-down forecast does not meet the requirements of this Standard at Level 2 or above because it does not demonstrate a mechanistic connection between the company's operational capacity and its projected revenue.

**DRIVER:** A quantified operational variable whose value, when changed, produces a calculable change in one or more financial outputs. Revenue per customer, monthly customer acquisition count, and average contract value are examples of revenue drivers. Headcount by function and average fully loaded cost per employee are examples of expense drivers. A driver must be stated as a discrete, modifiable assumption in the assumption layer defined in Book 1, Section 1.1.

**FORECAST PERIOD:** The number of future periods covered by a forecast from the date of preparation. The minimum forecast period required at each compliance level is stated in the compliance criteria of this section.

**FORECAST ACCURACY:** The ratio of actual financial results to forecast results for a completed period, expressed as a percentage. Forecast accuracy is a trailing measure used to assess the reliability of the forecasting methodology, not a measure of management performance.

**ROLLING FORECAST:** A forecast that extends a fixed number of periods forward from the current period, updated at each period end to maintain consistent forward visibility. The rolling forecast extends the forecast period by one period each time one period closes.

## **PRINCIPLES GOVERNING THIS SECTION**

Principle 2 from Book 0, Section 0.3 governs this section in its entirety: clarity of assumptions governs over precision of outputs. A bottom-up forecast built on documented, bounded, operationally grounded assumptions satisfies this Standard at a higher compliance level than a forecast built on undisclosed or optimistic inputs that produces granular outputs.

The methodological basis for every driver must be documented. Where a driver is derived from historical company data, that derivation must be traceable. Where a driver is derived from external market reference, the source must be identified. Where a driver is derived from management judgment, it must be labelled as such. These three categories of derivation are not ranked by reliability; they are distinguished to allow a third party to assess each assumption independently.

## **COMPLIANCE CRITERIA**

### **Level 1**

2.1.L1.1: The company maintains a financial forecast covering a minimum of twelve months from the date of preparation.

2.1.L1.2: The forecast includes projected revenue by period and projected total operating expenses by period.

2.1.L1.3: The key assumptions underlying the revenue projection are documented and can be stated on request.

2.1.L1.4: The forecast is updated at minimum quarterly to reflect changes in actual performance and changes in forward assumptions.

### **Level 2**

2.1.L2.1: The company maintains a bottom-up financial forecast covering a minimum of twenty-four months from the date of preparation.

2.1.L2.2: Every revenue line in the forecast is derived from at least one documented driver that is stored in the assumption layer defined in Book 1, Section 1.1.

2.1.L2.3: Every expense line in the forecast is derived from at least one documented driver; personnel expenses are derived from a headcount plan that lists roles, hire dates, and fully loaded costs per role.

2.1.L2.4: The forecast is updated monthly; and each update documents any assumption that has changed from the prior version and the reason for the change.

2.1.L2.5: The methodological basis for each driver is documented as one of the following: historical company data; external market reference with source identified; or management judgment labelled explicitly as such.

2.1.L2.6: The company maintains a forecast accuracy record for each completed period showing actual results against the forecast prepared for that period at least one month in advance.

2.1.L2.7: The company maintains a rolling forecast that extends the forecast period by one month each time a month closes.

### **Level 3**

2.1.L3.1: The company maintains a bottom-up financial forecast covering a minimum of thirty-six months from the date of preparation.

2.1.L3.2: Each driver in the forecast includes a documented sensitivity range stating the high and low values considered plausible by management, with the basis for each bound stated.

2.1.L3.3: The forecast is reviewed formally each month against actual results; material variances are documented with identified causes; and the forecast methodology is adjusted where a variance reveals a systematic error in driver assumptions.

2.1.L3.4: The company maintains a version-controlled forecast history such that any prior version of the forecast can be retrieved and compared to the current version with all changes identified.

2.1.L3.5: The forecast integrates with the three-statement model maintained under Book 1, Section 1.1, such that a change to any forecast driver updates the income statement, cash flow statement, and balance sheet without manual recalculation.

### **BENCHMARKS**

The following benchmarks represent observed norms as of Beta v0.5. Benchmarks marked with an asterisk are directional indicators where market data is insufficient for precise quantification in this version. They will be refined in v1.0.

Forecast period by stage:

Pre-Revenue: minimum twelve months

Early Revenue: minimum twelve months

Growth Stage: minimum twenty-four months

Scale Stage: minimum thirty-six months

Forecast update frequency by stage:

Pre-Revenue: quarterly minimum

Early Revenue: monthly minimum

Growth Stage: monthly minimum

Scale Stage: monthly minimum

Forecast accuracy tolerance:

- Growth Stage, revenue: plus or minus twenty percent of forecast for the relevant period is the outer bound of acceptable variance before a methodology review is required. Variance beyond this range for two consecutive periods requires documented investigation of driver assumptions.
- Scale Stage, revenue: plus or minus fifteen percent is the outer bound.
- All stages, operating expenses: plus or minus ten percent of forecast for the relevant period. Expense forecasts are generally more controllable than revenue forecasts; persistent variance indicates a structural error in the headcount plan or the cost driver assumptions.

These accuracy benchmarks apply to the period forecast prepared at least one month in advance. A forecast revised in the final week of a period is not the operative benchmark for that period.

### **COMMON DEFICIENCIES**

CD 2.1.1: The company maintains a revenue forecast based on a market size capture assumption. The forecast states that the company will capture a defined percentage of a defined market in each year of the projection. No operational mechanism is described. The forecast does not model the number of customers, the price per customer, the sales capacity required to acquire them, or the conversion rate from pipeline to close. The forecast cannot be interrogated at a driver level.

CD 2.1.2: The forecast is prepared once, at the beginning of the financial year, and is not updated. Actual performance diverges from the forecast over the year. No record is maintained of when assumptions changed or why. At the end of the year, the company cannot explain the variance between forecast and actual because no contemporaneous documentation of assumption changes exists.

