Master Degree in Engineering and Management

Drawdown or Permanent Capital? A Fund Structure Debate in Venture Capital

Candidate:

Pietro Paggio

******

Supervisors:

Prof. Marco Cantamessa

Politecnico di Torino

July 2026

Abstract

The ten-year drawdown fund structured as a limited partnership has dominated venture capital for over four decades. Yet something is shifting: evergreen funds, semi-liquid vehicles, and permanent-capital vehicles are growing at roughly twice the rate of traditional closed-end funds and now account for about one third of total private capital assets under management. Whether these vehicles address genuine structural shortcomings of the drawdown model or simply serve a different investor segment remains an open question in the academic literature. This thesis examines the inefficiencies of the drawdown model in venture capital and evaluates the extent to which permanent capital alternatives resolve, reframe, or displace them. Although the empirical literature focuses mainly on private equity buyout, the phenomenon has also emerged in venture capital. The analysis follows the three-phase framework of Gompers and Lerner (2001)[1]: fundraising, investing and monitoring, and exit, applying to each phase the same structure: the incentive logic of the drawdown model, the frictions documented in the literature, and the response offered by evergreen funds and investment holdings, both listed and non-listed. The research combines two approaches. Semi-structured interviews with managers of traditional and permanent capital vehicles operating in the Italian market provide practitioner evidence on how fund structure shapes incentives and decision-making. A Monte Carlo simulation extending Brown & Volckmann (2025)[2] models the LP's cash-flow dynamics, portfolio allocation, and terminal wealth across vehicle types, incorporating fee structures, borrowing costs, and liquidity constraints often treated as exogenous or omitted. The analysis shows that the drawdown model's shortcomings are not equally significant. The most relevant are the delayed exposure to illiquid assets, as the commitment-and- call mechanism keeps a large share of LP capital uninvested for years, and the finite fund life, often too short for sectors with long development cycles, particularly deep tech. Fundraising procyclicality matters less: strong track records, established LP relationships, and the ability to

invest at lower valuations during downturns reduce its impact. Permanent capital vehicles address both issues: capital is deployed immediately and the investment horizon is unconstrained. However, they introduce concerns around NAV-based performance measurement, liquidity buffer management (for evergreen funds), and the illiquidity investors bear without a termination date. Evidence from interviews and secondary sources suggests that most limited partners are not yet prepared for open-ended structures apart from particular cases, as they require periodic liquidity and greater visibility into portfolio values between rounds. The simulation further shows that, from an LP perspective, drawdown complexity favours large institutional investors, while the operational simplicity of permanent capital vehicles suits those seeking sustained exposure without active treasury management, raising the broader question of democratisation and wider access to private markets. These two types of structures are therefore best understood not as substitutes but as complementary vehicles, with suitability depending on investor type, sector characteristics, and tolerance for illiquidity.

Contents

List of Figures

List of Tables

Chapter 1

Thesis Outline

This thesis is organised in six chapters.

Chapter 2 introduces venture capital and its economic function, situating it within the broader private equity landscape. It explains the limited partnership model that underpins PE fund structures, traces the historical development of both PE and VC, and surveys recent industry data and trends. Among these trends, the chapter examines the emergence of evergreen and permanent capital vehicles, presenting early industry evidence on their adoption. The chapter closes by stating the research question, explaining the analytical approach, and previewing how the subsequent analysis is structured.

Chapter 3 analyses the fundraising phase of the traditional drawdown model, identifying its structural inefficiencies and discussing how permanent capital vehicles address or reframe them.

Chapter 4 extends the same analysis to the investing and monitoring phases, while

Chapter 5 focuses on exit dynamics.

Chapter 6 draws together the findings and offers concluding reflections.

Chapter 2

Introduction

2.1 The Venture Capital and its economic function

Over the years, the venture capital industry has grown into an important business within the financial world. Venture capital managers, better known as General Partners (GPs), act as agents. The story is simple. Macroeconomic theory tells us that the financial and economic ecosystem is populated by several actors that interact with one another: households, banks, institutions such as insurance companies, institutional investors, pension funds and investment funds, the government, and, we may add, the financial markets themselves (see Fig 2.1).

