REPORTS LAB · REPORT NO. 04 · 2026

Private Equity *and* Venture Capital

A Complete Professional Guide

38 chapters · 6 parts · ~6 hour read

Contents

Part I

The Landscape and Architecture of Private Markets

What PE/VC is, who participates, why private markets exist

CHAPTER 01

Why Private Markets Exist

The illiquidity premium and what stays off the public market

$13T+	8–11%	10–12 yr
global private markets AUM (Preqin, 2025)	long-run public-equity premium typically targeted	standard PE fund life from first close

Public markets get the headlines. Private markets quietly run the longer half of capitalism — the half where companies are built, restructured, recapitalised, and held long enough for fundamentals to actually compound. To understand why PE and VC exist at all, you have to start with what the public markets cannot do.

Where the public market stops

A listed company has to print its financials every quarter, host a call for analysts who will be merciless if guidance slips, and live with a price that moves on macro headlines unrelated to its business. None of that is friendly to the kind of work that is needed to rewire pricing, replace a sales leader, integrate two tuck-ins, and re-segment a customer base. The patient capital that does that work has, for fifty years, mostly lived inside private partnerships.

There is a structural reason the public market does not absorb this work. Public investors are paid to maintain liquidity for everyone else, which means they cannot tolerate the multi-year drawdowns that good operating change frequently requires. They also cannot price an unaudited, owner-run business with a customer concentration problem — the information asymmetry is too large. Private investors solve both problems by accepting illiquidity and by buying the right to put their own people on the board.

The illiquidity premium, honestly stated

The classical defence of PE/VC returns is the illiquidity premium: investors should earn extra for accepting that they cannot exit on demand. The empirical evidence — Kaplan and Schoar (2005), Harris-Jenkinson-Kaplan (2014, updated 2024) — shows a real but modest premium for top-quartile buyout funds and a wide dispersion across all of them. The PME (Public Market Equivalent) test we examine in Chapter 6 is the honest yardstick.

The premium is not free money. It is paid for in real costs: a 10-year lockup, the GP's 2% management fee, the 20% carry, and the genuine probability of dealing with a bad vintage. The professional investor's job is to price all of those against the spread the strategy actually delivers.

What private markets buy that public markets don't

Three structural differences matter. First, governance: a PE owner usually controls the board, hires the CEO, and signs every material capex decision — public-market shareholders can only vote against management once a year. Second, capital structure: a buyout layers leverage that no public board

would accept on its balance sheet, because the equity holder has direct control over de-risking decisions. Third, time horizon: the median PE hold has been 5–7 years, which is longer than most public CEOs survive in their seat.

These advantages compound. They are also the reason PE/VC funds can credibly underwrite operating plans the public market would never finance — turnarounds, carve-outs from corporate parents, founder buyouts, growth equity rounds in companies whose IPO is still 4 years away.

Who actually pays for it

The capital that flows into PE/VC comes from a small number of sophisticated allocators: pensions (state and corporate), endowments and foundations, sovereign wealth funds, insurance general accounts, family offices, and (since the 2020s) the high-net-worth wirehouse channel via interval funds and feeder structures. We profile each in Chapter 12.

This investor base shapes everything downstream — fund size, hold period, sector preferences, and the degree to which a manager can or cannot deviate from the LPA they sold. The 'fund' is not just a pool of money; it is a contract between people who could not have written it directly themselves.

CHAPTER 02

The Private Equity Universe

Buyouts, growth, venture, credit — and why the labels matter

60%+	$200B+	4 distinct
of PE AUM in buyout strategies	deployed annually in venture (Pitchbook 2024)	return architectures — and you should know them

PE is not one strategy. It is a label that covers buyouts, growth equity, venture capital, private credit, infrastructure, and real assets — each with different deal sizes, return profiles, and professional cultures. Mixing them up is the surest sign of an outsider.

Buyouts: control + leverage + time

A buyout fund acquires a controlling stake — usually 100% — in a mature, cash-generative business, finances most of the purchase price with debt, and underwrites a 4–7 year operating plan. The targets are companies whose unit economics already work but whose growth, margin, or balance sheet can be improved under new ownership. The return architecture is the famous returns bridge: multiple expansion, EBITDA growth, and debt paydown in some combination. This is the Bain, Blackstone, KKR, and Apollo flagship strategy.

Within buyouts there is a sharp distinction between **large-cap** (deals over $1B), **middle-market** ($100M–$1B), and **lower-middle-market** (under $100M). Each has its own sourcing patterns, its own debt providers, and its own competitive dynamics.

Growth equity: minority, cash-flow positive, scaling

Growth equity sits between late-stage venture and small-cap buyout. The companies are profitable or near-profitable, growing 20–50% per year, and the investor takes a minority position with significant governance rights but not formal control. Insight Partners and General Atlantic built the modern playbook here. Returns come from revenue compounding rather than leverage.

Venture: power-law returns from earliest signals

Venture capital invests minority equity in pre-product or pre-scale companies in exchange for the right to participate in a power-law outcome. The math is structurally different — most investments fail, and the fund is carried by one or two outliers per vintage. Chapter 3 walks the venture sub-ecosystem in detail.

Venture economics are different too: shorter operating involvement, lighter governance, and a willingness to lose 100% on bets where the upside is 50–200x.

Private credit: the lender's seat at the same table

Private credit funds (direct lending, mezzanine, distressed) lend to many of the same companies PE funds buy. Returns are coupon-driven (8–12% gross unlevered for senior direct lending; higher for mezzanine and unitranche). The growth of private credit since 2010 has reshaped buyout debt markets — the largest private credit shops now underwrite buyout financings the broadly-syndicated loan market can't or won't price.

Infrastructure and real assets

Infrastructure and real assets — toll roads, midstream pipelines, fibre networks, renewable generation, real estate — are managed inside PE platforms because they share a key feature: long-term contracted cash flows that institutional LPs want for inflation hedging and liability matching. The return profile is more bond-like than equity-like, with target IRRs in the 8–14% range and longer fund lives (12–25 years for some core infrastructure).

CHAPTER 03

The Venture Capital Sub-Ecosystem

Pre-seed to crossover, and what each stage actually buys

Pre-seed → IPO	~3%	$200M
the seven-stage capital staircase	of seed-stage companies that reach a $1B+ exit	typical Series A → Series B step-up in a hot vertical

Venture capital looks like one industry from outside and like five from inside. The pre-seed angel writing $50K cheques and the crossover hedge fund leading a $400M Series E both call themselves "VC". The economics, governance, and risk profile of each could not be more different.

The capital staircase

From earliest to latest, the venture stages are: **pre-seed** (idea/team, often pre-product), **seed** (early product, first customers), **Series A** (product-market fit, early scale), **Series B** (growth-stage scale-up), **Series C+** (mature growth, often profitable or near it), **late-stage / pre-IPO**, and **crossover** (capital that can hold the position post-IPO). Each stage has typical check sizes, typical ownership targets, and typical risk profiles.

The fund-size logic explains why funds specialise. A $50M seed fund cannot lead a $200M Series C; a $5B late-stage fund cannot economically write a $500K seed cheque. Stage-specific funds form because the math forces them to.

How fund size drives strategy

A useful identity: target ownership × fund size ÷ check size = number of investments. A $200M seed fund targeting 15% ownership writes ~30 cheques of ~$3M each over a 3-year deployment window. A $1B Series-B fund writes ~15 cheques of ~$50M each. The fund's strategy is dictated by this arithmetic, not by the partners' preferences.

This is also why every venture firm eventually faces the same strategic question: do we raise a bigger fund and move later, or stay disciplined and stay early? Almost every firm has answered "bigger" since 2014, with consequences we examine in Chapters 35 and 41 of this report.

Crossover capital and the public-private blur

Tiger Global, Coatue, and the Series-D-to-IPO crossover funds reshaped venture from 2018 onward by pricing private growth-stage rounds at near-public multiples. The 2021 peak collapsed in 2022; 2024–

2025 saw a partial return with much sharper underwriting. The crossover layer is the most macro-sensitive part of the venture stack.

Sector thesis as a structural choice

Modern venture is increasingly thesis-driven — funds market themselves around sectors (AI infrastructure, climate tech, biotech, defence-tech, fintech). A clear thesis attracts founders in that vertical and gives the fund informational edge. The trade-off is concentration: a sector thesis can underperform a generalist by a vintage if the sector is mistimed.

CHAPTER 04

Key Participants

GPs, LPs, advisors, and the supporting cast

8–12 parties	0.5–2%	3 layers
typical roster on a $1B buyout	advisor fee load per side	of fiduciary obligation to map

Every PE/VC transaction has the same supporting cast — the same lawyers, the same accountants, the same placement agents. Knowing who they are and what they want is half of running a deal.

On the GP side

The **general partner** entity is the legal manager of the fund. The **management company** employs the deal team and receives the management fee. The **investment committee** approves transactions; on most funds it includes two to four senior partners and is the binding decision body. The deal team itself is layered — partner, principal, vice president, associate, analyst — with most diligence work done by the bottom three layers and the partner owning the relationship and the IC pitch.

On the LP side

An **LP** can be a public pension, a corporate pension, an endowment, a foundation, a sovereign wealth fund, an insurance general account, a family office, a fund-of-funds, or (since the 2020s) a high-net-worth wirehouse channel. Each has different time horizons, regulatory constraints, return targets, and fee sensitivities. We profile each in detail in Chapter 12.

Intermediaries and advisors

Placement agents raise capital on behalf of GPs in exchange for a fee (typically 1–2% of commitments raised). **Investment banks** run sell-side processes, advise on financing, and underwrite IPOs. **Diligence advisors** — Big Four QofE teams, sector consultants, technical experts — are paid by the buyer to stress-test the seller's claims. **Legal counsel** drafts every document; the GP-side firm and target-side firm typically do not change across deals.

Fund administrators calculate NAV, process capital calls, and produce LP reporting. **Auditors** sign the annual financials. **Custodians** hold the fund's securities. **Technology providers** (Carta, eFront, Allvue, iLEVEL) run the data infrastructure.

Inside the portfolio company

The CEO and CFO of the portfolio company are technically employees of the company, not the fund — but in practice they answer to the GP via the board. **Operating partners** at the GP firm are senior ex-CEOs or functional experts who advise across the portfolio. **Independent directors** sit on the board to provide governance and (in the case of a sale) Revlon-duty cover.

CHAPTER 05

Historical Development and Market Cycles

From the 1980s LBO wave to 2024 rate normalization

1981	$30B	475 bps
when KKR closed its first institutional fund	RJR Nabisco — the deal that made the industry famous	Fed funds rate move, 2022–2023

PE's history reads as a series of regime shifts triggered by capital markets, regulation, and macro. The strategy that won in 1985 was illegal in 1995, the strategy that worked in 2005 ate itself in 2008, and the strategy that compounded from 2010–2021 ran into a wall when rates moved.

The 1980s: the LBO becomes an industry

Before 1980, leveraged acquisitions were a niche craft. The institutional LBO fund — pioneered by KKR, Forstmann Little, and a handful of others — turned that craft into an industry by raising third-party capital in 10-year limited partnership vehicles. The 1989 RJR Nabisco buyout, immortalised in Barbarians at the Gate, marked the cultural peak; the high-yield bond market crash that followed marked the first regime change.

The 1990s: institutionalisation

Through the 1990s the industry built the infrastructure: standardised LPAs, professional placement agents, secondary markets, and benchmarking. Returns were strong; capital was relatively scarce. Venture capital, separately, rode the dot-com wave to a 1999 vintage that produced both legendary winners (Google, Amazon comp set) and the worst losses of the asset class's history.

The 2000s: leverage abundance

The 2003–2007 buyout boom was driven by cheap, abundant credit — the era of covenant-lite term loans, second-lien debt, and 7x EBITDA leverage on cyclicals that should have carried 4x. The 2007–2008 vintages took the worst of the GFC; the 2009–2010 vintages, deploying into a discounted market, delivered some of the best returns in industry history.

The 2010s: yield compression and abundance

The post-GFC era of zero rates pushed institutional capital aggressively into private markets. AUM tripled from 2010 to 2021. Multiples expanded; entry valuations on quality buyouts ran from 8–10x EBITDA to 12–14x. Venture capital had its own boom, peaking in late 2021 with crossover capital pricing growth-stage rounds at unsustainable multiples.

2022–2025: the rate regime returns

The 2022–2023 rate move broke the 2010s playbook. Buyout debt repriced from L+275 to S+550+, exit multiples compressed, and the IPO window closed. Many 2019–2021 vintages now carry assets at marks the market will not validate. Distributions to LPs collapsed; secondary discounts widened to 15–25% on quality books and 30%+ on stressed ones. Continuation vehicles (Chapter 33) emerged as the dominant 2024–2025 liquidity tool.

CHAPTER 06

Performance, Benchmarks, and the Case for Alternatives

IRR, TVPI, DPI, PME — and what the data actually says

IRR	PME	~25%
the dominant metric — and the most easily gamed	the honest test of alpha	of buyout funds that beat the S&P 500 (Kaplan-Sensoy-Strömberg)

PE/VC performance is one of the most-debated empirical topics in finance. The headline numbers are flattering; the honest numbers, after fees and benchmarking against a public alternative, tell a more nuanced story.