CD 2.1.3: The company updates its forecast monthly but does not document changes to assumptions between versions. The current forecast reflects significantly different assumptions from the forecast prepared six months earlier. The basis for those changes cannot be established from the available records. Investors reviewing the forecast alongside earlier versions identify unexplained discontinuities.

CD 2.1.4: Personnel expenses in the forecast are modeled as a single aggregate figure rather than as a headcount plan. When actual payroll differs from forecast, the company cannot identify which roles drove the variance. The forecast does not distinguish between current headcount cost and the cost of planned future hires.

## **SECTION 2.2: THE UNIT ECONOMICS STANDARD**

### **PURPOSE**

The unit economics standard governs how a company calculates, documents, and tracks the financial performance of its fundamental unit of commercial activity: the customer relationship. Unit economics are the analytical foundation on which the viability of a business model is assessed. A company that cannot state its unit economics with methodological precision cannot demonstrate that its business model is viable at scale, regardless of its aggregate revenue trajectory.

Unit economics must be calculated from cohort-level data wherever retention and revenue data by customer vintage is available. Unit economics calculated from aggregate averages without cohort segmentation are acceptable only at Level 1 and must be labelled as average-based when reported in any investor document.

### **DEFINITIONS**

UNIT: The fundamental element of commercial activity from which the company's revenue model derives. For Recurring Revenue companies, the unit is the customer relationship. For Transactional Revenue companies, the unit is the transaction or the active user depending on the revenue model. For Project Revenue companies, the unit is the engagement. For B2B Enterprise companies, the unit is the account. For Consumer companies, the unit is the active user. For AI-Native companies, the unit is

defined by the company's pricing model, which may be the inference call, the active seat, or the output delivered; the definition must be documented and applied consistently.

**COHORT:** A group of customers or users acquired within the same defined period, typically one month or one quarter, tracked as a group over subsequent periods to measure retention, revenue development, and lifetime value.

**COHORT ANALYSIS:** The measurement of retention, revenue, and gross profit contribution for each cohort across the periods following acquisition. Cohort analysis is the required basis for lifetime value calculation at Level 2 and above.

**CUSTOMER ACQUISITION COST:** Defined in Book 0, Section 0.4. The fully loaded cost of acquiring one new customer or user, including all sales and marketing personnel costs, technology costs, agency fees, and directly attributable overhead, divided by the number of new customers or users acquired in the period. Customer acquisition cost must be calculated separately by acquisition channel at Level 2 and above.

**LIFETIME VALUE:** Defined in Book 0, Section 0.4. The total gross profit generated by a single customer or unit over the full duration of their relationship with the company, calculated from cohort-level retention and revenue data. Lifetime value is a gross profit figure. Lifetime value calculated from average retention assumptions rather than observed cohort data must be labelled as estimated lifetime value.

**LTV TO CAC RATIO:** Defined in Book 0, Section 0.4. The ratio of lifetime value to the fully loaded customer acquisition cost for the relevant acquisition channel or period. This ratio is the primary indicator of unit economics viability for Recurring Revenue and Consumer companies.

**PAYBACK PERIOD:** Defined in Book 0, Section 0.4. The number of months required for the cumulative gross profit contribution of a customer to equal the fully loaded cost of acquiring that customer. Payback period is measured in months and calculated at the cohort level at Level 2 and above.

**NET REVENUE RETENTION:** Defined in Book 0, Section 0.4. The percentage of recurring revenue retained from an existing customer cohort over a defined period, including expansion revenue and excluding revenue lost to churn and contraction.

**GROSS REVENUE RETENTION:** The percentage of recurring revenue retained from an existing customer cohort over a defined period, excluding expansion revenue and including only the revenue from customers who remained active. Gross revenue retention and net revenue retention are distinct measures and must not be reported interchangeably.

**CHURN RATE:** Defined in Book 0, Section 0.4. Customer churn rate and revenue churn rate must be calculated and reported as separate figures. They are not interchangeable. Revenue churn rate is the more operationally significant measure for Recurring Revenue companies; customer churn rate is the more operationally significant measure for Consumer companies.

**CONTRIBUTION MARGIN:** Revenue minus all variable costs directly attributable to the unit, including cost of goods sold and any variable sales costs directly attributable to serving that unit. Contribution margin is not equivalent to gross margin. A company that reports contribution margin must label it distinctly and must not present it as a substitute for gross margin in financial statements.

## **PRINCIPLES GOVERNING THIS SECTION**

Unit economics calculations must use the cost of goods sold definition established in Section 2.5 of this Book and governed by Book 1, Section 1.4. Lifetime value is a gross profit measure, not a revenue measure. Any unit economics calculation that uses revenue rather than gross profit as the numerator for lifetime value overstates the economic value of the customer relationship and does not satisfy this Standard.

The unit definition for AI-Native companies must be established at the point of first unit economics calculation and applied consistently across all subsequent periods. A change in unit definition constitutes a change in methodology and must be disclosed, with prior periods restated under the new definition or clearly labelled as calculated under a different definition.

## **COMPLIANCE CRITERIA**

### **Level 1**

2.2.L1.1: The company has identified and documented its unit as defined in the definitions of this section.

2.2.L1.2: The company calculates and records its average customer acquisition cost for each completed quarter using a consistent methodology.

2.2.L1.3: The company calculates and records an estimated lifetime value figure for its unit, with the methodology used to derive it documented.

2.2.L1.4: The company calculates its LTV to CAC ratio for each completed quarter.

2.2.L1.5: The company calculates its churn rate, distinguishing customer churn rate from revenue churn rate, for each completed month.

### **Level 2**

2.2.L2.1: The company calculates customer acquisition cost separately for each acquisition channel through which it has acquired more than ten percent of customers in any quarter in the preceding twelve months.

2.2.L2.2: The company maintains cohort analysis for all customer cohorts acquired in the preceding twenty-four months, tracking revenue and gross profit contribution for each cohort by period.