In the last instance, it is households that finance corporations: their savings are channelled into investment through the intermediation of these various economic agents. [3] A venture capital firm is one such agent. It carries out its intermediation through investment funds that typically raise money from other financial institutions, pension funds above all, especially in the United States, where the market is most active, and then use it to invest in corporations. Not in just any corporation, however, but only in those that are privately held rather than publicly listed. These are usually small businesses looking for capital to grow and scale, the kind of companies we call “startups”. Startups do not generate stable cash flows from the outset. Their distinctive feature is that the product or service they offer is new to the market in which they operate; in this sense, they move hand in hand with innovation and tend to drive forward the progress of society. Bringing the final product to maturity and establishing it as a standard takes time and, above all, money. This means costs

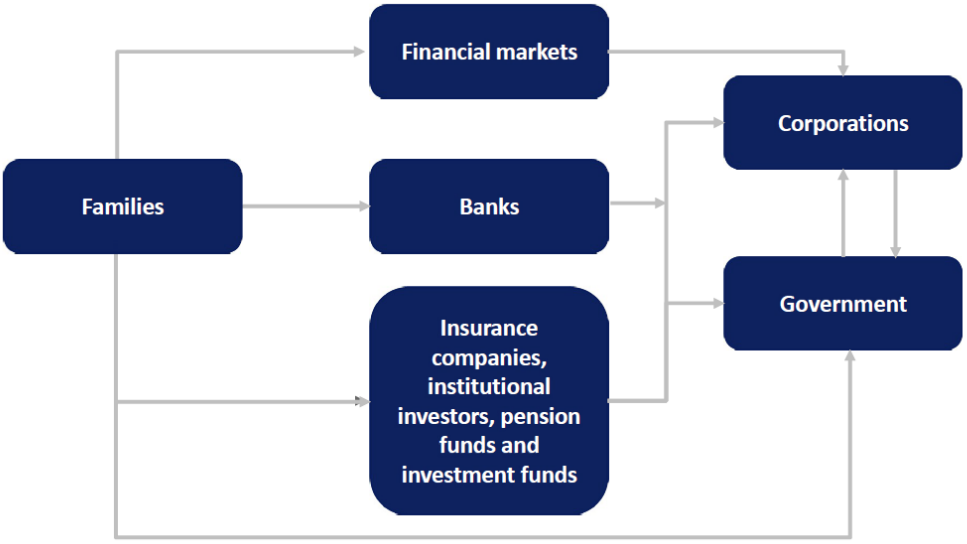

Fig. 2.1 Schematic representation of economic agents and its interaction.

without stable cash flows, and therefore no steady profits, while success is never guaranteed. It is for this reason that we speak of risk capital. Given the nature of these investments, startups do not normally receive direct financing from banks, which require some form of guarantee collateral on assets and covenants. This is all the more true once we consider that, in the so-called digital era, most startups of the past two decades have been engaged in developing SaaS (Software as a Service). Such firms held mainly intangible assets, which were difficult to value; if anything, their intrinsic worth depended on whether the project ultimately succeeded. These assets could not serve as tangible, real assets to be sold off in the event of the company's failure and liquidation.

2.2 Inside the broad context: The Private Equity

To understand the economic function performed by venture capitalists, we should first say something about private equity and set out a few considerations.

The definition of private equity rests on two aspects:[1]

[1]The introduction was made using notes on the Lesson 15 of Applied Corporate Finance & Private Capital course

The first concerns the type of investor: the private equity investor (PEI), a financial institution that may be a bank, a foundation, a social-security fund or a high-net-worth individual. This investor buys equity in a non-listed company with the explicit intention of selling it at a higher price after a predetermined period (typically four to ten years), aiming for a relatively high internal rate of return (IRR) to compensate for the illiquidity and the high risk of the investment. As a reminder, the IRR is the discount rate that sets the net present value to zero, where the net present value is the sum of the present values of all the cash flows generated by the project (including the negative initial outlay for the acquisition); it is the percentage rate of return based on incremental, time-weighted cash flows [4].

The second aspect concerns corporations: for them, private equity is a source of financing, an alternative to others such as a bank loan, a public bond or a public offering (IPO, etc.).

Private equity, in turn, is part of private capital. The latter covers both private equity and private debt (debt extended to private corporations), which may also grant stronger governance rights than those usually attached to public debt. Private debt (sometimes called leveraged finance when the leverage is at least medium, if not high) may come alongside private equity or stand on its own through an "equity kicker" mechanism: this happens, for instance, in turnaround situations, where the value-creation logic is the one typical of private equity but the technical form is mainly private debt.

Private equity encompasses both venture capital, early-stage financing (seed and start-up capital) for young, risky firms, and buyout, replacement, expansion/growth capital and turnaround investment. Different private equity players therefore finance different stages of a company's life cycle (see Fig. 2.2), each with its own characteristics (see Table 2.1).