The five metrics every LP knows

IRR annualises returns and is the industry's default metric — but it is sensitive to deal pacing and can be inflated by subscription-line financing that delays the first capital call. **TVPI** (total value to paid-in) is the cleaner total-return multiple. **DPI** (distributions to paid-in) measures realised cash returns and is the LP's true scorecard once a fund matures. **RVPI** (residual value) is the unrealised mark, contestable and increasingly so since 2022. **Net** versions of each subtract fees and carry — that is the number LPs actually earn.

A useful sanity check: a fund with a 25% gross IRR and a 17% net IRR is paying an 8-point spread to the GP. A fund with 25% gross and 12% net is paying a 13-point spread, which is a lot.

PME: the public-market yardstick

The Kaplan-Schoar Public Market Equivalent (PME) test asks the most important question in the asset class: did the LP do better in this fund than in a public-equity index over the same period? A PME of 1.0 means the fund tied; above 1.0 means it added value; below 1.0 means it didn't. The empirical result, across many studies, is that the median buyout fund has a PME of ~1.0–1.1 — barely positive — while top-quartile funds have a PME of 1.3+ and bottom-quartile funds destroy value relative to the index. **Manager selection matters more than asset-class allocation.**

Dispersion and persistence

PE/VC returns show the widest dispersion of any institutional asset class — top decile to bottom decile in venture often spans 30+ IRR points. There is also persistence: top-quartile managers in one vintage are more likely than chance to be top-quartile in the next, a result first documented by Kaplan-Schoar and confirmed in subsequent updates. This is why LP manager-selection effort is concentrated at the top of the distribution.

Biases in the published data

Reported PE/VC benchmarks suffer from several construction biases worth naming. **Survivorship bias**: failed funds may not report. **Selection bias**: voluntary reporting tilts toward better outcomes. **Mark-to-model bias**: unrealised values are GP-determined and tend to be smoothed. **Vintage timing bias**: fund-year cohorts have different underlying market exposures. Honest LP analysis adjusts for all four.

Part II

Fund Economics, Legal Structures, and Fundraising

LPA terms, economics, GP/LP dynamics, capital raising

CHAPTER 07

The Limited Partnership Structure

GP entity, management company, parallel funds, and offshore feeders

4–6 entities	LP, GP, MgmtCo	10+1+1
in a typical fund structure	the three core actors	standard fund life — 10 years plus two 1-year extensions

The PE fund is a legal artifact, not a thing in nature. It is a limited partnership designed to do one job: pool committed capital, deploy it through a small group of decision-makers, and return proceeds to a larger group of investors — under terms negotiated up front.

The legal anatomy

The **fund vehicle** is a limited partnership, usually formed in Delaware (US funds) or the Cayman Islands (international funds). Its only purpose is to hold investments. The **general partner entity** is a separate Delaware LLC owned by the senior partners; it is the legal manager of the fund and the recipient of carry. The **management company** is a third entity that employs the investment professionals and receives the management fee. **Parallel funds** are clones that invest pari-passu alongside the main fund; they exist to accommodate ERISA-regulated LPs, offshore investors, or LPs with tax-driven needs. **Feeder funds** aggregate non-US investors into a single LP slot in the master fund.

This separation is deliberate. The fund vehicle isolates investments. The GP entity isolates carry. The management company isolates payroll and lease obligations. Parallel and feeder vehicles isolate regulatory and tax exposures. None of this is overhead — it is the architecture that lets sophisticated LPs participate at all.

Why a limited partnership?

Two features of the LP form make it the dominant vehicle. First, **pass-through tax treatment**: LPs report their share of fund income as if they had invested directly, avoiding entity-level tax. Second, **limited liability**: LPs cannot lose more than their committed capital and have no liability for fund debts. The GP, by contrast, has unlimited liability for the GP entity — which is why GPs almost always sit behind a Delaware LLC.

Fund life and term

A standard fund has a 10-year life with two 1-year GP extensions, usually unilateral (the first) and LPAC-approved (the second). The first 4–5 years are the **investment period**, during which new investments can be made. The remaining 5–6 years are the **harvest period**, during which the GP exits investments and returns capital. After year 12, the fund is theoretically wound down — though in practice continuation vehicles (Chapter 33) or LP secondaries handle the tail.

CHAPTER 08

The Limited Partnership Agreement

Investment period, key-man, removal rights, and the LPAC

100–200	ILPA	75%
pages in a typical LPA	the standard against which terms are benchmarked	common LP supermajority threshold for no-fault GP removal

The Limited Partnership Agreement — the LPA — is the constitution of the fund. Every economic outcome, every governance right, every exception to a rule someone forgot to anticipate, is written down here. Reading an LPA fluently is the most valuable single skill on the LP side of the table.

How an LPA is organised

A standard LPA has a recognisable structure: definitions; commitments and capital calls; investment policies and limitations; management and operations; fees and expenses; the waterfall; transfers; defaults; reports; the LPAC; key-man and removal; term and dissolution. Most LP analysts read in this order, and the negotiated provisions (key-man, removal, fee offsets, LPAC composition) get the heaviest attention.

What a sophisticated reader looks for is not just what each section says but what each section doesn't say. An LPA without a clear no-fault removal provision is a different fund from one with a 75% supermajority threshold — even if neither is ever invoked.

Investment period and key-man

The **investment period** (typically 5 years) is the window during which the GP can call capital for new investments. After it ends, the GP can only call for follow-ons, fees, and expenses on existing investments. **Key-man clauses** automatically suspend the investment period if a named partner dies, departs, or stops devoting time to the fund. Suspension can be lifted by LPAC vote; absent a lift, the fund essentially stops investing. The list of named key persons is one of the most-negotiated parts of any LPA.

Removal rights and no-fault divorce

LPs can remove the GP **for cause** (typically defined as fraud, gross negligence, or material LPA breach) by simple-majority vote. **No-fault removal** — termination without cause — typically requires a 75–80% supermajority and triggers carry-forfeiture or carry-haircut consequences for the removed GP. These rights are rarely invoked but profoundly shape GP behaviour.

The LPAC

The Limited Partner Advisory Committee is a small body (typically 5–9 members, drawn from the largest LPs) that votes on conflicts, valuations, and any matters the LPA explicitly delegates to it. The LPAC is not management — it cannot direct investments — but it is the LPs' voice on conflicts, fee disputes, and continuation-vehicle proposals (Chapter 33). LPAC composition and the precise list of LPAC consents are heavily negotiated.

CHAPTER 09

Fund Economics

Management fees, carried interest, hurdles, and the waterfall

2 and 20	8%	1.5–2%
the canonical fee structure	standard preferred return / hurdle rate	modern range on management fees

Carried interest is the deepest, weirdest, most-litigated piece of compensation in finance. Done right, it aligns LP and GP interests beautifully. Done wrong, it pays the GP for things the LP didn't agree to. Every clause in a fund waterfall exists because someone, in some prior fund, tried to game the previous version.

The management fee

Annual fee of 1.5–2.0% during the investment period, calculated on **committed capital**. After the investment period, the fee typically steps down to 1.0–1.5% on **invested capital** (or, in some funds, on net invested capital — committed less write-offs and realisations). The step-down matters: a $1B fund with a 5-year, 2.0% / 5-year 1.0% structure pays $150M in fees over its life, not $200M.

Modern fee offsets — required by the 2023 SEC Private Fund Adviser Rules and standard in ILPA-aligned LPAs — credit transaction fees, monitoring fees, and break-up fees collected from portfolio companies back to the management fee.

Carried interest

The GP receives 20% of fund profits above a hurdle. The 20% is near-universal; what varies is how it is calculated. Two structures dominate: **European (whole-fund) waterfall**, where LPs must receive their full committed capital plus the hurdle before any carry is paid; and **American (deal-by-deal) waterfall**, where the GP earns carry as each deal exits, subject to a clawback at fund end if cumulative returns fall short. European is LP-friendly and standard in venture and Europe; American is GP-friendly and standard in US buyout.

Hurdle and catch-up

The **hurdle** (or preferred return) is typically 8% IRR. Below 8%, all distributions go to LPs. Above 8%, a **catch-up** sweeps further distributions (often 100%) to the GP until the GP has caught up to a 20% share of total profits. After catch-up, distributions are 80% to LPs and 20% to the GP — the steady-state carry split.

Worked example: $100M fund returns $200M (2x). LPs first receive $100M back. Hurdle of 8% over 5-year hold ≈ $47M, taking LPs to $147M. Catch-up: GP receives $11.7M (20% of $58.7M). Steady state: remaining $41.3M splits $33.0M to LPs and $8.3M to GP. Total: LP $180M (1.8x), GP $20M (20% of $100M profit). The math always reconciles, given a clean structure.

Clawback

If a fund pays carry on early winners and later losers reduce the fund's overall return below the deal LPs were promised, the GP must **claw back** previously distributed carry. Clawback is typically secured by an interim escrow on each carry distribution and by the personal guarantees of the senior partners. Clawback machinery is more important in American waterfalls (where the GP is paid as deals exit) than in European waterfalls (where the GP is only paid after LPs are made whole).

GP commit and recycling

The **GP commit** is the capital the GP itself contributes to the fund. ILPA-aligned funds expect 1–5%. **Recycling** allows the GP to re-invest realised proceeds back into new investments, deploying more than the committed capital base over the fund's life — a form of leverage on the LP commitment that benefits both sides if used disciplinedly.

CHAPTER 10

GP/LP Dynamics

Alignment, governance, and the ILPA Principles

ILPA Principles 3.0	3 alignment levers	$100M+
the de-facto LP-side standard	fees, carry, and removal	scale of GP commit at top funds — alignment, made tangible

The principal-agent problem at the heart of PE/VC is older than the asset class. The LP gives money to the GP to invest on their behalf, knowing the GP has more information, more expertise, and entirely different risk preferences. Every term in an LPA is, in some way, a response to this problem.

The conflict, named

GPs and LPs share an interest in maximising the fund's net return. They differ in almost everything else. GPs prefer larger funds (more fees), longer fund lives (more management fee periods), looser restrictions (more strategic flexibility), and softer reporting (less embarrassment). LPs prefer the opposite of each. The LPA is where this is negotiated.

The ILPA Principles

The Institutional Limited Partners Association published the first ILPA Principles in 2009 and updated them through Version 3.0 (2019). The Principles set norms for fee transparency, GP-LP alignment, and governance. They are not law, but they are the dominant LP-side benchmark — a fund whose terms diverge meaningfully from ILPA must explain why.

ILPA also publishes standard **reporting and fee templates** (the ILPA Reporting Template, the Quarterly Capital Call and Distribution Notice templates). Adoption is widespread among institutional LPs.

The three alignment levers

Carry structure: a European waterfall with a 1x-and-hurdle preferred return aligns GP economics with LP economics in a way an American deal-by-deal waterfall does not. **GP commit**: a 5%+ GP commit means a partner's personal balance sheet rides the same returns as the LPs'. **Removal rights**: a 75% no-fault removal threshold is rarely invoked but creates an out-of-the-money option that disciplines GP behaviour.

Trust as the binding glue

Even with the best LPA, fund governance ultimately depends on multi-decade relationships between named individuals. LPs allocate to GPs they have re-upped with for three or four vintages; GPs raise larger funds because the same anchor LPs come back. The negotiated terms set the floor; the relationship sets the ceiling.

CHAPTER 11

Fundraising

PPM, roadshow, DDQ, and how a close actually happens

12–24 months	100+	3–6
typical fundraise duration	LP meetings before a $1B fund closes	closes between first and final

Raising a fund is the hardest sales process in finance. The product takes 10 years to deliver. The buyer has to underwrite the seller's people for that decade. And every other GP in the market is calling on the same pool of LPs simultaneously.

Preparation: the year before the road

Before the first LP meeting, the GP team prepares a Private Placement Memorandum (PPM), a track-record analysis, an organisational deck, a market research pack, and a draft LPA. They run a **pre-marketing** phase with anchor LPs to test interest and refine positioning. They appoint legal counsel (Kirkland, Debevoise, Simpson Thacher are the dominant US shops) and often engage a **placement agent** for distribution reach.

Track-record presentation is the single highest-stakes piece of the materials. It must reconcile to audited numbers, attribute returns to attribution-eligible deals, and survive a Big Four diligence call. Cosmetic adjustments to the track record are the fastest way to lose institutional credibility.

The Due Diligence Questionnaire

Most LPs send a standardised DDQ — often the ILPA template — covering organisation, investment process, deal track record, ESG, compliance, technology, and conflicts. Answering a 200-question DDQ is a 4–6 week exercise and shapes the LP's first analytical view of the fund.

The roadshow

Over 6–12 months the GP team meets every plausible LP in a sequence of 1-hour meetings, deeper due diligence sessions, on-site visits, and reference calls. Sophisticated LPs will request access to portfolio company management, prior-fund LPAC members, and the firm's internal data. Final decisions go through the LP's investment committee.

First close to final close

Funds typically have a **first close** 6–9 months into the raise (anchor LPs fund), interim closes, and a **final close** 12–24 months in. LPs joining at later closes pay an **equalisation interest** to compensate first-close LPs for the time-value of capital already deployed. The final close locks the fund size; the GP must then deploy.