2.2.L2.3: Lifetime value is calculated from cohort-level retention and gross profit data, not from average assumptions; where cohort data is insufficient to calculate a full lifetime value, the company states the observation period used and labels the figure as a partial cohort lifetime value.

2.2.L2.4: The company calculates payback period at the cohort level for each acquisition channel, using fully loaded customer acquisition cost and cohort gross profit contribution.

2.2.L2.5: Net revenue retention and gross revenue retention are calculated separately for each completed quarter and are documented as distinct figures.

2.2.L2.6: Unit economics figures used in investor materials are calculated under the same methodology as the figures maintained in the company's internal records; any difference between internally tracked figures and investor-presented figures is documented and explained.

2.2.L2.7: The company maintains a unit economics model that projects forward unit economics under the current cost structure and the current cohort retention profile, covering a minimum of twelve months.

### **Level 3**

2.2.L3.1: The company maintains cohort analysis for all customer cohorts from inception, segmented by acquisition channel and customer segment where the company serves more than one identifiable segment.

2.2.L3.2: The unit economics model projects forward figures under at minimum three sets of assumptions: current retention and cost structure; a scenario in which retention improves by a defined increment; and a scenario in which customer acquisition cost increases by a defined increment.

2.2.L3.3: The company tracks the evolution of unit economics over time, maintaining a record of LTV to CAC ratio, payback period, and net revenue retention by quarter for the preceding twelve quarters or since inception if shorter.

2.2.L3.4: The company can demonstrate, from its cohort data, whether unit economics are improving or deteriorating as the company scales, with documented analysis of the drivers of any identified trend.

### **BENCHMARKS**

The following benchmarks are segmented by company type. Benchmarks marked with an asterisk are directional indicators where market data is insufficient for precise quantification in this version.

### **RECURRING REVENUE COMPANIES**

LTV to CAC ratio:

- Below 1.0: the unit economics of the business model are not viable at the current cost structure and retention rate. The company is generating less gross profit from each customer over their lifetime than it costs to acquire them.
- to 2.0: marginal. Sustainable only with rapid improvement in retention or reduction in acquisition cost.
- 3.0 and above: the threshold widely observed among institutional investors as the minimum for a viable recurring revenue business at Growth Stage. Companies below 3.0 at Growth Stage face investor scrutiny of the business model.
- 5.0 and above: indicates strong unit economics. Observed in recurring revenue businesses with high net revenue retention and low churn.

Payback period:

- Under twelve months: strong. Indicates that acquisition cost is recovered quickly, reducing the capital intensity of growth.
- Twelve to eighteen months: acceptable at Growth Stage with strong net revenue retention.
- Over twenty-four months: requires examination of whether the business model is capital efficient at scale.

Net revenue retention:

- Below 90 percent: the customer base is contracting in revenue value. Growth in revenue requires new customer acquisition to offset the contraction from existing customers.
- 90 to 100 percent: stable. Existing customers are retained but not expanding materially.

- 100 to 120 percent: the existing customer base grows in revenue value without new customer additions. Frequently observed in B2B SaaS businesses with strong expansion revenue.
- Above 120 percent: high. Observed in companies with strong product-led growth and high expansion revenue from existing accounts.

## **TRANSACTIONAL REVENUE COMPANIES**

LTV to CAC calculation note: for transactional companies, lifetime value must account for purchase frequency and average transaction gross profit over the expected customer lifespan, not a single transaction gross profit. Applying a single-transaction gross profit as lifetime value and calculating LTV to CAC on that basis understates the true unit economics and does not satisfy this Standard.

Payback period norms: transactional companies typically exhibit shorter payback periods than recurring revenue companies because there is no annual contract to be renewed. A payback period of under six months is common in high-frequency transactional businesses.

## **CONSUMER COMPANIES**

LTV to CAC benchmarks for consumer companies vary materially by monetisation model and engagement frequency. Benchmarks for advertising-monetised consumer companies, in-app purchase companies, and subscription consumer companies are distinct and cannot be applied interchangeably. The company must document which monetisation model its LTV calculation reflects.

## **B2B ENTERPRISE COMPANIES**

Payback period: B2B Enterprise companies typically exhibit longer payback periods than other company types due to extended sales cycles and high customer acquisition costs. Payback periods of eighteen to thirty-six months are observed among Growth Stage B2B Enterprise companies. The viability of these payback periods depends on net revenue retention and contract duration; a company with a thirty-month payback period and a three-year minimum contract term may have more defensible unit economics than a company with a twelve-month payback period and monthly contracts.

LTV to CAC: the high contract values characteristic of B2B Enterprise businesses typically produce high lifetime values. LTV to CAC ratios above 5.0 are common. The relevant scrutiny in this company type is of the payback period and of customer concentration rather than of the LTV to CAC ratio itself.

## **DEEP TECH AND HARDWARE COMPANIES**

Unit economics for Deep Tech and Hardware companies must account for the cost of goods sold structure, which includes manufacturing costs, component costs, and direct fulfilment costs that are absent from software-based businesses. Gross margin benchmarks for hardware-inclusive products are materially lower than for pure software products. The specific benchmarks for this company type will be stated in v1.0 following practitioner review.

## **AI-NATIVE COMPANIES**

Unit economics for AI-Native companies require distinct treatment because the cost of goods sold structure is dominated by computational inference costs rather than headcount, and because the relationship between usage, cost, and revenue is non-linear in ways that are not adequately represented by traditional unit economics frameworks.

The unit definition for AI-Native companies must specify whether the unit is a discrete inference call, an output delivered, an active seat, or a usage volume metric. The definition must be chosen to reflect the pricing model of the business and applied consistently.

Cost of goods sold for AI-Native companies includes computational infrastructure costs attributable to delivering the product, including API costs, model hosting costs, and inference compute costs. These costs may scale non-linearly with usage; the unit economics model must capture the relationship between usage volume and marginal inference cost, which may decline at scale as model efficiency improves or increase at scale as complex use cases consume more compute.