Kaplan and Stromberg [5] define private equity by focusing on leveraged buyouts (LBOs), drawing a sharp line with venture capital. In a typical buyout transaction, private equity firms operate in mature, low-growth sectors with stable cash flows and acquire a majority stake unlike venture capital, which operates in young sectors alongside emerging entrepreneurs and, as a rule, does not take majority control. Metrick and Yasuda [6] use the term private equity as an umbrella, containing several sub-categories: chiefly venture capital and buyout, but also others such as mezzanine and distressed investing. What these share, and what sets them apart from public

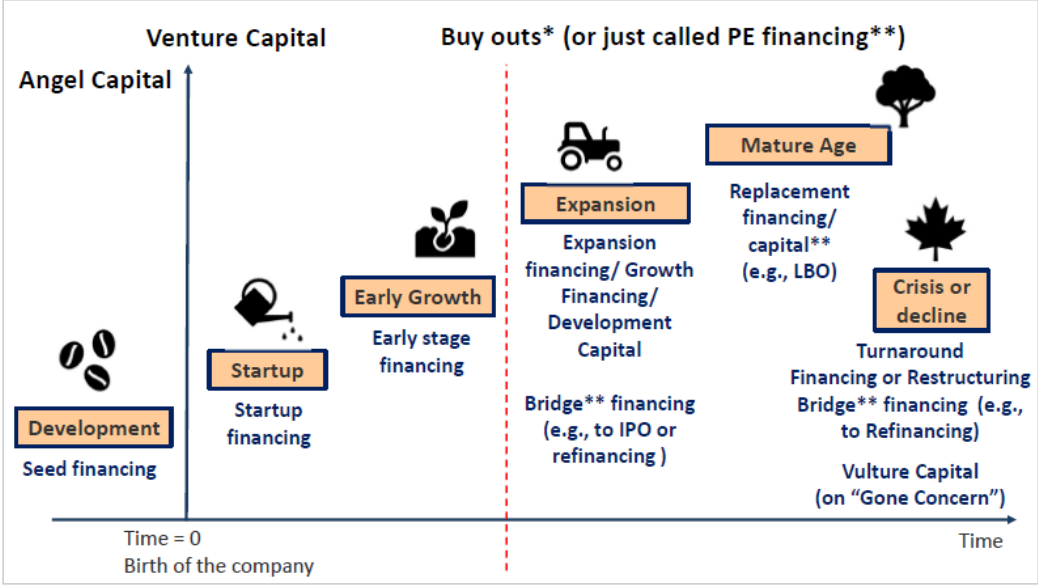

Fig. 2.2 Representation of different stages of a company's life cycle per PE players.

Life cycle phase	Type of investment	Stake acquired	Leverage involved	Risk	Rationale	Goal
Initial	Seed	Minority/ Majority	No	Very high	New business creation	Testing the business idea
	Start up/ Early growth				New company launch	Starting up sales
Development	Expansion	Minority	No	High	New business development	Increasing production capacity
						M&A activity
						Internationalization
Maturity	Buy out	Majority	Yes	Moderate	Generational turnover	Succession plan solutions
					Willingness of shareholders to sell	Company reorganization
Maturity	Replace-ment	Minority	No	Moderate	Shareholder no more interested	Shareholding in line with business development needs
Decline	Turnaround	Majority	No	Very high	Crisis situation	Getting business profitability back

Table 2.1 Table that shows different characteristics per life cycle phase.

markets, is precisely the investment in private, illiquid companies and the strong information asymmetry between insiders and outsiders.

2.2.1 A bit of history[2]

The first private capital entities were 3i (Investing in Industry), founded in 1945 by the Bank of England, and the ARD (American Research and Development) founded in 1946 in the United States by one MIT and one Harvard professor together with a number of private investors.

Venture capital has a long track record across past decades, throughout which it has performed its role of intermediation between flows of capital in various ways and through different models. The different organisational forms emerged as responses to the failures of the models that preceded them: the history of the industry reveals a continuous search for structures capable of aligning the interests of those who provide the capital with those of the professionals who manage it.

Venture capital began just after the end of the Second World War. The American Research and Development (ARD) was also the first venture capital. The investors involved at this early stage were predominantly individuals, since institutions showed little appetite for an unproven investment style regarded as excessively risky. ARD was structured as a closed-end fund listed on the stock exchange a form that allowed it to invest in illiquid assets while still giving investors the chance to trade their shares on the public market. What happened, however, was that brokers sold the fund's shares to unsuitable investors: typically elderly people in need of current income rather than long-term capital gains. When the promised profits failed to materialise quickly, the investors' frustration fell on the managers, forcing Georges Doriot to devote considerable time to handling shareholder discontent.

The turning point came in 1958 with the formation of Draper, Gaither, and Anderson, the first venture capital limited partnership. This organisational form, which soon became the industry standard, made it possible to narrow the investor base to a limited number of institutions and high-net-worth individuals. Another relevant category of investors emerged with the SBIC programme, launched by the United States federal government in 1958 to stimulate technological competitiveness. SBICs were federally licensed pools of risk capital that received government matching funds or loan guarantees against a modest contribution of private capital. Despite their initial proliferation during the 1960s, the SBIC model suffered from poor design and excessive regulation, which discouraged the more experienced venture capitalists.

[2]About the VC evolution through years see Sahlman (1990); for the rise of buyouts see M. C. Jensen (1989).

