The 19-Method Valuation Handbook — free sample

How one stock gets nineteen fair values — and how to read them

Equity Rank

2026

> **Free sample — the front matter and Part I of *The 19-Method Valuation Handbook* (2026 edition).** The complete book (Parts II–VIII and the appendices: every method chapter, the blend, the band, the margin of

> safety, seven dated walkthroughs, four post-mortems, the formula sheet, glossary, and knowledge check) is available at equity-rank.com/handbook.

Before you begin

About this book

This is a handbook, not a manifesto. It explains the nineteen valuation methods that the Equity Rank engine runs on roughly three thousand U.S.-listed stocks every trading day — what each method measures, the formula as the engine actually implements it, the inputs it needs, where it fails, and a worked example with real arithmetic. It then explains what happens *after* the nineteen numbers exist: how they are weighted by sector, trimmed, clustered, and blended into one model fair value estimate; why that estimate carries an uncertainty band; what "margin of safety" does and does not mean; and how to read the whole thing without fooling yourself.

Two fictional companies run through the book so that every method can be applied to the same inputs and the results compared side by side: **Northwind Tools Inc.**, a mature, profitable industrial, and **Helix Cloud Corp.**, a fast-growing, thin-margin software company. Neither exists. The numbers were chosen to be realistic and to make the methods' disagreements visible, which is the point. Smaller examples cover a REIT, a gold miner, a regulated utility, and a bank, because those are the sectors where general-purpose methods break and specialised ones take over. Part VIII reproduces

seven dated walkthroughs and four post-mortems on real companies from Equity Rank's education library, including three in which the model was wrong — a library's credibility rests on publishing misses with the same prominence as everything else.

You do not need to be a subscriber to use this book. Every formula can be worked by hand or in a spreadsheet, and Appendix A is a one-page formula sheet for exactly that.

How to read it

- **If you are new to valuation**, read Part I (Foundations) first. It defines fair value, the margin of safety, the discount rate, and the handful of inputs every method shares. The rest of the book assumes it.
- **If you already value stocks** and want to know what a specific method does, go straight to its chapter. Each method chapter has the same shape: *What it measures · The formula · Inputs · Where it applies · Where it breaks · Worked example · What the engine does with it.*
- **If you use Equity Rank** and want to understand the fair value on a stock page, Part VII is the chapter set for you: the blend, the band, the margin of safety, the scores, and the reading habit.
- Appendix C is a short knowledge check with answers.

What this book is not

It is not investment advice, and it does not tell you what to own. Leek Ventures LLC (dba Equity Rank) is not a registered investment adviser, broker-dealer, or financial planner under

federal or any state law. Every fair value figure in this book — whether from the fictional examples or from the dated real-company walkthroughs in Part VIII — is a *model estimate of a company's intrinsic worth under stated assumptions*. It is not a prediction, forecast, or target of where a security will trade. Fair value estimates describe what a model considers a business to be worth today under its assumptions; market prices are set by supply and demand and by factors no model captures, and they can differ from any estimate materially and indefinitely. All financial models contain assumptions that may be wrong. Past behaviour of any security, method, or model is not indicative of future results. You may lose money on any investment. Conduct your own due diligence and consult a licensed financial adviser, tax professional, or attorney before making any investment or financial decision.

Nothing in this book is personalised. It does not know your circumstances, goals, risk tolerance, tax situation, or time horizon, and it was written without any attempt to learn them.

A note on how it was written

The text of this book was generated with AI tools working from Equity Rank's published methodology, its valuation engine's formula reference, and its education library, and was then reviewed, corrected, and edited by a human before publication; the cover was likewise machine-generated. Amazon has been given the same disclosure. Where a figure in a real-company walkthrough is a model output, the date of the snapshot is stated; figures recompute daily and will differ today. The arithmetic in the fictional examples is reproducible from the inputs given in each chapter.

Conventions

- "Fair value" (FV) always means a *model estimate* per share unless the text says otherwise.
- "Price" means the market price per share at the time of the example.
- "Margin of safety" (MoS) is defined in Chapter 2; the book uses the sign convention *positive = price below the estimate*.
- Sector multiples, discount rates, and growth rates in the fictional examples are illustrative. The engine's live values change with the market and are recalibrated against the live universe; the chapters say where.
- Formulas are written in plain text so that they copy cleanly into a spreadsheet: `FV = EPS × P/E_sector`.

Part I — Foundations

Chapter 1 — What a fair value is (and is not)

Every valuation method in this book answers one question: *given what this business earns, owns, pays out, or is likely to generate, what would a share of it be worth if the market priced it on those fundamentals?* The answer is a **fair value estimate** — a number per share, produced by a formula from inputs you can see.