LTV to CAC norms for AI-Native companies are not yet established with sufficient market data to state as benchmarks in Beta v0.5. The company must calculate and track LTV to CAC using the methodology stated in this section and document the assumptions underlying its inference cost projections. These benchmarks will be stated in v1.0.

Payback period norms for AI-Native companies: where customer acquisition costs are low due to product-led growth and self-service distribution, payback periods may be significantly shorter than in sales-led businesses. Where enterprise sales motions are required, payback periods may extend. The company must calculate payback period using actual fully loaded acquisition costs and must not exclude sales costs from the calculation on the basis that they are viewed as a fixed cost.

## **COMMON DEFICIENCIES**

CD 2.2.1: The company calculates lifetime value using average monthly revenue per customer multiplied by an assumed average customer lifespan, without reference to cohort data. The assumed lifespan is based on management judgment rather than observed retention. The resulting lifetime value is higher than the value that would be produced by a cohort-based calculation using actual retention data. The figure is used in investor materials without disclosure of the methodology or its limitations.

CD 2.2.2: The company calculates its LTV to CAC ratio using a lifetime value figure based on revenue rather than gross profit. The ratio appears favourable. Applying gross profit rather than revenue as the numerator, using the actual gross margin of the business, produces a ratio that is materially lower and reflects the actual economic return of the customer relationship.

CD 2.2.3: The company calculates a single blended customer acquisition cost across all channels. Channel-level acquisition costs differ materially; one channel acquires customers at three times the cost of another. The blended figure obscures the variation. Investment decisions about channel allocation are made without visibility into channel-level unit economics.

CD 2.2.4: The company reports net revenue retention as customer retention. The two figures are different. The company retains eighty-five percent of customers by count but generates one hundred and five percent of prior year revenue from retained customers due to expansion. Reporting the eighty-five percent figure as net revenue retention misrepresents the revenue dynamics of the existing customer base.

CD 2.2.5: For an AI-Native company: the cost of goods sold used in the unit economics calculation excludes inference compute costs on the basis that infrastructure is treated as a fixed overhead. The resulting gross margin and unit economics figures do not reflect the actual cost of serving each customer or processing each unit of output. Inference costs that scale with usage are variable costs and must be included in cost of goods sold.

## **SECTION 2.3: THE SCENARIO AND SENSITIVITY STANDARD**

### **PURPOSE**

The scenario and sensitivity standard governs the construction, documentation, and maintenance of multiple-outcome financial projections. A single-case financial forecast does not meet the requirements of this Standard at Level 2 or above because it does not represent the range of plausible outcomes available to any early-stage company operating under material uncertainty. The scenario architecture required by this Standard is not a presentational device for investor meetings. It is an operational tool for decision-making under uncertainty.

### **DEFINITIONS**

**SCENARIO ANALYSIS:** The construction of multiple discrete financial projections, each based on a coherent and internally consistent set of assumptions, representing distinct views of how the company may develop. Each scenario must be internally consistent: if a scenario assumes lower customer acquisition than the base case, it must also assume lower acquisition costs, lower headcount in sales, and lower marketing spend. A scenario that changes revenue assumptions without changing the corresponding cost assumptions is not internally consistent and does not satisfy this Standard.

**BASE CASE:** The scenario that represents the most likely outcome given current operational data and forward assumptions. The base case is not an optimistic projection. It is the probability-weighted central estimate of the company's forward performance.

**UPSIDE CASE:** A scenario representing a coherent set of conditions under which performance exceeds the base case. The upside case must be plausible, not aspirational. It must be based on identifiable operational conditions, such as an improvement in conversion rate from a specific initiative, a reduction in churn from a specific product change, or an acceleration in market adoption. An upside case that cannot be connected to an operational mechanism is not a scenario; it is a wish.

**DOWNSIDE CASE:** A scenario representing a coherent set of conditions under which performance falls below the base case. The downside case must be plausible and must reflect risks that management has identified as material. A downside case that reduces all revenue assumptions by a uniform percentage without identifying the operational mechanism producing the reduction is not internally consistent.

**SENSITIVITY ANALYSIS:** An analytical technique that measures how a defined financial output changes when a single input variable is modified while all other variables are held constant. Sensitivity analysis identifies which driver assumptions exert the greatest influence on defined outcomes, including cash runway, gross margin, and LTV to CAC ratio.

**SCENARIO ARCHITECTURE:** The full set of scenarios maintained by a company, including the base case, the upside case, and the downside case at minimum, together with the documentation of the assumptions differentiating each scenario from the others.

**MONTE CARLO SIMULATION:** A probabilistic modeling technique in which key driver assumptions are treated as probability distributions rather than point estimates, and multiple simulations are run drawing values randomly from those distributions to produce a probability distribution of outcomes. Monte Carlo simulation is a Level 3 requirement for companies at Growth Stage and above.

**PROBABILITY DISTRIBUTION:** The range of values assigned to a driver assumption in a Monte Carlo simulation, defined by a central estimate and a standard deviation or by stated minimum and maximum

bounds with an assumed distribution shape. Every probability distribution used in a Monte Carlo simulation must be documented with the basis for its parameters.

## **PRINCIPLES GOVERNING THIS SECTION**

The scenario architecture must be constructed before any fundraising process begins. A scenario architecture created under the time pressure of an investor conversation and not previously maintained as an operational tool does not reflect the ongoing financial management requirement of this Standard.

Each scenario in the architecture must be internally consistent. A competent financial professional reviewing the scenario must be able to confirm that each assumption in the scenario is logically connected to the others within the same scenario.

## **COMPLIANCE CRITERIA**

### **Level 1**

2.3.L1.1: The company maintains a base case financial forecast and a downside case financial forecast covering the same forecast period.

2.3.L1.2: The key assumptions that differentiate the downside case from the base case are documented for each scenario.

2.3.L1.3: Each scenario is internally consistent: a change in revenue assumptions in a scenario is accompanied by corresponding changes in the cost and headcount assumptions that are operationally dependent on that revenue assumption.