Weak oversight allowed unscrupulous operators to enter the field, investing in firms controlled by relatives or friends, or in projects with poor prospects. Many SBICs failed between the late 1960s and the 1970s, in a scenario that foreshadowed the savings-and-loan crisis of the 1980s.

Sahlman's (1990) work provides an essential quantitative basis for understanding how the composition of investors changed during the industry's decade of greatest expansion. One of his tables shows that between 1980 and 1988 the independent private sector (made up mainly of limited partners) went from controlling 40% of the total pool to 80%. Over the same period the SBICs' share collapsed from 28.9% to 7%, while corporate programmes fell from 31.1% to 13%. This trend suggests that investors saw the limited partnership structure as the best suited to managing agency problems and to maximising risk-adjusted returns.

The transformation in the type of investor was driven not only by the organisational form of the funds but also by the nature of the parties supplying the capital: the limited partners. The year 1979 marks a watershed in this process, owing to the amendment of the "Prudent Man Rule" by the United States Department of Labor. Before 1979, pension-fund managers had been reluctant to invest in venture capital for fear of breaching their fiduciary duty of prudence. In 1978, individuals represented the largest share of new commitments to the sector, while pension funds supplied only 15%. The clarification of the rule explicitly allowed managers to invest in high-risk assets, provided these formed part of a diversified portfolio. By 1987, only eight years after the regulatory change, pension funds had come to account for more than half of all contributions to the venture capital industry. This massive inflow of institutional capital radically changed the dynamics of the market, allowing the creation of much larger funds and increasing the pressure for professional performance.

Also of interest is the steady presence of corporations as investors in independent funds (around 12–14%), often motivated by the wish to monitor the development of new technologies without having to run an in-house R&D programme directly. On this point, some companies began to operate through their "own" venture capital. This is the phenomenon which likewise spread widely in those years of Corporate Venture Capital, that is, investments made directly by large companies in private start-ups. Although CVC funds have missions similar to those of independent funds, their organisational and incentive structures differ sharply. Most corporate programmes

are set up as subsidiaries and offer far lower incentive-based compensation than independent funds. Companies' motives may be financial or strategic: some seek to stimulate demand for their own products or to buy a "technological window" onto new markets. However, CVC programmes are often subject to economic cycles and to fluctuations in corporate budgets, which makes them less stable partners over the long run than independent funds.

While venture capital was consolidating around the limited partnership, a second branch of private equity rose to prominence in the 1980s: the leveraged buyout. Michael C. Jensen observed the arrival of these new, privately managed companies in the 1980s and noted how they began to substitute for traditional listed companies with fragmented equity, many shareholders and few professional managers. These companies purchased equity with a high level of debt. The heavy debt load entailed substantial interest payments and worked, for managers, as a disciplining device. In the same way, internal compensation was based on the performance achieved by the companies rather than on the mere ownership of a majority stake, as had been the case for the CEOs of the time. These new companies thus started to work by creating real value for investors (returns) and by improving internal operational efficiency, without destroying R&D levels which, in truth, were already low before they were financed . [7]

After the first wave in the 1980s, the sector went through a sharp decline following the collapse of the junk-bond market (high-yield bonds that made it possible to finance acquisitions with extremely high levels of debt) and major scandals such as RJR Nabisco; [5] it then experienced a revival in the 2000s with a second great wave. The market was interrupted again by the great crisis of 2008, before recovering once more [5]. The private capital industry has enjoyed extraordinary growth, rising from roughly USD 2.2 trillion before the 2008 crisis to a figure somewhere between 7 and 10 trillion today.

Some observers now see a clear decline of the public market in favour of the private one. The fall in IPOs is a fact, but it does not necessarily mean that companies are failing. Many young companies prefer to be acquired by large existing public firms. This is partly because modern enterprises are increasingly characterised by intangible capital (ideas, patents, software), which has gone from representing 20% of net assets in 1970 to roughly 90% in the 2000s. [8] For firms rich in intangible assets, going public is risky, because the transparency requirements could reveal

industrial secrets to competitors. Moreover, traditional accounting rules (GAAP) struggle to value these assets correctly, often leading to an undervaluation of listed companies. On top of this, the abundance of private capital makes private equity firms behave increasingly like incubators, bringing companies to maturity before deciding whether to list them or sell them to a strategic buyer. Today's start-ups can raise privately volumes of capital that were historically available only through an IPO, allowing them to stay private far longer and to become so-called "unicorns" before considering a listing.

It should be noted, however, that the two markets are symbiotic rather than competing, as theorised in the "oscillating domains" theory.[9] There is no linear trend towards full privatisation, but rather a boundary between public and private that shifts continuously as a result of investment cycles, technological change and regulatory responses. The private market needs the public market for two fundamental reasons. The first is the exit: despite the fall in numbers today, the IPO remains a crucial way out for liquidating a fund's investments. The second is that private investors use the prices and information of comparable public companies to value their own private assets.

