Venture Fundamentals and Fund Economics

Interview preparation summary · VC track

IN THIS SUMMARY

- How a venture fund works
- Stages and strategy

INT.VF.1.a · How a venture fund works

Explain how a venture fund is structured and paid, and work out what fees and carry take from gross returns.

A venture fund is usually a limited partnership. Investors (limited partners, LPs) commit capital that the manager (the general partner, GP) calls down as needed over an investment period, often three to five years, and the fund typically lives about ten years plus extensions. The management fee is often around 2% of commitments a year in the investment period, stepping down later, and it is paid out of commitments, so investable capital is smaller than fund size. Carried interest is typically 20% of profit, sometimes more for sought-after managers. It is paid either on the whole fund once LPs have their contributed capital back, or deal by deal with a clawback; a hurdle rate is less common in venture than in buyout. Because fees and early write-offs come first, net returns usually dip below 1x before rising, the J-curve, and funds are compared within a vintage year. Interviewers ask this to see whether you know that a 3x gross portfolio is a much lower net number. The usual mistake is treating fund size as money invested in companies.

WORKED EXAMPLE: GROSS TO NET ON AN 80M SEED FUND

A venture fund has 80m of commitments. It charges 2.0% of commitments a year for the first five years and 1.5% a year for the next five. Carry is 20% of profit after LPs have received all contributed capital back, with no hurdle, measured on the whole fund. Its companies return total proceeds equal to 3.0x the money invested in them. What multiple do LPs net? Convention: all 80m is called, ignore the GP commitment, tax and timing.

Answer: A 3.0x gross portfolio becomes about 2.2x net to LPs after 14.0m of fees and 23.6m of carry.

> **Interview tip:** Say it in this order: fees come out of commitments first, carry is on profit above contributed capital, so 3.0x gross is about 2.2x net. Then add that the J-curve means the early years will look worse.

INT.VF.1.b · How a venture fund works

Use power-law logic and fund-returner arithmetic to say what outcomes a fund needs and why.

Venture returns follow a power law: a few investments produce most of the value and many return less than cost. A fund returner is one investment whose proceeds equal the whole fund, so the exit value needed is fund size divided by ownership at exit. Ownership at exit is entry ownership multiplied by what survives dilution from later rounds, so a 12% seed stake that keeps 55% is only 6.6% at exit. For a 120m fund that means an exit near 1.8bn, a number you test against the market: is there a credible path to that outcome? The net target works backwards the same way, because 3.0x net after 20% carry needs about 3.5x gross on the fund. Good candidates therefore ask of each deal what the outcome is if it works, not what the average outcome is. The usual mistakes are averaging across the portfolio as if every company returned 3x, ignoring dilution, and praising a safe 2x to 3x outcome that cannot move the fund.

WORKED EXAMPLE: DOES ONE WINNER CARRY A 120M FUND?

A 120m fund has 100m investable and makes 25 equal cheques of 4.0m at 12% entry ownership, which dilution cuts to 55% of that stake by exit. Outcomes: 13 companies return nothing, 7 return their cost once, 4 return 3x cost, and one is sold for 2.5bn. What exit value would return the fund, and how much does the single winner contribute? Convention: no follow-on investing, ignore fees and carry for this step.

Answer: One 2.5bn exit at 6.6% ownership returns 165m, which turns a 0.76x portfolio into 2.4x gross.

> **Interview tip:** Lead with the returner number: fund size over ownership at exit. Then say the portfolio is built to find that one company, so you judge each deal by whether it can plausibly reach it.

INT.VF.1.c · How a venture fund works

Plan reserves, ownership and number of companies for a venture fund and show what pro rata costs.

Portfolio construction links three choices: how many companies, how big each first cheque is, and how much to hold back as reserves for follow-on rounds. Ownership at entry is the cheque divided by the post-money valuation, and each later round dilutes it, often by 15% to 25%. A pro rata right lets the fund buy its current percentage of the new shares, which costs current ownership times the new money raised and holds the stake. Seed-focused funds commonly hold back a third to a half of investable capital for reserves; later-stage funds hold less. More companies give more chances of finding an outlier but smaller cheques and lower ownership; fewer companies give higher ownership but more risk of missing. Reserves should follow evidence, backing the companies that are working, not rescuing the ones that are not. A candidate who can say what a pro rata cheque costs and what it does to the exit stake sounds like an investor. The usual mistake is to forget that follow-on money is bought at higher prices, so it earns a lower multiple than the first cheque.

WORKED EXAMPLE: A 100M FUND WITH 45% RESERVES

A 100m fund has 80m investable after fees and expenses. It reserves 45% for follow-ons and splits the rest into equal first cheques of 2.0m. The seed round values the company at 16m post-money. At the Series A, 10m of new money is raised at 50m post-money. How many companies does the fund back, what is its stake, and what does pro rata cost? Convention: ignore option pool changes and fees.

Answer: The fund backs 22 companies at 12.5%, pro rata at the A costs 1.25m each, and full pro rata would use 27.5m of the 36.0m in reserves.

> **Interview tip:** Give the three numbers: companies, ownership and the cost of defending it. Then say reserves go to the companies that are working, because the money cannot cover every round.