### **Level 2**

2.3.L2.1: The company maintains a scenario architecture comprising at minimum a base case, an upside case, and a downside case, each covering the full forecast period required under Section 2.1.

2.3.L2.2: For each scenario, every assumption that differs from the base case is documented with the operational mechanism that produces the difference.

2.3.L2.3: The company maintains a sensitivity table for at minimum the following drivers: revenue growth rate; gross margin; customer acquisition cost; churn rate or equivalent retention metric; and the resulting impact on cash runway and LTV to CAC ratio.

2.3.L2.4: The scenario architecture is updated when a material change in operating conditions or management assumptions occurs, and each update is documented with the date and the reason for the change.

2.3.L2.5: The downside case includes a scenario in which the primary revenue growth assumption is reduced by at least thirty percent from the base case, with all operationally dependent cost assumptions adjusted accordingly.

### **Level 3**

2.3.L3.1: The company maintains a Monte Carlo simulation model covering the forecast period, treating at minimum five key driver assumptions as probability distributions rather than point estimates.

2.3.L3.2: Each probability distribution in the Monte Carlo model is documented with the central estimate, the stated bounds, and the basis for both, derived from historical company data, external market reference, or documented management judgment.

2.3.L3.3: The Monte Carlo model produces a probability distribution of cash runway outcomes, allowing the company to state the probability that runway exceeds defined thresholds under current assumptions.

2.3.L3.4: The company reviews the output of the Monte Carlo model quarterly and updates the driver distributions where actual results provide new information about the plausible range of future outcomes.

2.3.L3.5: The results of the Monte Carlo model are presented to the board or equivalent governing body at minimum twice per year as part of the financial reporting package governed by Book 1, Section 1.3.

## **BENCHMARKS**

Scenario analysis benchmarks by stage:

Pre-Revenue: base case and downside case minimum

Early Revenue: base case, upside case, and downside case minimum

Growth Stage: full scenario architecture plus sensitivity table minimum; Monte Carlo at Level 3

Scale Stage: full scenario architecture plus sensitivity table plus Monte Carlo minimum

Sensitivity table driver coverage:

The following drivers must be included in the sensitivity table at Level 2 and above, with their impact on cash runway and on the primary unit economics ratio calculated for each:

- For Recurring Revenue companies: monthly recurring revenue growth rate; monthly churn rate; customer acquisition cost; gross margin percentage.
- For Transactional Revenue companies: transaction volume; average transaction gross profit; customer acquisition cost; repeat purchase rate.
- For B2B Enterprise companies: average contract value; sales cycle duration; win rate from qualified pipeline; annual churn rate.
- For Consumer companies: daily or monthly active user count; average revenue per active user; user acquisition cost; retention rate at ninety days.
- For AI-Native companies: usage volume per active unit; inference cost per unit of output; customer acquisition cost; gross margin at current and projected scale. The sensitivity of gross margin to inference cost at different usage volumes must be explicitly modeled because the non-linear relationship between compute cost and usage is the primary financial risk unique to this company type.
- For Deep Tech and Hardware companies: unit production cost; average selling price; production volume; gross margin at different volume levels.

## **COMMON DEFICIENCIES**

CD 2.3.1: The company maintains a single financial forecast and labels it the base case. No upside or downside scenario exists. When investors ask to see the downside case, a new version is produced under time pressure by reducing revenue assumptions by a uniform percentage. The downside case does not adjust cost assumptions to reflect the lower revenue. The resulting downside scenario overstates the downside cash burn because it assumes the company continues spending at base case levels while revenue declines.

CD 2.3.2: The company maintains three scenarios but the assumptions differentiating them are not documented. Reviewing the three models, the only visible difference is in the top-line revenue figure. The underlying driver assumptions are identical across all three scenarios. The upside and downside cases are not scenarios; they are the base case revenue projection with a percentage applied to the top line.

CD 2.3.3: The sensitivity table exists but is constructed by changing one variable across a range while holding all others constant, including cost variables that are operationally dependent on the revenue variable being changed. A sensitivity analysis that reduces revenue by fifty percent while holding sales headcount constant does not produce a plausible operational scenario and does not satisfy this Standard.

CD 2.3.4: For Growth Stage companies that have not yet implemented Monte Carlo simulation: the company has not assessed the probability distribution of its cash runway under current driver assumptions. Management holds a point estimate of runway derived from the base case forecast. The actual distribution of runway outcomes, given the variance observed in historical driver performance, is materially wider than the point estimate implies.

## **SECTION 2.4: THE GROWTH MODELING STANDARD**

### **PURPOSE**

The growth modeling standard governs how a company models its revenue growth mechanism, its sales capacity, and the operational assumptions that connect commercial activity to financial outcomes. A growth model is not a revenue forecast. A revenue forecast states projected revenue figures. A growth model describes the operational mechanism by which those figures will be reached and demonstrates that the company has the capacity to execute the mechanism.

### **DEFINITIONS**

**GROWTH MODEL:** A financial model that derives revenue projections from explicit modeling of the commercial activities required to generate that revenue, including sales capacity, marketing investment, pipeline conversion, and customer retention.

**SALES CAPACITY MODEL:** A component of the growth model that quantifies the revenue-generating output of the sales organisation, derived from the number of quota-carrying sales personnel, their individual revenue quota, their average quota attainment rate, and the ramp time required for new sales personnel to reach full productivity.

**QUOTA ATTAINMENT RATE:** The percentage of their assigned revenue quota achieved on average by quota-carrying sales personnel over a defined period. Quota attainment rate must be calculated from actual historical performance data at Level 2 and above.

**PIPELINE MODEL:** A model of the company's revenue pipeline that tracks the volume and value of opportunities at each stage of the sales process, the conversion rate between stages, and the average time required to progress from one stage to the next.

**RAMP PERIOD:** The number of months required for a newly hired quota-carrying sales employee to reach full productivity, defined as the period from hire date to the first month in which the employee achieves their full assigned quota on a sustained basis.