2.2.2 The need of Private Equity[3]

The distinction between the public and the private market is not merely one of size, but a structural one:

- **Asset pricing:** in the public market the price is driven by the daily fluctuations of the market; in private equity it is the outcome of a direct negotiation between the parties.
- **Liquidity:** very high in the public market, extremely low in private equity, where finding a new partner is a long and complex process.
- **Monitoring:** whereas in the public market protection is ensured by stock-exchange regulations, in private equity investors must protect themselves on their own, through formal contractual agreements.

[3]The paragraph was made using notes on the Lesson 15 of Applied Corporate Finance & Private Capital course.

A company turns to private equity not only to raise capital, but to face strategic challenges and structural transformations that the public market or bank credit could not handle with the same operational intensity. Recourse to this kind of capital stems from a clear and well-identified corporate need. Unlike debt, the company raises funds by issuing new shares on the private market and does not have to pay interest to the investor, which lightens the income statement. Typical uses of this capital include:

- **Growth and innovation:** commercial expansion, both domestic and international; investment in new technologies and the introduction of new products.
- **Extraordinary operations:** financing of acquisitions (M&A) and restructuring of the capital structure in order to reduce debt.
- **Change management:** support during generational handovers (successions), changes of governance, and Management Buy-Out or Buy-In operations.
- **Operational support:** financing of working capital through "permanent capital" and professionalisation of management, often turning family businesses into structured organisations.

These objectives are reflected in five key benefits: Certification (prestige), Network (strategic contacts), Know-how (managerial skills), Financial (fresh resources) and People (professionalisation).

Unlike a bank, the private equity investor is an insider who sits on the Board of Directors and contributes actively to the company's strategy. When the stake is a minority one, success depends on a relationship of close trust between entrepreneur and investor. To avoid the freezing of the investment (*investment freeze*), mechanisms such as voting syndicates and PUT and CALL clauses are used, which set out in advance the exit conditions and the relevant prices.

2.2.3 The Limited Partnership[4]

The organisational structure of a private equity fund today follows a well-established model, often described as the industry's "canonical form" and analysed in a number

[4]The paragraph was made using notes on the Lesson 15 of Applied Corporate Finance & Private Capital course.

of fundamental academic contributions. It rests on a contract, the *limited partnership agreement*, signed between the fund's investors (the limited partners, as they are called) and the managers of the management company (the general partners). It is structured as follows (Fig. 2.3):

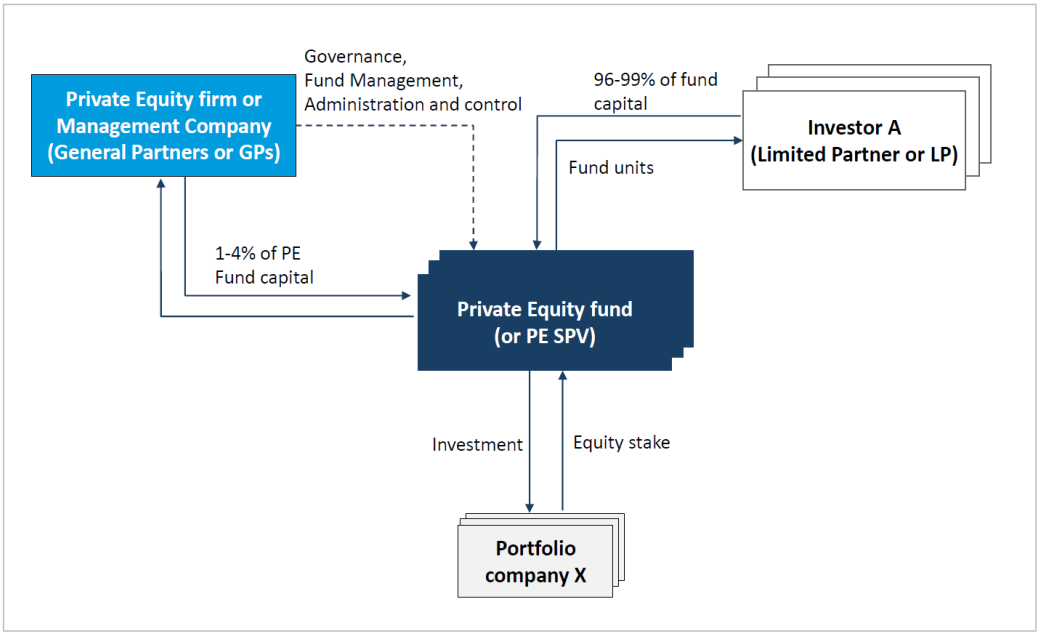

Fig. 2.3 Schematic representation of a limited partnership structure.

- **General Partners (GP):** this is the management company (or AMC, Asset Management Company), made up of the professionals who run the fund operationally. For our purposes, as we will see later, it is worth noting that in Italy this entity usually takes the form of an SGR (*Società di Gestione del Risparmio*), a financial institution supervised and authorised by the Bank of Italy. The GPs are responsible for governance, for the selection of investments and for administration.