CHAPTER 12

LP Types and Allocation Frameworks

Endowments, pensions, sovereigns, family offices — and pacing models

8 archetypes	5–15%	80% / 20%
of institutional LPs	typical PE allocation in a US public pension	the split between top-decile and everyone else, in any large LP's portfolio

From the GP's seat, the LP base looks like a list of names and commitment amounts. From inside, it is a small number of fundamentally different institutions with different mandates, time horizons, and tolerances. A GP who treats them all the same will lose all of them.

The eight LP archetypes

Public pensions (CalPERS, NYS Common, Texas Teachers): largest slice of US institutional capital, bound by fiduciary duty and FOIA disclosure rules. **Corporate pensions**: similar mandate, less disclosure. **Endowments** (Yale, MIT, Princeton): high allocation to alternatives, long horizon, top-quartile manager selection. **Foundations**: smaller scale of endowment-style. **Sovereign wealth funds** (GIC, ADIA, Mubadala, Temasek): largest cheque sizes, often co-invest. **Insurance general accounts**: liability-driven, prefer credit and infrastructure. **Family offices**: idiosyncratic mandates, fastest decisions. **Funds-of-funds and consultants**: aggregators on behalf of smaller institutional or HNW capital.

Pacing models

An LP committing $500M per year to PE over a 10-year vintage cycle does not invest $5B in year one. Committed capital is called over 4–5 years; distributions start in year 4 and peak in years 6–8. A **pacing model** projects the timing of calls, distributions, and NAV to keep the LP at its target allocation through the cycle. After 2022, when distributions slowed, many LPs found themselves over-allocated via the denominator effect (Chapter 35) and had to slow new commitments.

Co-investment and secondaries

Sophisticated LPs supplement primary commitments with **co-investments** (direct stakes alongside the GP, fee-and-carry-free) and **secondaries** (purchasing other LPs' fund interests). A modern $50B LP might have 60% primary, 25% co-invest, and 15% secondaries. The blended fee load is dramatically lower than 100% primary — a structural alpha that drives much of the post-2010 institutional capital reorganisation.

Part III

Sourcing, Diligence, and Valuation

Deal origination through investment decision

CHAPTER 13

Deal Sourcing

Proprietary flow versus the auction

80/20	3–5x	12–18 months
typical proprietary-to-auction split varies wildly by firm	valuation premium for hot-sector auctions vs. proprietary deals	courtship period for top proprietary deals

Every PE/VC return ultimately comes from one of two places: a deal won at a price the seller didn't realise was low, or a deal won at a normal price where the buyer's operating thesis was right. Sourcing — how the deal got onto your desk — determines which kind of deal you have before any model is built.

The intermediated auction

An investment bank runs a process. CIM goes to 50+ buyers under NDA. First-round bids in 4 weeks; selected buyers proceed to management presentations and data room access; second-round bids in 4–6 more weeks; signing on the highest credible bid. The seller's goal is to maximise price under disciplined process; the buyers' goal is to win without overpaying. **Auctions tend to clear at a price that reflects what the marginal buyer will pay, not what the deal is structurally worth.**

Auctions have a known cost structure to PE buyers: 2–4% of EV in advisor fees per side, plus the indirect cost of multiple expansion. Some firms deliberately avoid auctions; others compete in them but only at sectors and sizes where they have an informational edge.

Proprietary sourcing

Proprietary deals arrive without an advisor — direct from a CEO who knows the GP from prior diligence, a board member at a portfolio company, a lender flagging a balance-sheet stress, or a sector-specific search. Proprietary flow is the holy grail because it tends to clear at lower entry multiples and at terms closer to what the GP wants. The cost is patience and relationship investment.

Modern PE firms increasingly maintain dedicated **sourcing teams** — partners or principals whose only job is to map a sector, build relationships with every CEO and operator, and sit at the table when an opportunity arises. The premier case studies (Vista Equity in software, Thoma Bravo in cybersecurity) have built billion-dollar sourcing edges this way.

Hybrid sourcing models

Limited-process deals invite 5–8 strategic buyers and skip the broader auction. **Negotiated** deals start as proprietary and add a single competing buyer late, to validate price. **Sponsor-to-sponsor** deals (Chapter 30) are PE-to-PE transactions, usually advisor-led but with deeply prepared buy-side teams who know the asset from the prior fund's reporting.

Platform sourcing — WeFunder, AngelList, BizBuySell, and the SPV-syndicated round

Equity crowdfunding platforms — WeFunder, Republic, StartEngine — operate under Reg CF and Reg A+ exemptions that allow non-accredited investors to participate in private financings. These platforms rarely originate institutional-quality deal flow, but they are a live signal layer. A company raising $500K on WeFunder has disclosed a valuation, a business summary, and a funding target that any analyst can study without signing an NDA. Watching what surfaces on these platforms tells you about market demand at the earliest stage, and sometimes flags a company that will grow into a fundable deal years later.

AngelList syndicates have largely replaced the informal angel networks of the 2000s. A syndicate lead — typically a domain-expert operator or early-stage investor — invests directly and invites co-investors to participate in the same economics via an SPV. The economics are transparent: the lead carries the SPV at 15–20% carry, and each backer gets pro-rata exposure with no management fee. The important sourcing implication is that AngelList deal pages aggregate information — round size, round composition, and the identity of the lead — that is useful for building a map of who is active in a sector before the company ever runs an institutional process.

BizBuySell is the dominant US marketplace for small-business listings. For search-fund principals and lower-middle-market PE associates, BizBuySell is a practical sourcing starting point: thousands of businesses listed with asking prices, revenue, and cash-flow figures, many of them owner-operated businesses where the seller has no investment-bank advisor and no competing buyers. The entry multiples on BizBuySell listings are far below what a banker-run process clears — the trade-off is deal quality, business model, and management dependency. A sharp associate uses BizBuySell and AngelList together as a watchlist, not a deal pipeline: patterns in what gets listed — seller age, sector clusters, pricing behaviour — are proprietary data on what is for sale in a given market. The live deal evaluation tool in this report's lab can be applied to any listing to run a fast first-pass screen.

CHAPTER 14

Screening and Initial Assessment

How to read a CIM and decide quickly

90%	3 hours	5–8
of CIMs that should be killed in the first 30 minutes	time-budget for a real first-pass screen	core questions a screening framework needs to answer

An associate or principal will read 200+ teasers and CIMs in a year. Most of those deals are bad fits — wrong sector, wrong size, wrong dynamics. The discipline is not deciding which deal to do; it is deciding which deals to stop doing, fast.

Teaser → CIM

An anonymised **teaser** arrives first — 1–2 pages from the sell-side bank with sector, size, and broad financials. A buyer who wants more signs an NDA and receives the full **Confidential Information Memorandum**: 60–120 pages on the company, market, financials, and management. The CIM is a sell-side document. Read it as a marketing brochure, not a balanced presentation.

A disciplined screening framework

Five questions worth answering in the first read: (1) **Is this in our strategy?** Sector, size, geography, instrument. (2) **What does the business actually do?** Customers, competitors, unit economics. (3) **Is the trajectory real?** Growth quality, customer concentration, retention. (4) **What's the entry price likely to be?** Reverse-engineer from the CIM language and recent comp transactions. (5) **What is the value-creation thesis?** If you can't write it in 3 sentences after a 90-minute read, the deal probably isn't for you.

A useful internal habit: associates write a one-page screening memo on every CIM that gets past the teaser. The memo recommends pass / proceed / further diligence with the specific reasons. Memos are reviewed weekly. The discipline forces explicit reasoning where pattern-matching would otherwise rule.

The anchor problem

Sunk-cost reasoning is the most expensive bias in deal screening. A team that has spent two weeks on an opportunity has a strong gravitational pull to keep going, even when the early signal said no. The cleanest defence is a **kill threshold**: explicit criteria that trigger a deal-killer review, regardless of how much work has been done. Top firms enforce this culturally.

PE selectivity — saying no to 80%+ of deals

The typical PE or VC funnel looks roughly like this: 1,000 deals reviewed at the teaser or initial-contact stage; around 50 reach a first meeting; roughly 10 receive an indication of interest; around 3 proceed to a letter of intent; one closes. The numbers vary by firm size and strategy, but the shape is universal. The

implication is that most deals die at the screening stage, not at the investment committee. IC time is expensive; the associate and principal layer exists to filter down to the 10 deals that deserve it.

The discipline that separates good screening from pattern-matching is writing explicit **kill criteria before the meeting**. A kill criterion is a specific, measurable threshold: TAM floor (e.g. $2B+ serviceable addressable market), gross margin floor (e.g. 60%+ for software), growth floor (e.g. 20%+ revenue CAGR for the last two years), and founder-pattern criteria (e.g. domain experience in the target vertical). If a deal fails any one of these, the pass is written before the management team walks in the room. This discipline protects the team from the charm and narrative that management teams are, by professional necessity, very good at projecting.

The 80%+ pass rate is a feature of a well-run PE or VC practice, not a failure. Portfolio construction requires it. A fund that does 15 deals over a 3-year deployment period needs to be confident that those 15 were the best risk-adjusted opportunities from the 1,000+ reviewed. The most common pass language is "interesting business, not for us right now" — which should mean something specific: outside strategy, outside size band, outside margin profile, or outside sector thesis. Writing that sentence clearly, with the actual reason, keeps the firm honest and leaves the door open for a future approach when the company has matured. Use the pitch deck extractor in this report's lab to structure the screening pass more systematically.

CHAPTER 15

Commercial and Operational Due Diligence

Market sizing, cohorts, unit economics, and management

4–8 weeks	$200K–$2M	3 buckets
typical commercial diligence duration	fee load for a commercial DD study	to nail: market, position, management

Commercial diligence is the deal team's external study of a target. It answers: is this market what management says it is, is this company's position what management says it is, and is the management team capable of executing the plan we are about to underwrite?

Market sizing — TAM / SAM / SOM

The discipline begins with **TAM** (Total Addressable Market — the total demand for the product or service in the universe), **SAM** (Serviceable Addressable Market — the portion the company can credibly reach given geography, channel, and product), and **SOM** (Serviceable Obtainable Market — what is realistic to win against the actual competitive set in a 5-year window). Sell-side decks routinely conflate the three. The buyer's job is to separate them, sourcing each from independent data.

Competitive positioning and customer cohort analysis

Map the competitive set — who else does this, what is their share, what are their cost positions. Then test the company's win/loss data: customer interviews, churn cohorts, NPS, pricing power. **Cohort analysis** — tracking customers acquired in successive periods through retention, expansion, and lifetime value — is the most powerful diligence technique invented in the last twenty years. A SaaS company with 130% net revenue retention is structurally different from one at 95%, regardless of similar headline growth.

Management evaluation

Three diligence channels: structured interviews with the CEO/CFO/COO, on-site shadowing, and 360-degree references with prior employers, customers, and board members. Sophisticated firms employ **industrial-organisational psychologists** for structured assessments. The evaluation question is not are these people impressive; it is are they the right people to execute the specific plan we are underwriting.

How the workstreams divide

Commercial DD (sector consultants, e.g. LEK, Bain, Parthenon): market and competitive. Operational DD (Big Four, KPMG Strategy): operations, ops capability, technology stack. Legal DD (deal counsel): contracts, IP, litigation. ESG DD (specialised firms): material ESG risks. Each workstream has its own lead, its own deliverable, and its own integration touchpoint with the deal team.

CHAPTER 16

Financial Diligence and Quality of Earnings

Normalizing EBITDA and stress-testing the projections

$150K–$1M	3–6 weeks	8–15%
typical QofE cost	study duration	common EBITDA adjustment from reported to QofE-confirmed

The Quality of Earnings study is the financial X-ray of a target. It normalises EBITDA, stress-tests projections, and identifies the financial red flags that kill deals — or that merit a price re-trade if found late. No buy-side professional underwrites without one.

Normalising EBITDA

Reported EBITDA almost always overstates run-rate cash generation. The QofE adjusts for: **one-time items** (legal settlements, restructuring, founder bonuses, one-off contracts), **accounting policy choices** (revenue recognition timing, capitalised vs. expensed), **related-party transactions** (above- or below-market lease, founder salary), and **run-rate adjustments** (annualising recent contract wins, hires, price increases). The output is a Pro Forma EBITDA the buyer will quote in their LOI.

Worked example: a $30M reported EBITDA might normalise to $26M after $4M of legitimate adjustments. At a 10x multiple, the buyer just saved $40M on the entry price — assuming the seller accepts the analysis. Re-trade fights happen here.

Working capital and the peg

Working capital diligence sets the **working-capital peg** — the level of net working capital the seller is required to deliver at close. Below the peg, the buyer is owed dollar-for-dollar. The QofE study calculates a 12-month rolling average WC and proposes a peg. Negotiation here is dollar-for-dollar with the purchase price.

Projections — building a base case the seller's plan can't

Sell-side projections are always above what the company will actually do. The QofE process re-builds the projection: customer-level revenue build, retention and churn cohorts, hiring plan, capex schedule. The buyer then runs a base / downside / upside case from that build. The base case — not the seller's plan — is what the IC underwrites against.