INT.VF.2.a · Stages and strategy

Describe how the questions, evidence, cheque sizes and risk change from pre-seed through Series A to growth.

Stage decides what you are underwriting. At pre-seed and seed the evidence is the team, the insight, the size of the market and perhaps a prototype or early customers, cheques are small, ownership targets are higher, failure is common and a single winner pays for many losses. At Series A you look for product-market fit and a repeatable way to sell, shown in retention, cohorts and early growth. At Series B and growth the question moves to scaling efficiently: unit economics, burn, competitive position and the path to exit. Entry valuations rise, so the multiple a winner can deliver falls while the chance of survival rises. A seed fund can accept that perhaps one in ten reaches a large exit; a growth investor paying a high price needs far better odds. Interviewers ask this to see whether you apply the right tests at the right stage. The usual mistake is demanding growth-stage metrics from a pre-seed company, or paying a growth price for seed-quality evidence.

WORKED EXAMPLE: SAME EXIT, TWO ENTRY POINTS

A company is sold for 600m. A seed investor put in 2.0m at a 12m post-money valuation. A growth investor put in 20m at a 100m post-money valuation at the Series B. Later rounds dilute holders as follows: the seed stake keeps 0.8 x 0.85 x 0.88 = 59.84% of itself by exit, and the Series B stake keeps 88% through a final round. Each investor wants a 3.0x gross multiple on average. What chance of reaching this exit does each need? Convention: failed companies return nothing; ignore fees.

Answer: The same 600m exit pays about 30x at seed and about 5x at Series B, so seed needs about a 10% hit rate and growth about 57%.

> **Interview tip:** Say that earlier stages trade a lower survival rate for a much higher multiple, and name the evidence you would want at each stage.

INT.VF.2.b · Stages and strategy

Compare generalist, thematic and platform strategies and say what evidence shows that a claimed edge is real.

A generalist fund invests across sectors and relies on a wide network and pattern recognition. A thematic or specialist fund concentrates in one area, such as health technology or climate, and claims an edge in sourcing, judging and helping companies, but it carries correlation, because one theme cycle moves the whole portfolio, and a narrower set of buyers at exit. A platform describes services a fund provides to founders, such as recruiting, go-to-market help or introductions, used to win competitive deals and support winners. All three claim an edge; the test is evidence. Look at the share of deals sourced through founder referrals rather than auctions, the win rate in competitive rounds, what founders say when asked for references, and whether follow-on investors come back. Headcount and event numbers are costs, not evidence. Also size the concentration: a 40% fall in the theme's valuations should be mapped to the portfolio. The usual mistake is accepting a label (platform, thematic) as proof of advantage.

WORKED EXAMPLE: TESTING A THEMATIC EDGE

Fund T focuses on industrial software and reports, for one year: 400 companies reviewed, 24 reached a partner meeting, 6 term sheets were offered and 4 were won. A generalist, Fund G, reviewed 1,000 companies, held 30 partner meetings, offered 8 term sheets and won 3. Assume sector valuation multiples then fall 30%. Fund T is fully in that sector and Fund G has 25% of its value there. What do the funnels show, and how do the funds differ in exposure? Convention: marks move one for one with multiples.

Answer: Fund T wins 67% of offers against 38%, but a 30% sector fall cuts its marks by 30% against 7.5% for Fund G.

> **Interview tip:** Name the evidence (win rate in competitive rounds, referral share, founder references) and then name the cost, concentration in one theme.

INT.VF.2.c · Stages and strategy

Calculate and interpret DPI, RVPI, TVPI, gross and net multiples and loss ratio, and say how much to trust a markup.

A fund measures itself with cash and value ratios. DPI is distributions divided by paid-in capital and counts only cash returned. RVPI is the remaining net asset value divided by paid-in capital. TVPI is their sum, the total value per unit paid in. Net figures are after fees and carry and are measured on paid-in capital, which includes fees; a gross multiple is on invested capital before fees and carry and is therefore higher. IRR is sensitive to timing and is unreliable early. Early NAV is a mark, often the price of the latest round, and a markup is not cash: it may come from an insider-led round or a small extension, so ask who set the price. Loss ratio is the share of invested capital, or companies, that returned less than cost; in venture a high loss ratio is normal, so it says little alone. Look at how much of the value is in one company and how much is realised. The usual mistake is quoting a young fund's TVPI as if it were an achieved return.

WORKED EXAMPLE: READING A FUND AT YEAR SIX

A fund has paid in 70m, of which 12m was fees. It has distributed 14m and its remaining holdings are valued at 98m, one company being marked at 52m after a round led by an existing investor. Figures are before any accrued carry. What are DPI, RVPI, TVPI and the gross multiple, and what happens if the 52m mark is halved? Convention: gross multiple is on the 58m invested in companies.

Answer: DPI 0.20x, RVPI 1.40x, TVPI 1.60x and gross 1.93x, but halving one insider-led mark takes TVPI to 1.23x.

> **Interview tip:** Give DPI first because it is cash, then TVPI, then say how concentrated and how well tested the marks are.

Summary of the Charta interview-preparation topic. Companies and figures in the examples are invented. Draft content, not yet expert reviewed.