- **Limited Partners (LP):** these are the institutional or private investors (banks, pension funds, insurance companies, HNWIs) that provide the capital. They delegate the management entirely to the GPs, giving up the right to "tailor-made" investments in exchange for professional and diversified management.

- **The Fund (or SPV):** this is a separate pool of capital raised from the investors to be deployed in real or financial assets.

While the LPs provide almost all of the capital (96–99%), the GPs are required to invest a share of their own, generally between 1% and 4%, so as to align their incentives with those of the investors. The fund has a predetermined life, usually around ten years (with a typical range of four to ten), and operates on a fixed amount of capital, called *committed capital*. A distinctive feature of the fund as an investment vehicle is that it can never use debt to finance itself: financial leverage is used, if at all, only at the level of the individual portfolio companies (as in LBOs) and not at the level of the fund. The fund's activity is governed by an internal code and an investment policy set out in the Regulation, which in Italy must be approved by the Bank of Italy.

The economics of a fund rests on two main pillars of remuneration for the management company (the GP), designed to cover its operating costs and to align its incentives with the returns of the investors (the LPs). First come the so-called *management fees*. These are paid periodically by the investors to cover the costs of the GP's structure. Typically, and in principle (this qualification will be explored further later), they range between 1.5% and 2%. They are charged on the Assets Under Management (AUM), which correspond to the committed capital during the investment period and, thereafter, to the residual amount invested. Once the exits begin, the amount of assets under management decreases and, as a result, the management fees decline proportionally as well.

There is then the *carried interest* (performance fee). This is a "success fee" and is primarily a tool for aligning the interests of GPs and LPs. It usually amounts to 20% of the realised capital gain. The GP begins to earn the carried interest only if the fund's return exceeds a predetermined minimum threshold, called the *hurdle rate* (e.g. 8% IRR). The distribution of the proceeds from the exits follows a strict order of priority (the *waterfall*) as the following example:

1. **Class A & B – Capital:** repayment of the capital contributed by the LPs and the GP (pro rata, e.g. 98% and 2%).
2. **Class A – Yield:** payment of the preferred return (the hurdle rate) to the LPs until the established threshold is reached.
3. **Class B – Catch-up:** distribution to the GP to "recover" its share of the profit.
4. **Final split:** division of the remaining profits (the upside), typically 80% to the LPs and 20% to the GP.

The fund's performance is measured through the Internal Rate of Return (IRR), with a few clarifications:

- **Gross IRR:** the gross return generated by the fund's investments.
- **Net IRR:** the return actually received by the investors (the LPs), calculated net of the management fees and the carried interest paid to the managers.

2.3 The VC market: recent data and trends

The private market industry, venture capital included, has gone through a complex and difficult period in recent years, and only recently does it appear to be recovering within a fundamentally altered landscape.

The year 2020 brought a severe shock to the global economy with the outbreak of the SARS-CoV-2 pandemic. For many participants the crisis created deep uncertainty and fears of prolonged recession. In such an environment, venture capital managers found it difficult to value startups, conduct due diligence, and forecast revenues. Facing unstable demand, investors largely chose to protect their existing portfolios and reduce exposure to new deals, particularly in the sectors most affected by the pandemic. Yet venture capital did not come to a complete halt despite the unprecedented uncertainty. Gompers, Gornall, Kaplan, and Strebulaev [10] observed that VCs did slow their pace, but the decline was roughly 29%, meaning they invested at about 71% of their normal rate. This aggregate figure masks an important stage-level asymmetry: Howell, Lerner, Nanda and Townsend [11] show that early-stage deal activity fell by nearly 38% in the two months following the onset of the crisis, while later-stage investment remained far more resilient. This was a considerably softer contraction than the collapse of over 50% during the Dotcom bust in 2001 or the roughly 30% drop triggered by the 2008 financial crisis. The reasons behind the slowdown were largely practical: conducting due diligence without in-person meetings became a real constraint, especially for early-stage investors who rely heavily on evaluating the founding team. Still, only about 25% of surveyed VCs reported serious difficulties in assessing new deals or meeting new entrepreneurs. As for the companies already in their portfolios, the results were surprisingly positive. Around 52% of portfolio companies either experienced no negative effects or actually benefited from the pandemic, particularly those in

software, e-commerce, and healthcare. Another 38% suffered some impact but nothing fatal, and only about 10% were severely damaged [10].

In 2021, the market recorded new highs in both fundraising and investment, almost as if compensating for the partial slowdown of the previous year. The rebound was fuelled by abundant liquidity and historically low interest rates, introduced by central banks to pull the economy out of recession. KPMG reported a record year for global venture capital in 2021 [12].