The red flags that kill deals

Customer concentration above 25%. Top-customer churn risk. Aggressive revenue recognition. Cash-conversion below 70% of EBITDA. Working capital trends inconsistent with growth. Off-balance-sheet liabilities, especially uncapitalised leases and pending litigation. Rapid recent margin expansion that

can't be tied to specific operational change. Each of these is a re-trade conversation; some are deal-killers.

CHAPTER 17

Valuation I — Comparables and Precedent Transactions

Mean, median, control premium, and the right peer set

EV/EBITDA	EV/Revenue	8–15
the dominant PE multiple	the dominant venture multiple	typical peer-set size — fewer is suspect, more is noise

Comparable-company analysis is the workhorse valuation technique in PE/VC. Done with discipline it triangulates a defensible price; done sloppily it produces whatever number the deal team wants. The discipline is in the peer-set construction and the multiples chosen.

Constructing the peer set

A defensible peer set is small and tight. Each company should match the target on business model, scale band, growth profile, margin profile, and end-market. Eight close peers beat fifteen loose ones. The deal team prepares a **peer-comparison table** showing each peer's revenue, growth rate, EBITDA margin, EV, EV/Revenue, and EV/EBITDA — and the median and mean across the set.

Two failure modes are common: (1) including peers that are structurally different (a mid-cap with a strategic moat is not comparable to the target), and (2) chasing the multiple by adding favourable peers late in the process. Both are visible to a sharp IC.

Multiples that travel and multiples that don't

EV/EBITDA is the default for PE because it is capital-structure-neutral and approximates cash-flow potential. Use last-twelve-months EBITDA for trailing multiples, year-1-forward for forward. Forward multiples are usually 1–2x lower than trailing because of expected growth.

EV/Revenue is the default for venture and high-growth software where margins are not yet steady. The risk: revenue multiples expand and contract dramatically with the rate cycle (the 2021–2023 SaaS multiple compression is the canonical example). Use forward revenue, not trailing.

Specialty multiples — EV/Subscribers (telecom), EV/Beds (healthcare), EV/MW (renewables), EV/GMV (e-commerce) — exist because the underlying economics aren't well-captured by EBITDA or revenue. Use them when standard multiples produce nonsense.

Precedent transactions

Precedent transactions show what buyers actually paid for similar companies recently. They embed control premiums (typically 20–30% above public-trading multiples) and synergies (where the buyer is strategic). For a PE buyer, the relevant comp is precedent PE transactions, not strategic acquisitions, because synergies are not available to a financial buyer.

CHAPTER 18

Valuation II — DCF, LBO, and Venture Methods

Building the LBO from scratch and pricing the venture round

3 statements	4 cases	returns bridge
income, balance, cash — fully linked	downside, base, upside, management	the way every IC reads an LBO output

DCF is the framework everyone learns in business school and many practitioners distrust in private markets — too sensitive to assumptions, too vulnerable to the terminal-value flaw. The LBO model is the framework PE actually underwrites with. The VC method is a third thing entirely.

DCF — useful, with caveats

Build a 5–10-year cash-flow projection (unlevered free cash flow), select a terminal value (perpetuity growth or exit multiple — exit multiple is more honest in private markets), discount at WACC. The output is enterprise value; subtract net debt for equity value.

The DCF's value in PE is not as a price-finder — multiples do that — but as a **sanity check**. If the DCF says $300M and the market clears at $700M, the multiple is implying assumptions about exit price and growth that the deal team should make explicit. DCFs are also the primary tool in infrastructure and credit, where contracted cash flows make the technique more reliable.

The LBO model — the PE workhorse

Build assumptions: entry EV (e.g. 10x EBITDA), purchase price allocation (equity / debt / management rollover), debt schedule (term loan, mezzanine, revolver), 5-year operating projection (revenue growth, margin trajectory, working capital, capex), debt paydown waterfall (cash sweeps after mandatory amortisation), exit assumption (e.g. 9x EBITDA in year 5). Output: equity-IRR, MOIC (multiple of invested capital), and the **returns bridge** — the decomposition of equity value creation into multiple expansion, EBITDA growth, and debt paydown.

A representative bridge: $100M equity in, $250M equity out (2.5x). Of the $150M equity value created, $50M from EBITDA growth, $30M from multiple expansion, $70M from debt paydown. Reading the bridge tells you immediately what kind of deal it is: leverage-driven, growth-driven, or multiple-arbitrage-driven.

VC method and the alternative for early-stage

Pre-revenue companies have no meaningful EBITDA. The **VC method** works backwards from a hypothesised exit value: pick a 5–7 year exit value (revenue × multiple); apply target-IRR discounting

back to today; that is the post-money valuation. Subtract the new investment to get pre-money. Crude, but it forces the explicit conversation about what has to be true at exit.

Scenarios and stress

Every IC presentation has at least three cases: **downside** (what if growth misses by 30%, what if margins compress by 200bps), **base** (the underwriting case), **upside** (what does management say). The base case is what the IC underwrites; the downside is the IRR floor; the upside is the bull case for steel-manning.

CHAPTER 19

The Data Room

What it should contain, what it tells you about management

12–15	100s–1000s	3 colour-codes
standard data-room categories	documents in a mature room	the buyer's diligence team uses to track issues

The data room is the seller's curated library. What is in it tells you about the company; what is missing tells you about management. Both signals matter.

What a well-organised data room contains

A standard buy-side data room is organised in folders: Corporate (constitutional documents, cap table, board minutes), Financials (audited statements, monthly management accounts, budgets, projections), Tax (returns, audits, NOLs), Customer (top-customer contracts, retention data), Operations (org charts, KPIs), HR (key-employee contracts, equity plans), IT (architecture, security audits), Legal (litigation, contracts, IP), Real Estate, Insurance, Environmental and ESG, Regulatory. Index quality is itself a signal.

What's missing tells you what management hasn't done

A mid-market deal where there is no monthly board pack, no cohort retention analysis, no detailed customer-level revenue file — that is a different deal from one where all of those exist. Some absences are signals about deal complexity (tighter post-close finance integration); some are signals about management quality. Both are price-relevant.

NDA discipline

Sell-side data rooms are NDA-gated. Most NDAs include a **standstill** (the buyer cannot make an unsolicited offer for 12–24 months) and an **employee non-solicit**. Buyers' DD teams should know what they have signed; competing on a deal where you are restricted from approaching the team afterwards has its own cost.

CHAPTER 20

Investment Thesis Construction

A falsifiable hypothesis with explicit value levers

3–5	1 page	3
value-creation levers per thesis	target length for the thesis statement	specific things that would falsify it

An investment thesis is not a paragraph in a deck. It is a falsifiable hypothesis with explicit value levers and explicit failure modes. If you cannot state in one page what has to be true for the deal to make money — and what would prove the thesis wrong — you do not have a thesis yet.

A thesis statement, structured

Headline: one sentence on the deal — what it is, what we will do, what we expect to make. **Why this company**: the structural reasons this asset specifically. **Value levers**: 3–5 specific operating improvements with quantified impact. **Risks and mitigants**: the falsifiers. **Base-case returns**: equity in, equity out, IRR, MOIC. **What must be true**: the explicit assumptions baked into the base case.

A thesis that fails to specify which 3–5 operating levers will create the value is just a prayer. A thesis that promises 'multiple expansion' as a primary lever, in a market where multiples are at 15-year highs, is a prayer with bad odds.

Falsifiability — the hardest discipline

The most powerful question on any IC: what would have to happen for this deal to lose money? The answers are usually concrete and small in number — a major customer losing share, a product transition failing, a regulatory change. If the deal team cannot list them, they have not stress-tested the thesis. If they can, the IC then asks how each is mitigated, monitored, or priced.

Part IV

Term Sheets, Cap Tables, and Governance

Structuring, control, and ongoing portfolio management

CHAPTER 21

The Term Sheet

Valuation, control, and the option pool shuffle

$X pre / $Y post	1x non-participating	3:2
the pricing layer	the modern liquidation default	common board-seat ratio in early-stage VC

The term sheet is the negotiation. The 80-page Stock Purchase Agreement that follows is largely the term sheet's clauses written longer. If you are not winning the term sheet you are not winning the deal.

The price layer

Pre-money valuation: agreed value before the new round. **Post-money**: pre-money plus new investment. **Investor's ownership**: investment ÷ post-money. The **option pool shuffle**: increasing the option pool before the investment so that the dilution falls on existing shareholders rather than the new investor. A 5% pre-investment pool top-up shifts ownership materially.

Worked example: pre $40M, $10M raise, $50M post — investor owns 20%. If the founders agree to a $5M pre-money pool top-up, the pre is $40M including the new pool, founders own less, investor still owns 20%. The shuffle is worth real percentage points; founders should price it explicitly.

Liquidation preferences

1x non-participating preferred is the modern default — investors get either their money back or their pro-rata share of equity at exit, whichever is greater. **Participating preferred** ('double-dip') gives the investor their preference and their pro-rata share — far more aggressive. **Higher multiples** (1.5x, 2x) appear in down rounds and bridge financings. Chapter 22 walks the math.

Control: board, protective provisions, drag-along

Board composition: investors typically take 1–2 seats; founders typically retain 2–3 in early stages. **Protective provisions**: investor consent required for material acts (Chapter 26). **Drag-along**: lets a majority force a sale. **Pro-rata rights**: investor's option to participate in future rounds. **Information rights**: monthly financials, annual budget.

Anti-dilution and pay-to-play

Anti-dilution protects investors against dilution in down rounds — broad-based weighted average is standard, full ratchet is aggressive (Chapter 23). **Pay-to-play** converts an investor's preferred to common (or to a less-protected class) if they do not participate in future rounds — pushing investors to support the company through follow-ons rather than walk away.

CHAPTER 22

Preferred Stock and Liquidation Preferences

Modeling who gets paid what across exit prices

1x	Non-participating	$0–$200M+
modern default preference multiple	modern default structure	typical exit-price range over which structure matters

Liquidation preferences quietly determine who actually gets paid in a sale. Two cap tables that look identical in ownership percentages can produce wildly different founder outcomes at exit because the preferred stack is different. Modelling waterfalls is a non-optional skill.

Non-participating preferred

Holder gets either their preference back or their as-converted pro-rata share, whichever is greater. At low exit prices, the holder takes the preference. At high exit prices, the holder converts and takes the pro-rata share. There is a single break-even price above which conversion is preferred.

Example: $20M Series A at 25% post-money ownership ($80M post). At $40M exit, holder takes $20M (preference); at $200M exit, holder converts and takes $50M (25% pro-rata). The break-even is $80M — below it the preference dominates, above it the equity dominates.

Participating preferred

Holder takes their preference and their pro-rata share of remaining proceeds. At every exit price, the participating holder receives more than a non-participating holder would. The cost falls on common (founders and employees).

Modern markets typically resist full participating preferred outside specific situations (down-round bridges, distressed financings); **capped participating** — participation up to a 2x or 3x multiple, then converts to non-participating — is sometimes seen.

Stacked preferences across multiple rounds

After a Series A, B, C, D, the cap table has a stack of preferences. **Pari passu**: all preferences rank equally. **Senior**: the latest series ranks ahead of earlier ones (common in distressed bridges). The stack determines waterfall payment order at any exit price.

How to model it

A waterfall model: at each exit price, walk through the preference stack from senior to junior, paying each preference; then split the remainder between participating preferred (if any) and converting /

common. Build a chart from $0 to $1B exit price showing the payout to each share class. The chart immediately reveals where founder economics are real and where the structure has hollowed them out.

CHAPTER 23

Anti-Dilution Mechanics

Broad-based weighted average versus full ratchet

Broad-based WA	Full ratchet	0%–full repricing
modern default	aggressive — usually distressed-only	the spectrum of dilution protection

Anti-dilution provisions are the rate-of-return protection clause. They protect a preferred investor against dilution if the company raises future capital at a lower price. Two structures dominate; the math behind each tells you what to negotiate.

Broad-based weighted average

If a down round happens, the conversion price of the prior round resets to a weighted average of the old and new prices, weighted by share count. The result: the prior investor's effective ownership increases by a moderate amount that scales with the size of the dilutive issuance.

Formula: New CP = Old CP × (A + B) / (A + C). Where A = shares outstanding pre-issuance, B = shares the new money would buy at the old price, C = shares actually issued. The further the new price is below the old price, and the larger the issuance, the more the CP resets.

Full ratchet

If a down round happens, the conversion price resets to the new (lower) price regardless of the size of the issuance. A small dilutive issuance can dramatically reset the prior round's CP. Founders and employees take all the dilution.

Full ratchet is structural-vandalism aggressive in a normal financing. It appears in bridge rounds, distressed scenarios, or where an investor has unusual leverage. Founders should resist it; if they cannot, they should negotiate a sunset (the ratchet expires after a defined period or a defined good-news event).

What's negotiable in practice

Pay-to-play conditions on anti-dilution (only investors who participate in the down round retain their protection); duration of the protection (full term vs. sunset); and what counts as a 'dilutive issuance' (carve-outs for employee options, M&A consideration, equipment financing). Each of these tilts the protection's bite.