Three properties of that sentence do almost all the work in the rest of the book.

It is an estimate of worth, not a statement about price. Price is what the last trade happened at. Worth is what a business's economics imply. They are different quantities measured by different instruments, and the gap between them — positive or negative — can persist for years. A fair value estimate that sits 30% above the market price is not a claim that the price will rise 30%; it is a claim that, under the method's assumptions, the economics justify a higher number than the market is currently paying. What the market does next depends on whether the market comes to share those assumptions, and on a great deal that no valuation model contains.

It is conditional on assumptions. Change the discount rate by one percentage point and a discounted-cash-flow value moves by double digits. Change the sector multiple and a price-to-earnings value moves proportionally. A fair value without its

assumptions attached is a number without a unit. Every chapter in Part II through Part VI therefore states the assumptions explicitly, and Part VII explains how the engine publishes them alongside every estimate.

It comes from a formula you can inspect. This is the difference between a model estimate and an opinion. An opinion can be right for reasons you cannot check. A formula is right or wrong for reasons you can — and when it is wrong, you can usually say *which input* was wrong. That inspectability is the whole reason to learn nineteen methods instead of memorising one number.

Intrinsic versus relative

Valuation methods fall into two families, and knowing which family a method belongs to tells you most of what you need to know about when it will mislead you.

Intrinsic methods derive value from the business's own cash flows, earnings power, dividends, or assets. Discounted cash flow, the dividend discount model, earnings power value, residual income, and net asset value are intrinsic. They need no comparison to other companies. Their weakness is that they depend heavily on forward assumptions — growth rates, discount rates, terminal values — which are exactly the things nobody knows.

Relative methods derive value by applying a multiple observed in comparable companies to one of the business's own metrics: earnings, book value, sales, EBITDA, cash flow. P/E, P/B, P/S, EV/EBITDA, and their cousins are relative. Their strength is that they are anchored to what the market actually pays for similar businesses today. Their weakness is the same

thing: if the whole sector is mispriced, a relative method inherits the mispricing faithfully.

Neither family is superior. A good valuation process runs both and pays attention when they disagree. That disagreement is information, and Part VII is largely about how to read it.

The unit of analysis: per share

Everything in this book is expressed per share because that is the unit a market price comes in. Enterprise-level methods (the EV family, EPV, EVA) compute a value for the whole firm and then bridge to equity by subtracting net debt and dividing by the share count. Chapter 4 covers the bridge. The share count matters more than it looks: use the wrong one (basic versus diluted, or a pre-split figure) and every per-share value in the book is off by the same factor. The engine's own history includes a share-count trap worth remembering — data vendors sometimes report a single share class or ignore depositary-receipt ratios — which is why the engine carries guards for exactly this, and why a human reading a filing should check the diluted count on the face of the income statement.

What "nineteen methods" means

Equity Rank's engine implements nineteen named valuation methods. It does not run all nineteen on every stock. A method fires only when its inputs exist and are meaningful: the dividend discount model needs a dividend; funds-from-operations needs a REIT; net asset value needs reserves. On a typical stock, between eight and fifteen methods fire. When a method cannot fire, the engine records that it *abstained* and why, and redistributes its weight to the methods that did. Abstention with a stated reason is a design choice you will see

throughout this book — a blend that quietly changes its ingredients is harder to trust than one that tells you its coverage shrank.

Beyond the nineteen, the engine carries a further set of specialised or experimental methods — residual income, economic value added, CAPE, bank price-to-tangible-book, a two-stage dividend model, AFFO for REITs, an options-implied forward, and a risk-adjusted NPV for clinical-stage biotech. Part VI covers them. They are real, they run, and some of them carry consensus weight in the sectors they were built for; they are simply not the headline nineteen.

Chapter 2 — Price, value, and the margin of safety

The margin of safety is the oldest idea in value investing and the most misread number on any stock page. It measures the distance between price and a fair value estimate. It does not measure the probability that the distance closes.

The two conventions

There are two ways to express the gap, and they give different percentages for the same pair of numbers. The engine computes both and is explicit about which it is showing.

The Graham convention divides by fair value:

```
MoS (Graham) = (FV − Price) / FV
```

This reads as “the discount to fair value”. A $50 stock against a $53.78 estimate has a Graham MoS of (53.78 − 50) / 53.78 =

7.0%. It is bounded: it cannot exceed 100% (price cannot go below zero) and it can be very negative. Benjamin Graham's famous insistence on a "margin of safety" used this framing — pay a price well below the estimated value so that errors in the estimate are absorbed before capital is impaired.