**PRODUCT-LED GROWTH:** A growth mechanism in which the product itself drives user acquisition, expansion, and retention without primary reliance on a sales organisation. Companies whose growth mechanism is primarily product-led must model growth through user acquisition rates, activation rates, and expansion from existing users rather than through a sales capacity model.

**MAGIC NUMBER:** A metric calculated as net new annualised recurring revenue divided by the sales and marketing expenditure in the prior quarter, used to assess the efficiency of the growth investment. A Magic Number above 0.75 is commonly observed among efficient Growth Stage recurring revenue companies. A Magic Number below 0.5 indicates that the growth investment is producing returns below the threshold commonly observed in institutionally funded businesses.

**RULE-OF-40:** A heuristic applied to Growth Stage and Scale Stage recurring revenue companies, calculated as the sum of the annual revenue growth rate percentage and the operating profit margin percentage. A combined figure of forty or above indicates that the company is balancing growth and profitability at a level commonly expected by growth equity investors. This heuristic is not a requirement of this Standard; it is a benchmark for contextualising the trade-off between growth investment and margin.

## **PRINCIPLES GOVERNING THIS SECTION**

The growth model must demonstrate that the company has the operational capacity to execute its revenue projections. A revenue projection that assumes growth rates that the company's current sales capacity, marketing investment, and conversion performance cannot support does not meet this Standard at Level 2 or above, regardless of whether the projection is produced using a bottom-up methodology.

For companies whose primary growth mechanism is product-led rather than sales-led, the sales capacity model is replaced by a user acquisition model that derives active user counts from acquisition channel performance, activation rates, and retention rates.

## **COMPLIANCE CRITERIA**

### **Level 1**

2.4.L1.1: The company has documented its primary growth mechanism, identifying whether revenue growth is driven primarily by a sales organisation, by marketing investment, by product-led acquisition, or by a combination.

2.4.L1.2: The company's revenue forecast is connected to at least one documented operational assumption about the commercial activity that will produce the projected revenue.

2.4.L1.3: The company has documented the conversion rate it observes between its primary lead or pipeline stage and closed revenue, based on historical data from the preceding six months or since inception if shorter.

## Level 2

2.4.L2.1: For companies with a sales-led growth mechanism: the company maintains a sales capacity model that derives projected revenue from the number of quota-carrying sales personnel, their assigned quota, their observed quota attainment rate, and their ramp period.

2.4.L2.2: For companies with a product-led growth mechanism: the company maintains a user acquisition model that derives projected active user counts from documented acquisition channel performance metrics, observed activation rates, and observed thirty-day and ninety-day retention rates.

2.4.L2.3: The company maintains a pipeline model that tracks opportunity volume and value at each stage of the sales or acquisition process, with documented conversion rates between stages based on historical performance.

2.4.L2.4: The company's revenue forecast does not project revenue growth rates that exceed what the documented sales capacity model or user acquisition model can support under the stated assumptions.

2.4.L2.5: The company calculates its Magic Number for each completed quarter and maintains a record of this metric for the preceding four quarters.

2.4.L2.6: The headcount plan underlying the growth model documents the hire date, role, quota or equivalent productivity measure, and ramp period for each planned sales or growth hire in the forecast period.

## Level 3

2.4.L3.1: The sales capacity model or user acquisition model is integrated with the three-statement model such that a change in a growth capacity assumption updates revenue projections, headcount costs, and marketing spend assumptions without manual recalculation.

2.4.L3.2: The company maintains a documented analysis of the leading indicators of its growth mechanism, identifying the metrics that, when observed two to three months in advance, are predictive of revenue outcomes, and tracking those metrics monthly.

2.4.L3.3: The company calculates its Rule of 40 figure for each completed quarter and maintains a record for the preceding four quarters, where the company has reached Growth Stage as defined in Book 0, Section 0.5.

2.4.L3.4: The growth model is stress-tested under a scenario in which the planned sales hires for the next two quarters are delayed by one quarter, with the revenue impact documented.

## BENCHMARKS

Sales capacity benchmarks by company type and stage:

RECURRING REVENUE, Growth Stage: quota attainment rates below sixty percent across the sales team indicate either that quotas are set above achievable levels or that the sales team is underperforming. Quota attainment rates above one hundred percent across the team indicate that quotas may be set below achievable levels. Both conditions warrant examination of quota-setting methodology.

RECURRING REVENUE, Growth Stage: ramp periods of three to six months are typical for inside sales roles. Six to nine months are typical for mid-market account executives. Nine to twelve months are typical for enterprise sales roles. A growth model that assumes full productivity from new sales hires from the month of hire does not satisfy this Standard.

Magic Number benchmarks:

Above 0.75: efficient growth investment. The company is generating new recurring revenue efficiently relative to its sales and marketing spend.

0.5 to 0.75: moderate efficiency. Acceptable with improving trend.

Below 0.5: the efficiency of growth investment requires examination.

AI-Native companies with product-led growth: the growth model for AI-Native companies whose distribution is primarily self-service or product-led must model the relationship between marketing investment, inbound discovery, free-to-paid conversion rates, and expansion from initial usage. The traditional sales capacity model does not apply. Usage-based revenue models require a growth model that projects usage volumes, conversion from free to paid tiers, and revenue per unit of usage at different volume levels.

## **COMMON DEFICIENCIES**

CD 2.4.1: The revenue forecast projects growth rates that imply a doubling of the customer base within twelve months. The sales capacity model shows a team of two quota-carrying sales personnel. At their current quota attainment rate, the two-person team can acquire approximately forty percent of the customers implied by the revenue forecast. The remaining sixty percent of the projected customer acquisition has no identified source in the growth model. The forecast and the growth model are not internally consistent.

CD 2.4.2: The company models new sales hires as contributing full quota revenue from their first month of employment. No ramp period is assumed. The headcount plan adds four sales hires in the first quarter of the forecast. The revenue model reflects full revenue contribution from all four from month one. Actual ramp periods in the company's historical data are seven months. The forecast overstates revenue for the first seven months following each hire.