CHAPTER 24

Cap Table Mechanics

Common, preferred, options, SAFEs, and conversion math

Fully diluted	Common, Preferred, Options, SAFEs, Notes	3 conversion mechanics
the only basis that matters in negotiations	five classes to track	to model right

The cap table is the company's ledger of who owns what. A founder running a $30M Series B with no fully-diluted cap table model is flying blind — and most founders, before professional CFO infrastructure, are doing exactly that.

The five classes

Common stock: founders, employees, ex-employees post-vest. **Preferred**: each financing round (Series Seed, A, B, C…) gets its own series with negotiated rights. **Options**: outstanding employee options, both vested and unvested, plus the unallocated reserve. **SAFEs** and **convertible notes**: pre-priced instruments that convert into preferred at the next priced round, with caps and discounts. Track each class separately and reconcile to share count.

Fully diluted share count assumes every option, warrant, SAFE, and note has converted. This is the basis on which ownership percentages are computed. Quoting ownership on issued-and-outstanding (which excludes options and convertibles) is misleading and most LP/diligence teams will catch it.

SAFEs and convertible notes

A **SAFE** (Simple Agreement for Future Equity) converts into preferred at the next priced round. Two key terms: **cap** (the maximum effective valuation at which the SAFE converts) and **discount** (a percentage discount to the next-round price). Y Combinator's post-money SAFE is the modern standard. SAFEs do not accrue interest; convertible notes do.

Conversion math: if a SAFE has a $10M cap and the next round prices the company at $20M, the SAFE converts at $10M — the SAFE-holder gets twice the shares per dollar that new investors get. The post-money cap structure (vs. pre-money) determines whether dilution from SAFEs falls on founders or on new investors.

Modelling a new round

To model a new round: take the pre-money valuation, add any convertibles at their conversion terms, add the option-pool top-up, then compute new investor ownership as (new investment) ÷ (post-money including all of the above). Re-balance the table; reconcile every share class to a final share count. Sophisticated cap-table tools (Carta, Pulley) automate this; the underlying math is something every Series-A and later founder must understand.

CHAPTER 25

Board Composition and Fiduciary Duties

Care, loyalty, Revlon, and the failure modes that destroy value

3–7	Care, Loyalty, Revlon	Caremark
typical private-company board size	the three fiduciary duties	the standard for board-level oversight failure

Boards are the formal corporate authority of a company. PE and VC investors buy board seats because the board — not the shareholder vote — controls every material decision a private company makes between exits.

Board composition

Early-stage VC: 3-person board, founder + investor + independent. Series B: 5-person, two founder, two investor, one independent. Series C+: 5–7-person boards with multiple investor seats and independents. PE buyout: GP-controlled boards, with 2–4 GP seats, 1 management seat, and 1–2 independents (often industry operators). Independents matter — they provide governance and Revlon-duty cover in a sale process.

Fiduciary duties

Duty of care: directors must inform themselves of material facts before deciding. **Duty of loyalty**: directors cannot self-deal or favour one constituency at the expense of another. **Revlon duties**: in a sale of control, directors' duty is to maximise short-term value for shareholders. **Caremark**: directors must establish reasonable systems to monitor risk; failure to do so is itself a breach.

These duties run to the corporation and its shareholders, not to the appointing investor. A VC partner sitting on a portfolio company board cannot favour their fund over other shareholders — including in the most contentious cases (a down round in which the fund leads, a sale at a price that pays the preferred but zeros common).

Governance failure modes

Three patterns recur. (1) **Asleep boards** — directors who don't read the pre-read, don't ask hard questions, ratify management. (2) **Captured boards** — too aligned with founders or with one investor, unable to resist a bad strategic call. (3) **Conflicted boards** — directors voting on transactions where their fund or their portfolio is on the other side of the table. Each is a specific Caremark or loyalty issue and each has been litigated.

CHAPTER 26

Protective Provisions and Control Rights

Consent rights and shareholder agreements in practice

8–15	Series-class vote	Drag-along + ROFR
standard protective provisions	the typical consent mechanic	two more rights worth knowing cold

Protective provisions are the minority investor's veto rights. They define the actions a company cannot take without the investor's consent — even if the founders, the board, and a majority of shareholders agree. They are where minority investors actually have control.

The standard set

Standard NVCA-template protective provisions require investor consent to: amend the certificate of incorporation; create a senior or pari-passu series; redeem or repurchase shares; pay dividends; merge or sell the company; liquidate; change the board size; incur debt above a threshold; alter the protective provisions themselves. Each is designed to prevent a controlling shareholder from acting against the minority's interests.

These are consent rights, not information rights. They are voted as a series-class — Series A holders vote separately as a class on Series A protective provisions, even when they are a minority of overall shareholders.

Drag-along, tag-along, ROFR

Drag-along: lets the majority force minority holders to sell their shares in a sale of the company. Without it, minority hold-outs can block a sale. **Tag-along**: lets minority holders join a sale on the same terms as a controlling holder. Without it, minority holders can be left behind in a partial sale. **ROFR** (Right of First Refusal): the right to match a transfer of shares to a third party. Each manages a distinct conflict between majority and minority.

How they function in practice

Most protective provisions are never invoked. Their value is the option they create — the implicit threat that lets the minority investor be heard at the board, in negotiations, and in side-letter discussions. A board that knows a major action requires Series B consent will ensure Series B is comfortable before bringing it up.

CHAPTER 27

Employee Equity

ISOs, NSOs, 83(b), vesting, and the alignment design problem

ISO vs. NSO	83(b)	4-year cliff
the two US option types — and they are different	the election that saves founders six-figure tax bills	the modern default vesting schedule

Employee equity is one of the most powerful retention and alignment tools in private companies and one of the most poorly explained. Done right, it makes employees part-owners. Done wrong, it produces angry tax surprises and litigation.

ISOs vs. NSOs

Incentive Stock Options (ISOs): tax-favoured, US-only, with strict eligibility rules (employees only, $100K vesting cap per year, 10-year expiry, 90-day post-termination exercise). Properly held, ISO gains are long-term capital gains. **Non-qualified Stock Options (NSOs)**: ordinary-income tax at exercise on the spread between strike and FMV. Anyone can receive NSOs (including contractors and advisors). For most US employees, ISOs up to the $100K cap are preferable; the rest of the grant is NSO.

The 83(b) election

When restricted stock vests, the holder owes ordinary income tax on the spread between strike and FMV at vesting. For early-stage founders this can be small at grant and very large at vesting. The **83(b) election**, filed within 30 days of grant, accelerates the tax to the grant date — paid on a tiny spread, locking in long-term capital-gains treatment for any subsequent appreciation. Missing this 30-day window is the most common preventable tax mistake in early-stage founders' lives.

Vesting schedules

Modern default: **4-year vest with 1-year cliff** — no shares vest in the first year, then 1/4 vests at month 12 and the rest monthly over the remaining 3 years. The cliff protects against early-quitter dilution. **Acceleration** can be single-trigger (full or partial vest on a sale) or double-trigger (vest on a sale and termination by the buyer). Double-trigger is the modern default for senior employees; single-trigger is sometimes negotiated by C-level.

409A valuations and strike pricing

A **409A valuation** is an IRS-blessed independent valuation of the company's common stock, used to set the strike price on options. Without a 409A valuation, options can be deemed under-priced and trigger a 20% federal penalty on the holder. Modern startups refresh 409As annually, after major financings, and at material business changes. Strike-price discipline keeps the option program compliant and the resulting taxation clean.

QSBS — Section 1202 and the founder/early-employee tax shield

Qualified Small Business Stock (QSBS) under IRC §1202 allows a non-corporate taxpayer to exclude from federal income tax up to 100% of the gain on the sale of qualifying stock — up to $10 million or 10 times the taxpayer's adjusted basis in the stock, whichever is greater. The exclusion applies to stock acquired after September 27, 2010 and held for more than five years. For a founder who bought shares at a nominal price and held through a meaningful exit, this provision can shield tens of millions of dollars of capital-gains income at the federal level. Note: state tax treatment varies — California, for instance, does not conform to the federal QSBS exclusion. This is general information, not tax advice; consult a qualified tax advisor for your specific situation.

To qualify, the stock must have been issued by a domestic C-corporation whose aggregate gross assets did not exceed **$50 million** at the time of issuance (and immediately after). The company must be engaged in an active qualified trade or business — certain professional services (health, law, finance), hospitality, and farming businesses are excluded. These requirements mean QSBS is most commonly applicable to venture-backed technology, life sciences, and software companies. When a financing round is planned, founders and early employees should confirm the company's gross-asset position relative to the $50M cap before the round closes; a round that pushes gross assets over the cap does not retroactively disqualify prior stock, but new shares issued post-cap are ineligible.

Practical execution requires attention to three administrative steps. First, track the **exact issuance date** of each tranche of stock — the five-year holding clock starts at issuance, and partial dispositions that trigger the clock on a later-issued tranche are a common pitfall. Second, file the **83(b) election** within 30 days of grant on any restricted stock; this establishes the grant date as the holding-period start and locks in the basis for the exclusion calculation. Third, keep **clean cap-table records** with documentation of the company's gross assets at each issuance date — this is the IRS audit trail. Secondary sales of QSBS before the five-year mark do not qualify for the exclusion, and acquisitions can be structured to preserve or terminate QSBS status depending on how consideration is paid; both scenarios require tax counsel review before execution.

Part V

Value Creation, Exits, and Liquidity Pathways

Operating playbooks, exit routes, secondaries, continuation vehicles

CHAPTER 28

The Operating Playbook

Multiple expansion, leverage paydown, and EBITDA growth

3 levers	4 EBITDA levers	100-day plan
multiple expansion, EBITDA growth, debt paydown	revenue, margin, mix, M&A	where the operating playbook becomes specific

PE returns come from three places, and only three. The buyer pays X today, sells for Y in five years, and the difference came from some mix of multiple expansion (the market revalued the asset), EBITDA growth (the company got bigger or more profitable), and debt paydown (the equity grew because the debt shrank). Reading the returns bridge tells you what kind of fund you are looking at.

The returns bridge

Decompose the equity-value creation: multiple expansion = (exit multiple − entry multiple) × entry EBITDA. EBITDA growth = (exit EBITDA − entry EBITDA) × exit multiple. Debt paydown = (entry debt − exit debt). The three sum to total equity value created (with a small cross-term). Plot it as a stacked bar; every IC sees this chart.

A buy-and-build SaaS deal might bridge: 70% EBITDA growth, 20% debt paydown, 10% multiple expansion. A 2009 vintage industrial deal might bridge: 40% EBITDA growth, 40% multiple expansion (recovering from the downturn), 20% debt paydown. The bridges look like the strategies they came from.

Operating levers, named

Revenue growth: pricing increases, new product launches, geographic expansion, channel expansion, sales-force productivity. **Margin improvement**: procurement, manufacturing efficiency, SG&A reduction, technology automation, working-capital reduction. **Mix**: shifting toward higher-margin product or customer segments. **M&A**: bolt-ons that expand revenue, expand margins, or both. Each lever is specific, measurable, and assigned to an owner in the 100-day plan.

Multiple expansion as strategy

The cleanest source of multiple expansion is positioning a sub-scale asset into a category for which the public or strategic-buyer market awards a higher multiple. Examples: building a software portfolio out of services + tech tuck-ins (services trade at 1x revenue, software at 6x); aggregating regional providers into a national platform (regional 6x EBITDA, national 10x); transforming a B2B mid-market into a recurring-revenue B2B SaaS multi-tenant platform.

But multiple expansion as a standalone thesis is a dangerous bet at most points in the cycle. Buyout multiples have ranged from 7x in 2002 to 14x in 2021. A thesis whose primary lever is 'we will sell at a higher multiple than we paid' is a thesis on the cycle, not on the company.

CHAPTER 29

The 100-Day Plan and Portfolio Governance

Quick wins, reporting cadence, and management upgrades

100 days	4 workstreams	Weekly
the canonical window	people, finance, ops, strategy	the discipline cadence in the first quarter

The 100-day plan is the PE owner's operating contract with the company. It is a written, calendared plan covering the first 100 days — the period during which most of the value-creation thesis is locked in or lost.

The plan, structured

People: assess the management team, decide on upgrades within 60 days, retain key talent, design the equity rollover. **Finance and reporting**: install the monthly reporting pack, the budget process, and the cash-flow forecast. **Operations**: identify quick wins (procurement consolidation, SG&A trims, working-capital tightening) and execute them. **Strategy**: validate or refine the value-creation thesis based on the first 90 days of internal data.

The plan is a written document, owned by the deal partner and the new CFO/CEO jointly. Each workstream has a named owner, a list of milestones, and a weekly review.

The management decision

The most consequential 100-day call is on management. PE-backed companies frequently change CEO or CFO within the first 6 months, often within the first 100 days. This is a hard call when the existing CEO sold the deal. The discipline is to assess against the value-creation thesis specifically — not against general capability — and to act decisively if the answer is no.

Reporting cadence

Modern PE companies operate on monthly closes within 8–10 business days, monthly board packs, weekly KPI dashboards on revenue and key operational metrics, quarterly strategy reviews. The cadence itself enforces discipline: a CEO who has to report monthly on margin progress acts on margin progress.