The return-to-fair-value convention divides by price:

```
MoS (return-to-FV) = (FV − Price) / Price
```

This reads as "the differential between price and the estimate, measured from where you are". The same $50 / $53.78 pair gives (53.78 − 50) / 50 = **7.6%**. The engine clamps this version to the range [−100%, +250%] because a formula can produce absurd values on tiny prices, and an unbounded number on a stock page invites exactly the misreading this chapter is about.

Both are descriptions of a gap. Neither is a forecast. A return-to-FV of +7.6% is emphatically *not* a statement that the stock will return 7.6% — it says that if, and only if, the price moved to the estimate, the move would be 7.6%. Whether it moves, when, and whether the estimate was right in the first place are three separate questions the number does not answer.

Why the sign matters more than the size

A positive margin of safety means the price sits below the estimate; a negative one means it sits above. At small magnitudes — single digits either way — the sign is usually inside the estimate's own uncertainty and should be read as "roughly fairly priced under these assumptions". At large magnitudes the sign becomes a prompt for a specific question: *what would have to be true for the market to be right and the model wrong?* Part VIII's walkthroughs show this question

being asked of a −41% reading on a bank and a +91% reading on a software company, with very different answers.

What a margin of safety cannot see

The margin of safety is computed from a fair value estimate and a price. Anything the estimate does not contain, the margin of safety does not contain either. It cannot see pipeline optionality in a biotech, a pending regulatory decision, fraud, or a data error in the price feed. One of the post-mortems in Part VIII is about a stock that showed an 83% "margin of safety" on a dead ticker whose price feed had kept serving stale prints after the company was acquired. The number was arithmetically correct and completely meaningless. The first question to ask of any extreme margin of safety is whether the inputs are even real.

The required margin of safety

The engine does not treat every margin of safety alike. Because estimates for some companies are inherently less certain than for others — fewer methods fired, wider disagreement, thinner data, more cyclical earnings — it computes a *required* margin of safety that scales with that uncertainty. A stock whose methods agree tightly needs less of a discount to be called attractively priced than one whose methods span a factor of three. Chapter 27 explains the mechanics. The principle is Graham's own: the margin of safety you need is a function of how sure you can be, and the estimate deserves the least trust exactly where it is widest.

Chapter 3 — Why nineteen methods instead of one

If one method were reliably best, this book would be one chapter long. None is, and the reason is structural rather than a failure of effort.

Every method is right about something

Each valuation method is an answer to a specific question:

Method family	The question it answers well	The question it answers badly
P/E, forward P/E	What does the market pay for a dollar of earnings like these?	What if earnings are temporarily distorted or negative?
P/B, justified P/B	What are the net assets worth, and does the return on them justify a premium?	What if value lives in intangibles the balance sheet does not carry?
P/S, EV/Sales	What is a dollar of revenue worth when profit is deliberately suppressed?	What if the revenue never converts to profit?
EV/EBITDA, EV/EBIT	What is the operating business worth regardless of how it is financed?	What if depreciation is a real economic cost (EBITDA) or the company is a bank (no meaningful EV)?
P/FCF, EV/FCF	What is real cash generation worth?	What if cash flow is lumpy, or flattered by under-investment, or ignores stock compensation?
PEG	How much growth premium is the price paying for?	What if growth is low, negative, or about to decelerate?
Graham Number	What is a conservative floor under a profitable,	Anything asset-light.

Method family	The question it answers well	The question it answers badly
	asset-backed business?	
DCF family	What are the future cash flows worth today under explicit assumptions?	Whenever the assumptions are wrong — and the terminal value is most of the answer.
DDM	What is the dividend stream worth to a holder?	Non-payers, and anyone who changes the discount rate by a point.
EPV	What is the business worth if it never grows again?	Growth businesses, where the answer is structurally low.
FFO, NAV, mid-cycle P/E	What is this REIT / miner / cyclical worth on the metric its sector actually uses?	Everything outside that sector.

Read down the right-hand column and you will notice that each method's blind spot is another method's strength. That is the whole argument for running many: not that the average of nineteen wrong answers is right, but that the *pattern* of disagreement tells you which assumptions are doing the work.

Dispersion is information

When nine methods applied to the same company produce values between $22 and $64, a naive reaction is that the

methods are useless. The more useful reaction is to ask why. In the Northwind example that runs through this book, the low values come from methods that capitalise only the current dividend or assume zero growth; the high values come from methods that apply a sector revenue multiple; the middle comes from earnings and cash-flow multiples. The spread is not noise. It is a map of which assumption — growth, payout, sector pricing — a reader has to believe to arrive at each number. Part VII turns that map into a single estimate and an uncertainty band, and the band is wide precisely when the map is.