CD 2.4.3: The pipeline model tracks the number of opportunities at each stage but not the value of those opportunities. Conversion rates are calculated by count rather than by value. The conversion rate from first meeting to proposal is sixty percent by count but the opportunities that convert are systematically smaller in average contract value than those that do not. The revenue implied by the pipeline model overstates actual expected revenue.

CD 2.4.4: The company does not calculate its Magic Number. Marketing and sales expenditure increases each quarter as the company invests in growth. Revenue growth is also increasing. The company has not assessed whether the incremental revenue generated is proportional to the incremental investment. Without the Magic Number calculation, the efficiency of the growth investment is not visible.

## **SECTION 2.5: THE COST STRUCTURE STANDARD**

### **PURPOSE**

The cost structure standard governs how a company models, categorises, and monitors its cost base. A cost model is not a list of expected expenses. It is a structured representation of the relationship between

the company's operational activities and the costs those activities generate, built from documented drivers and organised to support gross margin analysis, departmental performance assessment, and capital efficiency measurement.

## **DEFINITIONS**

**FIXED COST:** A cost that does not change in response to changes in the volume of output or revenue within a defined range of operational scale. Office rent, base salaries, and software licence fees with flat pricing are examples of fixed costs. Fixed costs create operating leverage: as revenue grows, fixed costs remain constant and an increasing proportion of revenue is available for operating profit.

**VARIABLE COST:** A cost that changes proportionally with changes in the volume of output or revenue. Inference compute costs in AI-Native businesses, transaction processing fees in payment businesses, and raw material costs in hardware businesses are examples of variable costs. Variable costs must be modeled as a rate per unit of output or revenue, not as a fixed absolute amount.

**STEPPED COST:** A cost that remains fixed within a defined range of operational scale but increases in discrete steps when the company's scale crosses a defined threshold. Customer support headcount that must increase at defined customer count thresholds is an example of a stepped cost. Stepped costs must be modeled explicitly in the cost structure rather than treated as either purely fixed or purely variable.

**OPERATING LEVERAGE:** The relationship between the rate of revenue growth and the rate of operating cost growth. A company with high operating leverage generates a disproportionate increase in operating profit as revenue grows, because a significant portion of its cost base is fixed. Operating leverage must be demonstrated in the cost model by showing how gross margin and operating margin evolve at different revenue levels.

**BURN MULTIPLE:** Net cash consumed in a period divided by net new annualised recurring revenue generated in that period. Burn multiple measures the cash cost of generating each incremental dollar of new recurring revenue. A burn multiple below one indicates that the company is generating more new recurring revenue than cash it is consuming. A burn multiple above two at Growth Stage indicates that the growth investment is capital-intensive relative to the new revenue being generated.

**HEADCOUNT MODEL:** A component of the cost structure model that lists every current and planned role, the department it belongs to, the hire date or planned hire date, the fully loaded cost including salary, employer taxes, benefits, and equipment, and the functional category to which the cost is allocated.

**FULLY LOADED COST:** The total cost of employing one person for one period, including salary, employer-side taxes and social contributions, healthcare and other benefits, equipment, software licenses required for the role, and an allocated share of facilities costs. Salary alone does not constitute fully loaded cost. A cost model that uses salary as a proxy for headcount cost understates total personnel expense.

## **PRINCIPLES GOVERNING THIS SECTION**

The cost classification policy established in Book 1, Section 1.4 governs the categorisation of costs in this section. Cost of goods sold is defined in Section 2.1 of this Book. Operating expenses are the costs that are not cost of goods sold. These definitions apply consistently across the cost structure model.

Variable costs must be modeled as rates, not as fixed amounts, because their purpose in the model is to show how total cost changes as the business scales. A variable cost that is modeled as a fixed absolute amount does not allow the model to project cost behavior accurately at different revenue levels.

## **COMPLIANCE CRITERIA**

### **Level 1**

2.5.L1.1: The company maintains a cost model that distinguishes cost of goods sold from operating expenses for each completed period.

2.5.L1.2: The company calculates gross margin for each completed period.

2.5.L1.3: The company identifies its primary cost categories and maintains actual expenditure data by category for each completed period.

### **Level 2**

2.5.L2.1: The company maintains a cost structure model that categorises all costs as fixed, variable, or stepped, with the classification of each material cost item documented.

2.5.L2.2: Variable costs are modeled as rates per unit of output or per unit of revenue, not as fixed absolute amounts; the rate for each variable cost is documented and sourced from historical data or documented market reference.

2.5.L2.3: The company maintains a headcount model listing every current and planned role, the department allocation, the hire date, and the fully loaded cost per role.

2.5.L2.4: The cost structure model projects how gross margin, operating expenses, and operating margin evolve at the revenue levels projected in the base case forecast, the upside case, and the downside case.

2.5.L2.5: The company calculates its burn multiple for each completed quarter and maintains a record for the preceding four quarters.

2.5.L2.6: The cost structure model is updated monthly to reflect actual cost performance; any cost that has exceeded its modeled level by more than fifteen percent for two consecutive months is reviewed and the model is updated with a revised rate or fixed amount and the basis for the revision is documented.

### **Level 3**

2.5.L3.1: The cost structure model is integrated with the three-statement model such that a change in any cost driver updates the income statement, cash flow statement, and balance sheet without manual recalculation.

2.5.L3.2: The company models operating leverage explicitly, producing a schedule that shows gross margin percentage, operating expense as a percentage of revenue, and operating margin at revenue levels representing fifty percent, one hundred percent, and two hundred percent of current annualised revenue.

2.5.L3.3: The headcount model projects the planned headcount for each department twelve months forward, with the financial impact of each planned hire on gross margin, departmental operating expense, and total burn modeled at the time the hire is planned.

2.5.L3.4: The cost structure model includes a scenario in which variable costs per unit of output increase by twenty percent from current levels, with the impact on gross margin and cash runway documented.