CHAPTER 30

Exit Planning and Route Selection

Strategic, sponsor-to-sponsor, IPO, dividend recap, and partial sale

4 routes	3–7 years	18 months
strategic, sponsor-to-sponsor, IPO, dividend recap	typical PE hold; longer if exit windows close	lead time for a properly run process

Exit is the only thing that matters for an LP's return. A PE deal that compounds 25% IRR for 4 years and exits at 1x preference because the market closed is not a 25% IRR deal — it is a zero. Choosing the exit route, and the exit timing, is the most consequential strategic decision in any portfolio company's life.

Strategic acquirer

An operating company buys for synergies and capability — typically the highest-priced exit, because the strategic can pay above standalone DCF based on cost synergies and revenue synergies. The downside: strategic buyers move slower, are more selective, and often want a clean integration (which can mean management departures and culture loss).

When to pursue: when there is a credible strategic with a synergy story; when the company has a capability the strategic cannot build organically; when the deal team can run a process narrow enough to attract strategics without it becoming a public auction.

Sponsor-to-sponsor

Selling the company to another PE firm. The dominant exit route in many vintages: roughly 30% of buyout exits in 2024 were sponsor-to-sponsor. The next sponsor is paying for the next leg of value creation; the price reflects what they think they can do, not what the strategic synergy is. Process speed and certainty are the advantages.

IPO

Listing the company on a public exchange. Pricing on the basis of public-comparable multiples; sponsor remains a major shareholder through lockup (typically 180 days) and beyond. Best executed when public-market multiples are favourable, the company has the scale and reporting maturity for public-life, and the lockup overhang is something the market will absorb. Chapter 31 walks the process.

Dividend recapitalisation

Distributing cash to equity holders by issuing new debt against the portfolio company. A partial monetisation that does not exit the position. Used when the company can support more leverage than at acquisition, exit windows are closed, and the GP wants to return cash to LPs. The new debt re-clocks the leverage profile and re-positions for a future exit.

SPV mechanics for partial exits and direct co-investment

A **Special Purpose Vehicle (SPV)** is a single-purpose legal entity — typically a Delaware LLC — formed to hold a single investment. In the late-stage and co-investment context, the SPV is the dominant wrapper for allowing LPs or third-party investors to concentrate exposure into one portfolio company without joining the main fund. The two most widely used administration platforms are **AngelList Stack** and **Sydecar**, each of which automates the SPV formation, subscription, and K-1 generation workflow. Lead economics are typically 10–20% carry on the SPV (with 20% being common for high-conviction leads) and an administration fee of 0.5–2% of committed capital, paid once at close. LPs participating in an SPV accept concentrated single-name exposure in exchange for the right to size up beyond their pro-rata in the main fund.

SPVs are increasingly used as a **partial-exit tool** when a fund needs to show distributions per invested capital (DPI) but a portfolio company is not ready for a full sale. The GP sells a portion of the fund's position into a newly formed SPV — either to existing LPs at a negotiated price or to a secondary buyer — generating cash distributions to the main fund while retaining upside through the remaining position. This structure is also used in continuation vehicles (Chapter 33), but at the individual SPV level it is simpler: the main fund receives cash, the SPV holder receives the residual equity, and both parties have consented to a valuation that cleared in a market process.

Two conflict-of-interest watchpoints govern SPV use at institutional-quality managers. First, **cross-fund SPVs**: when a GP uses a separately capitalized SPV to invest alongside a fund into the same company, the GP must manage the allocation between vehicles carefully — any pricing difference between the SPV and the fund entry is a conflict that LPAC review should catch. Second, **GP-affiliate SPVs**: if the GP or a GP-related party is a significant investor in the SPV (beyond the standard carry interest), the conflict is structural and requires disclosure. The SEC's Form ADV Part 2 and the LPA's related-party provisions are the primary disclosure framework; fund counsel and LPAC chairs are the practical enforcement layer.

CHAPTER 31

The IPO Process and Sponsor Roles

S-1, roadshow, pricing, lockup, and overhang

S-1	180 days	12–24 months
the registration document	standard lockup	typical IPO sell-down completion timeline

An IPO is a public-market financing event that the sponsor structures, prices, and lives with for two years afterward. The sponsor's role is dual: seller of the IPO shares, and ongoing controlling shareholder responsible for governance through the lockup and beyond.

S-1 and the registration process

The S-1 is the SEC registration statement, 200+ pages of business description, MD&A, audited financials, risk factors, and detailed disclosures. Drafting takes 4–6 months with the sponsor's underwriting bank, the issuer's legal counsel, and the auditor. The first amendment incorporates SEC comments; the final amendment becomes the prospectus on pricing day.

Roadshow and pricing

Once the S-1 is on file, the lead underwriters take the management team on a 2-week global roadshow — meeting institutional investors who will anchor the offering. Indications of interest build the order book; the underwriters and issuer set the price range, then narrow to a final price the night before trading. Underpricing — first-day pop — is structural in IPOs (issuer leaves money on the table; underwriters allocate to favoured clients).

Lockup and overhang

Insiders (sponsor, founders, employees) typically agree not to sell shares for 180 days. The lockup expiration is a known supply event. Sophisticated sponsors plan the lockup expiration around earnings cycles and market conditions. The overhang — public knowledge that a large pre-IPO holder will eventually sell — depresses the post-lockup share price. Modern best practice: structured selling programs (Rule 144), block trades to crossover funds, secondary offerings with market-friendly pricing.

CHAPTER 32

The Secondary Market

LP-led portfolio sales and how a 90¢-on-the-dollar deal gets priced

$130B+	85–95¢	70–80¢
2024 LP-led secondary volume	typical pricing on quality buyout books	stressed-book pricing

The secondary market is where LP fund interests trade hands. A 2024 secondary market did over $130B in volume. Pricing is a function of NAV, vintage, sector exposure, and the buyer's view of the underlying portfolio.

LP-led portfolio sales

An LP selling its fund interest to another investor. The seller wants liquidity; the buyer wants exposure to a particular vintage or strategy at a price below NAV. The discount reflects: (1) the buyer's underwriting of the underlying portfolio, (2) the time-value of the remaining fund life, (3) the expected fee load over the remaining life, (4) the buyer's risk appetite for the specific GP and strategy.

How pricing actually works

A secondary buyer underwrites every position in the fund, models the expected exit value and timing, applies their target IRR (typically 12–18% gross), and computes the maximum bid that hits the target. This produces a number expressed as a percentage of NAV: 90¢ on the dollar, 75¢ on the dollar. A clean book of 2018 vintage buyout funds might trade at 92–95¢; a stressed book of 2020-vintage venture funds might trade at 60–70¢.

Why the market grew to $130B

Three drivers. (1) **LP rebalancing**: the denominator effect (Chapter 35) and pacing-model adjustments push LPs to manage exposure actively. (2) **Portfolio cleanup**: large LPs prune tail-end positions to free up team time and reporting load. (3) **The exit drought**: 2022–2024 closed the IPO and M&A windows, choking distributions; LPs increasingly turned to secondaries for synthetic liquidity. The market is now a structural feature of the industry, not a niche.

CHAPTER 33

Continuation Vehicles and GP-Led Secondaries

Mechanics, conflicts, LPAC consent, and the $100B-a-year market

$100B+	LPAC	3 conflicts
2025 CV volume	the consent body for every CV	to manage cleanly

Continuation vehicles — the GP-led secondary structure — are the most important structural innovation in PE since the 2009 fund-vehicle reforms. They let a GP hold a high-conviction asset past the original fund's life, with new LP capital and reset economics. Done right, they extend the runway on the best assets. Done badly, they are conflicted and lawsuit-prone.

How a CV is structured

The existing GP forms a new fund (the continuation vehicle). The CV buys one or a few assets from the original fund at an agreed price. Existing LPs choose: **cash out** at the agreed price, or **roll over** into the CV with reset terms (often a fresh hurdle and a new carry crystallisation). New LPs (typically secondaries-focused funds) provide most of the CV's capital. Independent valuation and LPAC consent are universally required.

The CV is run by the same GP team that managed the original fund's investment. The thesis is that they know the asset best and the additional 3–5 year hold can compound value at a different rate than a forced sale would.

The three core conflicts

Pricing: the GP sits on both sides of the trade — selling for the old fund, buying for the new. Independent valuation, LPAC consent, and a third-party fairness opinion are the standard safeguards. **LP fairness**: existing LPs who cash out and existing LPs who roll over are taking different bets on the same assets. The 'status quo' option (rolling on the same terms) is increasingly mandated. **GP economics**: the GP can crystallise carry on the original fund's exit price, then re-earn carry on the CV. ILPA's 2023 guidance on CVs sets norms for how this should be structured.

The 2024–2025 CV market

CVs grew from a niche $5B market in 2018 to over $100B in 2025, driven by the same exit drought that grew the LP-led secondary market. Single-asset CVs — concentrated bets on one trophy asset — now account for a meaningful share of volume. Independent fairness opinion providers (Houlihan Lokey, Lazard) have built dedicated CV practices. The market is now mature enough that LPs treat well-structured CVs as a normal liquidity option, not an exotic event.

Part VI

Advanced Topics and Expert Casework

Market cycles, sector investing, IC memos, integrated cases

CHAPTER 34

Sector Investing — AI, Deep Tech, and Space

Three sector theses and what makes each different

37%	10–15 yr	3 sector cases
of 2024 global VC funding that went to AI (CB Insights)	typical deep-tech development cycle	covered here: AI/ML, deep tech, space

Sector investing is a structural choice. A sector-thesis fund concentrates capital and brand into a narrow domain in exchange for informational edge and founder-side reputation. The trade-off is concentration risk — and getting the timing of a sector wrong can cost a vintage.

AI / ML — the 2024–2025 megacycle

The 2024 venture market saw extraordinary concentration: AI-related companies received roughly 37% of all global venture funding, the highest sector concentration since the 1999–2000 internet wave. Within AI, capital split between **infrastructure** (foundation-model labs, GPU cloud, training data) and **application** (vertical AI products, agentic systems). The infrastructure layer is dominated by a small number of capital-intensive players; the application layer is fragmented and competitive.

Diligence questions specific to AI: (1) **Model moat** — does this company have a structural advantage that survives the next foundation-model release? (2) **Data network effects** — does usage create proprietary training data that improves the product? (3) **Cost structure** — what is gross margin after model inference cost, and how does it evolve as token prices fall? (4) **Go-to-market** — is this a product the customer's existing IT can adopt, or does it require a multi-quarter integration?

Deep tech — long timelines, non-dilutive capital

Deep-tech companies (advanced manufacturing, biotech, novel materials, fusion, quantum) require 10–15 year development cycles and capital structures designed for that. The capital stack often includes **government grants** (DARPA, ARPA-E, NSF), **strategic partners**, **tax credits** and **venture debt** alongside primary equity. The diligence is technical: dedicated technical advisors, lab visits, and milestone-based financing rounds rather than time-based ones.

Commercial space

From 2014 onward, the commercial-space ecosystem has grown along three vectors: **launch** (SpaceX dominant; Rocket Lab, Blue Origin, others competing), **earth observation** (Planet Labs, BlackSky, Maxar), and **satellite internet** (Starlink, OneWeb, Amazon Kuiper). Each has different capital intensity, different revenue models, and different defence/dual-use dynamics. The sector's investability has expanded as launch costs fell from ~$10K/kg to ~$2K/kg.

Four worked sector theses — embedded finance, defense tech, applied AI infrastructure, and value-based healthcare

Embedded finance. The secular driver is the shift of financial services delivery from bank-owned channels to the software layer inside non-financial products — payroll, fleet management, e-commerce, and vertical SaaS platforms that now distribute lending, insurance, and payments natively. The structural advantage today is that Banking-as-a-Service (BaaS) regulatory infrastructure has matured enough to make sponsor-bank partnerships commercially viable at scale without a full bank charter. The target archetype is a vertical SaaS company with 5,000+ SMB customers that has not yet monetised its payment or credit surface — where the GP can fund the embedded-product build and underwrite the attach-rate expansion. The thesis is wrong if the BaaS regulatory environment tightens materially — specifically if sponsor-bank arrangements come under the kind of enforcement that has already hit several BaaS-heavy banks — in which case the cost and time to stand up compliant infrastructure rises significantly and the model does not close at the target margin.

Defense tech. The secular driver is the shift of US defense procurement from the traditional prime-contractor model toward commercially developed dual-use technology — autonomous systems, software-defined communications, AI-enabled ISR, and low-cost attritable platforms — accelerated by the lessons of the Ukraine conflict and the NDAA reform agenda. The structural advantage is that DoD Other Transaction Authority (OTA) contracting now provides a faster acquisition pathway that commercial-stage companies can navigate without a decade-long prime relationship. The target archetype is a founder-led company with a working prototype and an active DoD program-of-record sponsor at the GS-15 or SES level, raising a Series A or B on the basis of a signed CRADA or OTA. The thesis is wrong if procurement reform stalls and traditional prime incumbents successfully defend the acquisition process against commercial entrants — timeline risk is the dominant killer in this sector.