Beyond the immediate recessionary effect, the pandemic triggered a structural reallocation of resources [13]. found that venture capitalists increased the capital invested in pandemic-related sectors, healthcare, biotechnology, remote-work IT, and social-distancing technologies, by as much as 44%, while the number of deals in those specific areas rose by nearly 6%. Although this shift might appear to be driven by panic or hype, the evidence suggests otherwise: the authors show that it was mainly large, well-established vintage funds that moved capital most aggressively. This points to a deliberate reorientation rather than a reactive one, a structural change that laid the groundwork for new investment directions, new sectors, and ultimately new operational challenges for general partners, possibly including new ways of structuring their vehicles altogether.

The recovery was short-lived. In 2022, a global energy crisis set off a period of high inflation and aggressive interest-rate hikes that persisted over the following years. The consequences were predictable: risk appetite declined, investments fell, valuations came under pressure, and exits slowed. Venture capital returned to levels more consistent with a cautious, selective market. Within the broader private equity universe, venture capital fundraising suffered the steepest decline, contracting at a compound rate of roughly 11% per year between 2019 and 2024 — more than double the rate observed for buyout funds [14]

By 2025, dealmaking showed signs of revival. Global VC deal value increased by approximately 27% year-on-year, though the number of deals actually declined by about 10% (Fig. 2.4) [15]. In other words, deal sizes grew while deal count shrank. The operating environment, however, had changed profoundly. High inflation, elevated interest rates, and expensive assets had become the new baseline. Entry multiples reached record levels, with founders demanding prices that many GPs were unwilling to pay, since paying a premium at entry means expecting an exit return

large enough to justify it, which is difficult to count on given the inherent risk of venture investing.

Capital itself was not scarce: dry powder, committed but not yet deployed capital, continued to accumulate, with more than 40% of available dry powder sitting uninvested for over two years, creating mounting pressure to deploy. At the same time, GPs had to address growing LP dissatisfaction. Delayed exits and longer holding periods in preceding years had led to a significant shortfall in distributions back to investors. The result was a market characterised by large commitments flowing into fewer, higher-quality deals. Where financial engineering no longer sufficed, the approach shifted toward operational value creation: improving margins, integrating artificial intelligence, and driving efficiency[15].

In this environment, large funds with established reputations continued to thrive, while smaller and newer entrants increasingly struggled. At the same time, the industry began to experiment with structural models that around 2020–2021 had seemed like a niche curiosity but by the mid-2020s had become a fast-growing trend, driven in part by rising demand from retail investors for access to private markets. The most prominent example is the evergreen or semi-liquid fund. These alternative models, offering greater liquidity, portfolio diversification, and longer capital duration, grew from roughly 27–29% of total private capital in 2020 to an estimated 33–35% by 2025, expanding at a rate of 15–20% annually, well above the 5–10% growth of traditional closed-end funds, which nonetheless remain the dominant vehicle by far [15].

2.4 The rise of Evergreen funds

The growth of private capital has historically been driven by closed-end drawdown funds aimed at large institutional investors. In recent years, however, a new investment model has gained traction: the evergreen or semi-liquid fund. These structures allow continuous entry at net asset value (NAV) and periodic redemptions, removing the fixed maturity typical of traditional vehicles. Unlike drawdown funds, where capital is called gradually over several years, evergreen funds put investor capital to work almost immediately. To manage this, they typically maintain a liquid buffer, usually between 10% and 20% of fund value, in order to meet redemption requests and handle cash flows, while the remainder is allocated to private investments. The