## **BENCHMARKS**

Gross margin benchmarks by company type:

**RECURRING REVENUE** (pure software SaaS): gross margins of sixty-five to eighty-five percent are typical at Growth Stage and above. Gross margins below fifty percent indicate that cost of goods sold is elevated relative to peers, which may reflect inefficiency in infrastructure provisioning, high customer success costs included in cost of goods sold, or a service-heavy delivery model.

**TRANSACTIONAL REVENUE**: gross margins vary materially by transaction cost structure. Payment processing companies with high interchange costs typically operate at gross margins of forty to sixty percent. Marketplace businesses collecting a net commission typically operate at higher gross margins because the cost of goods sold is lower. The company must define its cost of goods sold clearly before gross margin benchmarks are applicable.

**PROJECT REVENUE**: gross margins of thirty to sixty percent are typical, reflecting the direct labour cost of delivering professional services. Gross margins above sixty percent in a project revenue business typically indicate that delivery is increasingly systematised and scalable. Gross margins below thirty percent indicate that delivery costs are high relative to the fees charged.

**DEEP TECH AND HARDWARE**: gross margins of twenty to fifty percent are typical for hardware-inclusive products at scale. Early-stage hardware companies often operate at negative gross margins during initial production runs as unit costs are high before manufacturing scale is achieved. The cost model must distinguish pre-scale unit costs from projected scale unit costs.

**AI-NATIVE**: gross margin benchmarks for AI-Native companies are not yet established with sufficient market data to state as benchmarks in Beta v0.5. The defining characteristic of AI-Native gross margin dynamics is the relationship between inference compute cost and usage volume. As model efficiency improves and as usage volume increases, inference cost per output unit may decline; this improvement must be modeled explicitly rather than assumed. Gross margin for an AI-Native company at Growth Stage using high-cost frontier model inference may be materially lower than for the same company at scale if it has developed more efficient inference. The cost model must make this trajectory explicit. Benchmarks will be stated in v1.0.

**B2B ENTERPRISE**: gross margins of sixty to eighty percent are typical for software-led B2B Enterprise companies. Where professional services revenue is included in total revenue, the blended gross margin will be lower; the company must present gross margin separately for the software and services components.

**CONSUMER**: gross margin benchmarks vary by monetisation model. Advertising-monetised consumer businesses typically operate at high gross margins because the cost of serving an additional impression is low. Subscription consumer businesses operate similarly to Recurring Revenue software businesses. In-app purchase businesses must account for app store fees in cost of goods sold.

Burn multiple benchmarks by stage:

- Growth Stage: a burn multiple below two is the commonly observed threshold among efficiently growing venture-backed companies. A burn multiple above three at Growth Stage indicates that the company is consuming significantly more cash than it is generating in new recurring revenue.
- Scale Stage: a burn multiple below one is typical at Scale Stage for companies approaching profitability. A burn multiple consistently above two at Scale Stage indicates that the growth investment is not producing proportionate new revenue.

## **COMMON DEFICIENCIES**

CD 2.5.1: The company's cost model lists expected monthly expenses as fixed amounts for every line item, including costs that are operationally variable. Inference compute costs for an AI-Native company are modeled as a fixed monthly amount based on current usage. As usage grows with customer acquisition, actual compute costs grow proportionally but the model does not reflect this. The cost model systematically understates future cost of goods sold and overstates future gross margin.

CD 2.5.2: The headcount model uses salary figures rather than fully loaded cost. The cost model understates total personnel expense by the employer tax, benefits, equipment, and facilities costs attributable to each employee. For a company with thirty employees, the difference between salary cost and fully loaded cost commonly represents fifteen to twenty-five percent of total personnel expense. The model produces a gross margin and operating expense figure that is materially more favourable than the actual cost structure.

CD 2.5.3: The company has not calculated its burn multiple. It tracks total cash burn and total revenue growth separately but has not related them to assess the capital efficiency of growth. The company has invested materially in sales headcount over the preceding four quarters. Revenue has grown. The burn multiple calculation would reveal that new recurring revenue generated per pound of cash consumed has declined over the period, indicating declining growth efficiency. This trend is not visible from cash burn and revenue growth figures observed independently.

CD 2.5.4: The cost structure model does not include stepped costs. Customer support headcount is modeled as a fixed amount regardless of customer count. In practice, the company adds a support employee for each increment of fifty customers. At three hundred customers, actual support costs are sixty percent higher than modeled. The model consistently understates operating expenses in periods of strong customer acquisition.

## **STAGE AND LEVEL APPLICABILITY**

The compliance level expectations for each section of Book 2 are established in the Financial Infrastructure Maturity Model in Book 0, Section 0.8. The operative mappings for this Book are as follows.

Section 2.1: Forecasting Methodology

Pre-Incorporation: no requirement

Pre-Revenue: no requirement

Early Revenue: Level 1

Growth Stage: Level 2

Scale Stage: Level 2

Section 2.2: Unit Economics

Pre-Incorporation: no requirement

Pre-Revenue: no requirement

Early Revenue: Level 1

Growth Stage: Level 2

Scale Stage: Level 3

Section 2.3: Scenario and Sensitivity

Pre-Incorporation: no requirement

Pre-Revenue: no requirement

Early Revenue: Level 1

Growth Stage: Level 2

Scale Stage: Level 3

Section 2.4: Growth Modeling

Pre-Incorporation: no requirement

Pre-Revenue: no requirement

Early Revenue: Level 1

Growth Stage: Level 2

Scale Stage: Level 2

Section 2.5: Cost Structure

Pre-Incorporation: no requirement

Pre-Revenue: no requirement

Early Revenue: Level 1

Growth Stage: Level 2

Scale Stage: Level 3

Stage definitions are as established in Book 0, Section 0.5. Compliance level expectations reflect the Maturity Model established in Book 0, Section 0.8.

Feedback on Book 2 may be submitted to [standard@ffistandard.org](mailto:standard@ffistandard.org) with reference to the specific section and criterion number. Feedback received before the close of the beta review period will be considered in the preparation of v1.0.