Applied AI infrastructure. The secular driver is the 2023–2025 explosion in enterprise demand for AI inference and fine-tuning capability outside of the hyperscaler APIs — driven by data-sovereignty requirements, latency constraints, and the margin math of running large workloads at scale against per-token pricing. The structural advantage is that the open-weight model ecosystem (Llama, Mistral, and successors) has created a supply of high-quality base models that infrastructure companies can build on without foundation-model-level capital. The target archetype is a company building the tooling layer — inference optimisation, model routing, fine-tuning pipelines, or vector retrieval infrastructure — that enterprises need to run AI workloads on their own compute at costs below hosted API pricing. The thesis is wrong if the hyperscalers discount their hosted API pricing aggressively enough to eliminate the unit-economics advantage of self-hosted inference — which is the single most important variable to monitor in diligence. See the sector thesis builder in this report's lab for a structured framework to stress-test any thesis against this kind of disqualifier.

Value-based healthcare. The secular driver is the long-run shift of US healthcare reimbursement from fee-for-service (pay per procedure) to value-based contracts (pay for outcomes) — a shift the CMS has been legislating since the ACA and which commercial payers have increasingly adopted for primary care, oncology, and chronic disease management. The structural advantage is that the enabling technology — risk stratification, care-gap identification, remote monitoring — has become cheap enough to deploy at the primary-care group level without hospital-system infrastructure. The target archetype is a primary care or specialty group with 50–300 attributed lives under value-based contracts already, a demonstrated MLR improvement over the fee-for-service baseline, and a management team willing to accept PE governance in exchange for the capital to expand the attributed-life base. The

thesis is wrong if the company's MLR improvement cannot be replicated across new patient panels — if the performance is physician-specific rather than model-specific — in which case the roll-up economics collapse as soon as the founding physicians depart.

CHAPTER 35

Market Cycles, Vintage Years, and Macro Sensitivity

Why the year you deploy matters more than the deal you do

J-curve	Dry powder	Denominator effect
the fund-life return signature	the supply side of deal pricing	what 2022 did to LP allocations

Vintage year matters more than almost any other portfolio decision. A 2009 buyout vintage and a 2007 buyout vintage are barely the same asset class. Capital deployed into a recovering market behaves entirely differently from capital deployed at a peak.

Why vintage year matters

The capital deployed in a fund's investment period is priced at the multiples and macro conditions of those specific years. A 2007 vintage bought at peak buyout multiples on cyclical assets just before the GFC had a ferocious time-to-recovery; a 2009 vintage deployed into discounted assets had a tailwind. The vintage is not a label; it is a macro snapshot that the fund cannot escape.

Sophisticated LPs deliberately diversify across vintages — committing every year, in roughly equal amounts, to smooth the entry-multiple cycle. This **vintage diversification** is a core LP-side discipline.

Dry powder and pricing

Dry powder is committed but uncalled capital across the industry. Going into 2024, global PE dry powder exceeded $2.5T. High dry powder pushes deal pricing up: more capital chases the same number of high-quality assets. Falling dry powder loosens pricing. The dry-powder cycle shapes vintage-year quality.

The denominator effect

When public-equity markets fall, an LP's PE NAV (as a share of total portfolio) rises mechanically — the denominator (total assets) shrank, the numerator (PE NAV) didn't. LPs constrained by allocation policies (5% PE, 10% PE, etc.) become over-allocated and must slow new commitments. This was the dominant force on 2022–2023 LP behaviour, and it propagated through the GP fundraising market for 18 months.

J-curve and the fund-life cash-flow signature

Early in a fund's life, fees are paid before exits arrive — the fund's net cash position is negative (the J's bottom). As exits begin, distributions exceed remaining calls, and the fund's net cash position turns sharply positive (the J's hook). The depth of the J-curve depends on fee load, deployment pace, and time-to-first-exit. Long-hold strategies (infra, growth) have steeper J-curves than fast-deploying buyouts.

CHAPTER 36

Legal and Regulatory Foundations

Advisers Act, Rule 506, ERISA, AIFMD, and the 2023 Private Fund Adviser Rules

Investment Advisers Act	Rule 506(b) / 506(c)	AIFMD
the foundational US framework	the exemptions PE funds rely on	the EU equivalent

PE/VC operates inside a thicket of US and global regulations. Most of the time these are background. Occasionally they are everything — a registration deadline missed, a CFIUS filing botched, a Rule 506 marketing rule breached, and the consequences are existential.

The Investment Advisers Act

GPs managing more than $150M in private fund assets must register with the SEC as Investment Advisers. Registration imposes obligations: a **Form ADV** public disclosure document, **compliance policies and procedures**, an annual **compliance officer review**, and book-and-records requirements. The 2023 **Private Fund Adviser Rules** (since partially vacated by the Fifth Circuit, but with much of the LP-protection content adopted in practice anyway) added: standardised quarterly statements, mandatory independent audit, required side-letter disclosure, and restrictions on certain LP-disadvantaging practices.

Rule 506 exemptions

PE funds raise capital under Regulation D's Rule 506 exemption from securities registration. **506(b)** permits private offering with no general solicitation. **506(c)** permits general solicitation but requires verification that all investors are accredited. Both require LPs to be qualified — typically **accredited investors** (income, net-worth, or institutional thresholds) and, for $5M+ fund-level pools, **qualified purchasers** under Section 3(c)(7) of the Investment Company Act.

ERISA and CFIUS

ERISA applies if 25% or more of fund commitments come from US employee-benefit plans. ERISA-covered plan assets impose fiduciary duties on the GP that go beyond standard partnership duties. Most US PE funds either stay below the 25% threshold (avoiding ERISA) or operate as 'Venture Capital Operating Companies' (an ERISA exemption). **CFIUS** (Committee on Foreign Investment in the United States) reviews foreign investment in US businesses with national-security implications. The 2018 FIRRMA expansion broadened CFIUS jurisdiction to more sectors and more investor types; non-US LPs in US PE funds are now routinely subject to CFIUS analysis on individual deals.

AIFMD and the EU regime

The EU's Alternative Investment Fund Managers Directive (AIFMD) regulates fund managers with European investors or assets. AIFMD requires authorisation, capital adequacy, depository requirements,

and detailed reporting. For PE managers, the practical operational requirement is the **private placement regime** (national private placement under each member-state's rules) or full **AIFMD passport**.

CHAPTER 37

The Investment Committee Memo

Structure, voice, and the difference between a recommendation and a wish

12–25 pages	1-page summary	3-section structure
typical memo length	the part that gets read first	is dominant — recommendation, body, appendices

An IC memo is the deal team's contract with the firm. It states the recommendation, the thesis, the price, the risks, and the open issues. Senior partners read for the things that aren't said. Reading bad IC memos is the single fastest way to learn what a good one looks like.

Standard structure

Recommendation on the first page: invest / pass / further work, with the proposed terms (entry multiple, equity cheque, debt structure, IRR target). **Executive summary**: 1 page on the company, market, thesis, returns, risks. **Body**: company overview, market analysis, financial review, valuation, value-creation plan, risks and mitigants, deal structure. **Appendices**: financial model output, comp tables, management assessments, reference reports.

The first page is the part that matters most. Senior partners decide whether to engage in the second hour of reading based on the first page. A muddled first page is a lost deal.

The voice that travels

IC memos earn their authority by being specific and concrete. A claim like 'we believe the company can grow EBITDA 15% per year' is weaker than 'pricing has trailed inflation by 200 bps annually for three years; modest re-pricing aligned with peers and the announced new SKU ladder produces 4–6% revenue and 8–10% EBITDA growth in the base case.' The first version is hope; the second is mechanics.

Avoid: passive voice ('it is believed'), weasel adjectives ('attractive market', 'compelling growth'), and unsupported quantitative claims. Senior partners are trained to read for these and they devalue the entire memo.

The risk section

The risks-and-mitigants section is the second part senior partners read carefully. A memo with three concrete risks and three concrete mitigants is more persuasive than a memo with twelve generic risks. Each risk should specify: what could go wrong, how likely it is, what the impact on returns would be, what the mitigant is. Generic 'macro risk' or 'execution risk' tells the IC nothing.

CHAPTER 38

Integrated Expert Casework

An LBO, a Series B, and an LP secondary — beginning to end

3 cases	Beginning to end	All chapters
LBO / VC Series B / LP Secondary	from sourcing to exit	applied

The previous 37 chapters covered components. This chapter integrates them across three end-to-end cases — an LBO, a Series B, and an LP secondary. Each case reuses concepts from prior chapters and shows what professional fluency actually looks like.

Case 1: B2B software LBO

Setup: $400M EV acquisition of a vertical-market SaaS company, 30% EBITDA margins, 12% revenue growth. Sourced via a sponsor-to-sponsor process. Diligence (Chapters 14–20): commercial DD with cohort retention analysis confirming 115% net revenue retention; QofE normalising EBITDA from $35M to $33M; comps of 12x EBITDA / 6x revenue; LBO model targeting 20% IRR over 5 years.

Structure (Chapters 21–27): $400M total ($150M equity, $200M term loan, $50M unitranche). Management rolls $20M; sponsor commits $130M. Standard PE governance package; 3 sponsor seats, 1 management seat, 1 industry-operator independent.

Operating plan (Chapters 28–29): 100-day plan focused on sales hiring (12 new AEs), product expansion (one bolt-on in adjacent vertical), and pricing (flat 5% YoY across base). Budget targets EBITDA growth from $33M to $58M over 5 years.

Exit (Chapters 30–31): Year 5 exit at $720M EV (12.4x exit EBITDA on $58M, with multiple expansion to 12.4x). Equity proceeds $470M to sponsor + management. Sponsor IRR ~26%. Returns bridge: 60% EBITDA growth, 25% multiple expansion, 15% debt paydown.

Case 2: Series B at $80M post-money

Setup: an AI-application company raising $20M at $80M post-money. Investor lead, 25% target ownership, 1x non-participating preferred. Diligence focused on data network effects, model-moat, gross margin trajectory.

Term sheet (Chapters 21–27): $20M / $80M post; 1x non-participating preferred; broad-based weighted average anti-dilution; 4-year founder vesting with 1-year cliff and 12-month double-trigger acceleration; 2 founder + 2 investor + 1 independent board; pro-rata rights, ROFR, drag-along, standard NVCA-template protective provisions.

Cap-table modelling (Chapter 24): pre-money $60M, with $5M pre-investment option-pool top-up. New investor 25%, founders 47%, employees (option pool) 18%, prior investors 10%. Reconcile fully diluted to ~5.6M shares.

What good looks like at exit: a 5x exit at $400M would return the lead investor $100M (5x on $20M), founders ~$190M, employees ~$72M, prior investors ~$40M. The 1x non-participating preference does not constrain because the pro-rata share dominates at this exit price.

Case 3: LP secondary on a 2018-vintage buyout fund

Setup: an LP wants to sell its $50M commitment / $44M called / $42M NAV interest in a 2018-vintage US mid-market buyout fund. The fund has 9 remaining portfolio companies and 4 years of fund life left. Sponsored by a Tier-1 GP with strong DPI track record.

Pricing (Chapter 32): underwriting buyer projects exit values across all 9 positions, applies 14% target IRR, computes net cash flows including remaining fees, and arrives at a price representing 91¢ on NAV — $38.2M for a $42M NAV. Discount reflects 4 years of remaining fees, exit-timing uncertainty, and 2018 vintage's full deployment into pre-2022 multiples.

Process: secondary advisor runs a process with 8 secondary buyers; bids range from 84¢ to 92¢; the LP transacts at 91¢ with the highest bidder. The transferring LP signs over its commitment; the buyer assumes the remaining unfunded $6M and the future fee/carry stack.

Career pivots into PE/VC — three patterns that work and one that doesn't

DoD / national security → defense-tech VC. Military and intelligence community alumni who reach the O-5/O-6 or GS-15/SES level carry something most VC investors lack: a working understanding of how the DoD actually buys things, who the real decision-makers are inside a program office, and what capability gaps look like from the operator's side. Firms like Shield Capital, Lux Capital, and In-Q-Tel have built explicit pipeline programs for this transition. The pattern works because the informational edge is structural — a former program manager can evaluate a defense-tech startup's government go-to-market faster and more accurately than most generalist investors — and because the dual-use technology wave means there is now a large enough deal universe to justify dedicated sector exposure.

Enterprise sales → growth-stage VC. A senior enterprise sales executive — VP of Sales or CRO at a public or late-stage private technology company — brings a network of CIOs and procurement decision-makers who are exactly who growth-stage portfolio companies need introductions to. The transition works best at the growth-equity or late-stage VC tier, where the value-add is commercial rather than analytical. The pattern has produced a number of successful platform partners and operating advisors at firms like Salesforce Ventures, Bessemer, and Iconiq. The entry point is often an operating partner or venture partner role before a full GP track, precisely because the background is commercial rather than financial — the modelling fluency usually needs to be built explicitly.

Operator / founder → seed VC. A founder who has built and exited a company — even at a modest outcome — carries credibility with other founders that is nearly impossible to manufacture from a banking or consulting background. At the seed stage, where the investment decision is almost entirely a people and thesis judgment, founder-investors have a natural sourcing and signaling advantage: their name on a cap table tells the next round's lead that someone who has done this before believes in the company. The most common entry path is a scout role or angel investing using the exit proceeds, followed by a fund-I raise on the back of a small but demonstrable portfolio. The pattern's constraint is

check size: founder-investors often raise underpowered first funds ($20–50M) and face the strategy question — stay seed-focused or move later — by fund II.