Sector routing: the engine's central idea

The engine does not weight the nineteen methods equally. Each sector — roughly ninety classifications, counting sub-industry overrides — carries a weight vector that says which methods speak loudest there. Utilities lean on the dividend discount model; banks on book-value and residual-income methods; software on revenue and growth multiples; REITs on FFO; energy on cycle-aware measures. A method with zero weight in a sector is still computed and shown (so you can see it), but it does not move the blend. Part VII gives the representative weights; each method chapter says where that method is weighted heavily, lightly, or not at all.

Two corollaries follow. First, the same method can be excellent in one sector and destructive in another — a DCF on a regulated utility produced a model value of $17.56 against an $85 price in one of Part VIII's walkthroughs, not because the utility was overvalued fivefold but because a cash-flow formula applied to a business that by design reinvests all its cash is the wrong tool. Second, when the engine's sector classification of a company is wrong, its blend inherits the wrong toolkit; the "Model Assumptions" panel on every stock page — which lists the

weights actually used, the adjustments applied, and every abstention — exists so that a reader can notice.

What blending does and does not fix

Blending reduces the influence of any one method's idiosyncratic error. It does not reduce *shared* error. If the whole market is paying 30 times earnings for a sector that historically fetched 18, every relative method inherits the 30. If interest rates move, every discounting method moves together. The methods that protect against shared error are the intrinsic ones with conservative assumptions — EPV, Graham, residual income with a full fade — and they are in the blend precisely as ballast. A reader who understands which methods share an error source reads the blend correctly; one who treats nineteen as nineteen independent votes does not.

Chapter 4 — The inputs every method shares

Nineteen formulas share a surprisingly small set of inputs. Get these right and most methods follow; get one wrong and it propagates into several at once. This chapter defines them once so the method chapters can refer back.

Per-share fundamentals (trailing twelve months)

- **EPS** — net income to common shareholders over the last four reported quarters, divided by diluted shares. "TTM" EPS is the default input for P/E, PEG, Graham, and mid-cycle methods. Its failure mode is contamination by one-time items: a restructuring charge, an asset sale, a tax

benefit. The engine carries an explicit TTM-EPS contamination guard that down-weights earnings methods when trailing EPS looks unrepresentative.

- **Forward EPS** — analyst consensus for the next fiscal year where available, otherwise TTM EPS grown at a modelled rate. Feeds forward P/E and PEG.
- **Book value per share (BVPS)** — common equity divided by shares. Tangible book (TBVPS) strips goodwill and intangibles and matters for banks.
- **Revenue per share (SPS)**, **EBITDA per share**, **EBIT per share**, **free cash flow per share (FCFPS)**, **dividends per share (DPS)** — each the TTM figure divided by diluted shares. FCF is operating cash flow less capital expenditure.
- **Net debt per share** — total debt minus cash and equivalents, divided by shares. It is the bridge between enterprise value and equity value (below). When net debt is missing the enterprise-value methods should abstain rather than assume zero. That is the engine's design intent and it is enforced for the gross-profit multiple; at the time of writing, the EV/EBITDA and EV/FCF paths still fall back to treating missing net debt as zero, which overstates equity value for a levered company whose debt did not come through. If you build these by hand, abstain.

The discount rate: cost of equity and WACC

Every intrinsic method discounts something — dividends, cash flows, residual income — and the rate it discounts at is the single most influential input in the book. The engine uses the Capital Asset Pricing Model for the **cost of equity**:

```
r_e = RFR + β_adj × ERP + illiquidity premium
β_adj = min(β, 1.5)
```

- **RFR** is the live risk-free rate (the 10-year U.S. Treasury yield). The engine reads it live; it does not hard-code a number.
- **ERP** is the equity risk premium — the extra return equity investors require over the risk-free rate. The engine reads a live estimate.
- **β** is the stock's beta, its historical sensitivity to market moves. The engine **caps beta at 1.5**. An uncapped raw beta on a volatile growth stock can reach 2 or 3 and produce discount rates near 20% that crush any future cash flow to insignificance; the cap is standard practice (Damodaran, *Investment Valuation*) and the engine adopts it.
- An **illiquidity premium** is added for small or thinly traded names, because a holder of an illiquid stock bears a cost the CAPM does not price.