The pattern that doesn't work: generalist consultant without domain edge. A management consultant from a brand-name firm, with no operating experience and no defined sector expertise, who targets a front-office role at a top-quartile buyout or growth fund directly — without an MBA or a clear operating story — is presenting a profile that these funds receive hundreds of times per year and select almost zero times. The pattern fails for a specific reason: consulting builds process fluency and a clear writing style, but it does not build the domain depth, the network, or the demonstrated judgment that PE/VC firms pay for at the senior associate and VP level. The path from consulting to PE almost universally requires an intermediate step — either an MBA from a target school (which resets the recruiting process) or 2–3 years of operating experience that creates a genuine domain story. Trying to skip that step by targeting funds outside the top tier and hoping to lateral is possible, but the landing point is usually a lower-quality platform where the trajectory is harder. Use the job fit tool in this report's lab to map your current profile against the specific role type you are targeting.

REFERENCE

Glossary

Terms organised by thematic cluster.

Fund Structure and Legal

LP
Limited partner — the investor in a fund. Provides capital, has limited liability and limited control.

GP
General partner — the entity that manages the fund and makes investment decisions. Typically a small group of partners with personal liability for the GP entity.

LPA
Limited Partnership Agreement — the constitutional document of the fund. Defines economics, governance, term, and the rules under which everyone operates.

Management Company
The operating entity that employs the investment professionals and receives the management fee. Distinct from the GP entity.

Parallel Fund
A second fund vehicle that invests pari passu with the main fund, used to accommodate LPs with regulatory or tax constraints (e.g. ERISA, offshore).

Feeder Fund
A vehicle that aggregates investors (often offshore) and invests as a single LP into the master fund.

LPAC
Limited Partner Advisory Committee — a small group of LPs that advises on conflicts and gives consents called for by the LPA.

Key-Man Clause
Provision that suspends the investment period if a named partner departs or stops devoting time to the fund.

No-Fault Divorce
LP right to remove the GP without cause, typically requiring a supermajority (e.g. 75%) and triggering carry forfeit consequences.

Clawback
The GP's obligation to return previously distributed carry if final fund returns fall short of the deal LPs were promised.

PPM
Private Placement Memorandum — the offering document used to market the fund to qualified investors.

Subscription Agreement
The contract by which an LP commits capital to the fund and makes regulatory representations.

Fund Economics

Committed Capital
The amount each LP has agreed to invest over the life of the fund.

Called Capital
The portion of committed capital that has actually been drawn by the GP.

Management Fee
Annual fee paid to the management company, typically 1.5–2.0% of committed capital during the investment period and stepping down on invested capital after.

Carried Interest
The GP's share of profits, almost always 20% above a hurdle rate. The economics that make the asset class.

Hurdle Rate
The minimum return LPs receive before the GP earns carry — typically 8% IRR. Also called the preferred return.

Catch-Up
After the hurdle is met, distributions go entirely or 80/20 to the GP until the GP has caught up to a 20% share of total profits.

American Waterfall
Deal-by-deal carry — the GP earns carry as each deal exits, subject to clawback. Standard in US buyout.

European Waterfall
Whole-fund carry — LPs must receive their full committed capital plus the hurdle before any carry is paid. Standard in venture and Europe.

GP Commit
The capital the GP itself contributes to the fund, typically 1–5% of total commitments. A signal of alignment.

Recycling
Re-investing realized proceeds back into new investments, allowing the fund to deploy more than its committed capital over time.

Performance Metrics

IRR
Internal Rate of Return — the annualized discount rate at which the NPV of cash flows equals zero. Time-weighted but distorted by deal pacing.

Net IRR
IRR after management fees, expenses, and carry. The number LPs actually earn.

TVPI
Total Value to Paid-In capital — the sum of distributions and remaining NAV divided by called capital. The total return multiple.

DPI
Distributions to Paid-In capital — cash actually returned to LPs divided by called capital. The realized portion of TVPI.

RVPI
Residual Value to Paid-In capital — unrealized NAV divided by called capital. Marked at fair value, contestable.

PME
Public Market Equivalent — Kaplan-Schoar style benchmark that compares fund cash flows against a public index. The honest test of alpha.

J-Curve
The early-life negative net cash position of a fund as fees are paid before exits arrive. Steeper for longer-hold strategies.

Deal Mechanics — PE

LBO
Leveraged buyout — acquiring a company using a mix of equity and (typically substantial) debt.

EBITDA
Earnings before interest, taxes, depreciation, and amortization. The default cash-flow proxy in PE valuation.

Quality of Earnings (QofE)
A diligence study that normalizes EBITDA for non-recurring items, accounting policies, and run-rate adjustments.

Entry Multiple
Acquisition enterprise value divided by EBITDA at close.

Returns Bridge
Decomposition of LBO returns into multiple expansion, EBITDA growth, and debt paydown.

Rollover Equity
Existing management equity that converts into the new capital structure rather than cashing out at close.

Dividend Recap
Distributing cash to equity holders by issuing new debt against the portfolio company. A partial monetization.

Deal Mechanics — VC

Pre-Money Valuation
The agreed value of the company before the new investment is added to the cap table.

Post-Money Valuation
Pre-money plus the size of the new round.

Option Pool Shuffle

Increasing the option pool pre-investment so the dilution falls on existing shareholders, not the new investor.

SAFE
Simple Agreement for Future Equity — Y Combinator's unpriced instrument that converts at the next priced round, often with a cap and/or discount.

Convertible Note
Short-term debt that converts into equity at the next priced round, typically with interest, a cap, and/or a discount.

Pro-Rata Right
An investor's right to participate in future rounds in an amount sufficient to maintain its ownership percentage.

Pay-to-Play
A provision that converts an investor's preferred shares to common (or to a less-protected class) if they fail to participate in a future round.

Security Types and Cap Table

Preferred Stock
Equity senior to common, with negotiated rights: liquidation preference, dividends, conversion, voting, and protection.

Liquidation Preference
The amount preferred holders receive ahead of common at exit. Quoted as a multiple of original investment (e.g. 1x, 1.5x, 2x).

Participating Preferred
Preferred that takes its preference and then also participates pro rata in remaining proceeds with common. 'Double-dip'.

Non-Participating Preferred
Preferred that takes the greater of its preference or its as-converted share — the modern default in venture.

Anti-Dilution — Broad-Based Weighted Average
Adjusts the conversion price modestly for a down round, weighted by the size of the dilutive issuance relative to the existing share count.

Anti-Dilution — Full Ratchet
Resets the conversion price to the lowest price of the new issuance regardless of size. Aggressive.

409A Valuation
An IRS-blessed independent valuation of a private company's common stock, used to set option strike prices.

Governance and Control

Board of Directors
The body with formal corporate authority. Investors typically buy seats as part of the term sheet.

Observer Rights
The right to attend board meetings without voting — common for non-lead investors.

Protective Provisions
Consent rights over material corporate actions (issue new senior security, sell, dissolve, change board).

Drag-Along
Right that lets a majority force minority holders to join in a sale — clears the path to an exit.

Tag-Along
Right that lets minority holders join a sale on the same terms as a controlling holder. Anti-freeze-out protection.

ROFR
Right of First Refusal — the right to match a transfer of shares to a third party.

Fiduciary Duty
Directors' obligations to act in the corporation's best interests: duty of care, duty of loyalty, and Revlon duties in a sale.

Diligence and Process

CIM
Confidential Information Memorandum — the sell-side marketing document that describes the company, market, and financials.

NDA
Non-Disclosure Agreement — required before a CIM or data room is shared.

Data Room
A secured online repository of diligence documents, organized by workstream.

LOI
Letter of Intent — non-binding (mostly) indication of price and key terms, often paired with exclusivity.

Reps and Warranties
Statements of fact about the company in the purchase agreement, with indemnification if proven false.

Earn-Out
Contingent purchase price tied to post-close performance, used to bridge valuation gaps.

Exit and Liquidity

Strategic Acquirer
An operating company buying for synergies and capabilities — often the highest-priced exit.

Sponsor-to-Sponsor (S2S)
Sale of a portfolio company from one PE firm to another. The dominant exit route in many vintages.

IPO

Initial Public Offering — listing the company's shares on a public exchange.

Lockup
Post-IPO period (typically 180 days) during which insiders agree not to sell. Its expiry is a known supply event.

LP-Led Secondary
An LP selling its fund interest to another investor, allowing early liquidity at a market-clearing price (often a discount to NAV).

GP-Led Secondary
A GP-initiated transaction — most commonly a continuation vehicle — that extends hold on selected assets.

Continuation Vehicle (CV)
A new fund structure formed by the existing GP to acquire one or a few assets from the prior fund, with new LP capital and reset economics.

Macro and Market

Vintage Year
The calendar year in which a fund holds its first close. The key cohort dimension for benchmarking.

Dry Powder
Committed but uncalled capital across the industry — the supply side of deal pricing.

Denominator Effect
When public market drawdowns inflate the relative weight of private holdings in an LP's portfolio, forcing PE selling or pacing changes.

Zombie Fund
A fund past its term that cannot exit assets and continues collecting fees — a governance failure.

Mark-to-Market
Carrying assets at fair value (ASC 820) rather than at cost. The accounting basis for NAV.

REFERENCE

Bibliography

Sources, organised by category.

Academic Foundations

- Gompers, Paul, and Steven N. Kaplan. Advanced Introduction to Private Equity (Edward Elgar, 2022).
- Gompers, Paul, Victoria Ivashina, and Richard Ruback. Private Equity: A Casebook (Imperial College Press, 2019).
- Metrick, Andrew, and Ayako Yasuda. Venture Capital and the Finance of Innovation (Wiley, 3rd ed., 2021).
- Kaplan, Steven N., and Per Strömberg. "Leveraged Buyouts and Private Equity." Journal of Economic Perspectives 23(1), 2009.
- Harris, Robert, Tim Jenkinson, and Steven N. Kaplan. "Private Equity Performance: What Do We Know?" Journal of Finance 69(5), 2014.
- Lerner, Josh, Antoinette Schoar, and Wan Wongsunwai. "Smart Institutions, Foolish Choices: The Limited Partner Performance Puzzle." Journal of Finance 62(2), 2007.

Industry Bodies and Standards

- ILPA. ILPA Principles 3.0 (2019). https://ilpa.org/ilpa_principles/
- ILPA. Reporting and Fee Disclosure Templates. https://ilpa.org/templates/
- NVCA. Model Legal Documents (2023 revision). https://nvca.org/model-legal-documents/
- CAIA Association. CAIA Level I and II Curriculum — Private Equity & Private Debt track. https://caia.org/
- CFA Institute. Private Equity Certificate Curriculum. https://www.cfainstitute.org/

Market Data and Research

- CB Insights. State of Venture 2024 Report. https://www.cbinsights.com/research/report/venture-trends-2024/
- Preqin. Global Private Equity & Venture Capital Report (annual). https://www.preqin.com/
- Pitchbook. Annual VC, PE, and M&A Reports. https://pitchbook.com/
- Bain & Company. Global Private Equity Report (annual). https://www.bain.com/insights/topics/global-private-equity-report/
- Jefferies. Global Secondary Market Review (semi-annual).
- Greenhill / Lazard. Secondary Market Reviews.

Legal and Regulatory

- U.S. SEC. Investment Advisers Act of 1940 and Private Fund Adviser Rules (2023). https://www.sec.gov/
- Dechert LLP. "GP-led Secondaries and Continuation Vehicles" OnPoint, 2025.
- Cooley LLP. Q4 2024 Venture Financing Report. https://www.cooley.com/
- Gunderson Dettmer. Annual Trends in Venture Capital Financing. https://www.gunder.com/
- Wilson Sonsini. Entrepreneurs Report — Private Company Financing Trends. https://www.wsgr.com/

Operating and Value Creation

- BCG. Value Creation in Private Equity (annual). https://www.bcg.com/
- McKinsey & Company. Private Markets Annual Review. https://www.mckinsey.com/
- Bain & Company. PE Operations Report. https://www.bain.com/

Sector and Frontier Investing

- a16z published thesis pieces on AI infrastructure and applications. https://a16z.com/
- Sequoia Capital. Adapting to Endure and successor partnership memos. https://www.sequoiacap.com/
- Bryce Space and Technology / SpaceWorks. Commercial Space Reports.
- Bessemer Venture Partners. State of the Cloud (annual). https://www.bvp.com/

Practitioner Handbooks

- Rosenbaum, Joshua, and Joshua Pearl. Investment Banking: Valuation, LBOs, M&A, and IPOs (Wiley, 3rd ed., 2020).
- Pearl, Joshua, and Brendan O'Sullivan. The Little Book of Private Equity.
- Ramsinghani, Mahendra. The Business of Venture Capital (Wiley, 3rd ed., 2021).
- Carta. Cap Table and Equity Management Guides. https://carta.com/
- Wall Street Prep. LBO Modeling Course Materials. https://www.wallstreetprep.com/