Enterprise-level methods (EPV, EVA, and the DCF terminal mechanics) discount firm-level cash flows at the **weighted average cost of capital**:

```
WACC = w_e × r_e + w_d × r_d × (1 − tax)
w_e = E / (D + E),  w_d = D / (D + E)
```

where r_d is the pre-tax cost of debt (the engine uses the risk-free rate plus a spread) and the weights come from the debt-to-equity ratio. The engine bounds WACC to a sane range and collapses it to the cost of equity when there is no debt.

Illustrative market inputs used throughout this book: risk-free rate 4.2%, equity risk premium 5.0%, corporate tax

rate 21%. These are stated once so that every example is reproducible; they are not the engine's live values, which move daily.

Growth

Growth enters three ways. **Historical growth** (revenue, EPS) describes what happened. **Analyst consensus** (next-year and year-after EPS estimates) describes what sell-side analysts assume; the engine anchors the first two years of its DCF growth on consensus when at least five analysts cover each year, clamping the implied growth to a sane band. **Terminal growth** is the perpetual rate assumed after the explicit horizon — 2.5% in the engine's DCF family, a deliberately modest number close to long-run nominal GDP growth. Any perpetuity formula **abstains when the discount rate is not comfortably above the growth rate**; the engine requires at least a one-percentage-point spread for the dividend model and a wider one for residual income, because as r − g approaches zero the formula approaches infinity and the output stops meaning anything.

Sector multiples

Relative methods multiply a per-share metric by a **sector multiple** — the P/E, P/B, P/S, EV/EBITDA and so on that comparable companies command. The engine maintains roughly ninety sector and sub-industry multiple sets, anchored to Damodaran's sector data and periodically **re-anchored to the live universe medians** and to a mean-reversion formula that pulls a sector's current multiple toward its ten-year average when the sector is dislocated. The examples in this book use illustrative multiples (stated in each chapter); the live values are recalibrated and will differ.

A sector multiple is an average, and averages under-fit the tails. The engine applies dynamic adjustments at runtime — when a stock's P/E exceeds 40, for example, the P/E and Graham methods are down-weighted by 85% and the DCF and PEG methods are boosted, because a sector-average P/E applied to a company the market prices at 60 times earnings says more about the gap than about the company.

The enterprise-value bridge

Several methods value the *firm* — debt and equity together — and then convert to a per-share equity value:

```
Equity value per share = (Enterprise value per share) −
(Net debt per share)
```

This is why the EV/EBITDA, EV/EBIT, EV/Sales, EV/FCF, EPV, and EVA chapters all end with "– net debt per share". A company with a lot of debt can have a healthy enterprise value and a small equity value; one with net cash has an equity value above its enterprise value. Forgetting the bridge is the single most common error in hand-built valuations. Treating a missing net-debt figure as zero is the second, and it is the more dangerous of the two because it fails silently and always in the same direction — upward.

The two example companies

Northwind Tools Inc. (fictional, Industrials): price $50.00 · EPS $3.20 (forward $3.39) · BVPS $18.00 · revenue/share $40.00 · FCF/share $2.80 · EBITDA/share $6.00 · EBIT/share $4.50 · net debt/share $5.00 (debt $9.00, cash $4.00) · DPS $1.20 (payout 37.5%) · beta 1.10 · ROE 17.8% · ROIC 14% · revenue growth 6% · ten-year terminal growth 2.5%. Cost of

equity = 4.2% + 1.10 × 5.0% = **9.7%**. WACC with market-value weights of 84.7% equity / 15.3% debt and a 6.2% pre-tax cost of debt = **9.0%**. Illustrative Industrials multiples: P/E 17 · P/B 2.5 · P/S 1.6 · EV/EBITDA 11 · EV/EBIT 14 · EV/Sales 1.8 · P/FCF 20 · EV/FCF 19 · fair PEG 1.0.

Helix Cloud Corp. (fictional, Software): price $120.00 · EPS $1.10 (forward $1.45) · revenue/share $30.00 · FCF/share $1.80 · revenue growth 28% · EPS growth 32% · beta 1.35 · ROIC 22%. Cost of equity = 4.2% + 1.35 × 5.0% = **11.0%** (beta below the 1.5 cap). Illustrative Software multiples: P/S 9.0 (growth profile) or 4.0 (mature profile) · P/E 32 · fair PEG 1.4.

Every worked example in Parts II–VI uses these inputs; Part VII blends Northwind's results into one estimate.

Continue reading

This sample ends with Part I. Parts II–VIII and the appendices — the remaining chapters — are in the full edition at **equity-rank.com/handbook**. Educational use only; not investment advice.