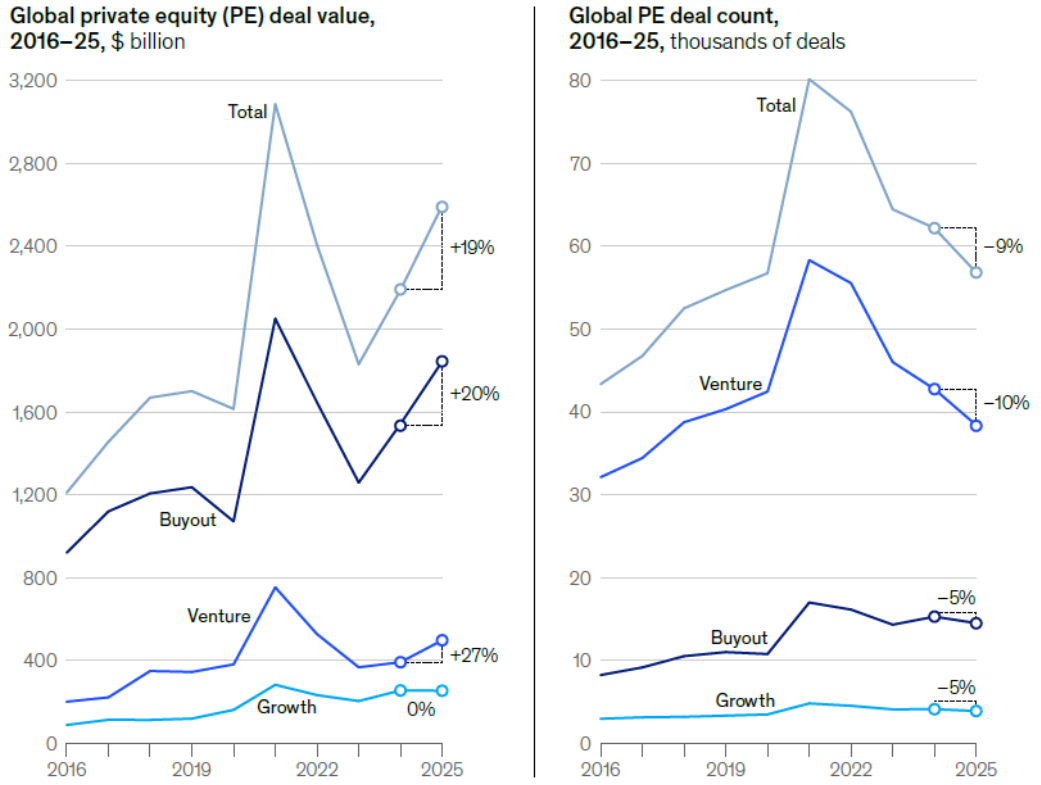

Fig. 2.4 The increase in 2025 deal value was driven by bigger deals, as overall deal value went up but the number of deals went down.

GP actively manages the balance between liquidity and private holdings, automatically reinvesting proceeds from exits, which reduces what is sometimes referred to as "reinvestment risk" for the LP [16].

As for fees, management fees in evergreen funds are generally calculated on the fund's NAV rather than on committed capital, aligning GP compensation with assets actually under management [17]. A qualitative analysis highlights the main trade-offs between these two vehicle types. Evergreen funds offer greater liquidity, albeit gated or limited, and are simpler to manage from the LP's perspective, since the fund manager handles both market exposure and capital reinvestment. Over time, they also tend to generate higher compounded returns for the individual investor, precisely because capital is invested from the outset rather than sitting idle while awaiting capital calls. Drawdown funds, on the other hand, benefit from a much larger existing universe and may offer a gross performance advantage, since they

do not suffer from the cash drag created by the liquidity buffer that evergreen funds must maintain [18].

According to Pitchbook data cited by [2], the expansion of the evergreen segment has been rapid: since 2019, an estimated 200-plus product launches have brought approximately $400 billion in assets under management. Much of this demand has come from individual investors through the wealth channel, representing a potential pool of over $50 trillion[17]. General partners are responding by designing alternative vehicle structures and expanding internal distribution teams to capture this channel. Within this landscape, Myers and Bill (2024)[19] has argued that SEC efforts to shield retail investors from risk through legal barriers are, in some respects, counterproductive. As the public equity universe continues to shrink relative to private markets, restricting retail access may end up excluding a large investor base from a growing share of value creation.

UHNW investors typically access traditional buyout and venture capital funds with high minimum tickets, making them natural partners for closed-end drawdown vehicles that require patient, non-panicked capital. Evergreen funds, by lowering minimum investment thresholds, aim to extend access to smaller investors, a form of market democratisation, but must also find ways to instil the same patience in a more reactive investor base, effectively acting as behavioural correctives. To function well, they need to manage the risk of mass redemptions during crises. This is why they incorporate gates, redemption limits, and early exit fees: mechanisms designed to protect the fund against the pro-cyclical instincts of its investor base. Moreover, by automating reinvestment and simplifying liquidity management, the fund structure effectively locks in a compounding mechanism that generates wealth accumulation without requiring constant active decisions from the investor.

Although the evidence discussed so far pertains primarily to private equity buy-out and private credit, the phenomenon is no longer confined to those segments. Within venture capital, a small but growing number of vehicles with open-ended or semi-liquid structures have begun to emerge, launched both by large fund-of-funds managers and by pure-play VC firms. The cases below illustrate different approaches from the radical restructuring of a legacy model (Sequoia) to new wealth-channel vehicles (Hamilton Lane, StepStone) and European pioneers (Redalpine) all of which share the partial or total abandonment of the ten-year fixed-life constraint typical of the drawdown model.

Sequoia Capital — The Sequoia Fund (2021). In 2021, Sequoia announced the largest restructuring in its fifty-year history, replacing the traditional closed-end fund cycle with a singular, permanent structure called The Sequoia Capital Fund. The new vehicle is an open-ended, liquid portfolio composed of public positions in a selection of the firm's most enduring companies; it allocates capital to a series of closed-end sub-funds for venture investments at every stage, from inception to IPO, with proceeds flowing back in a continuous loop. As the firm stated on its website, the ten-year fund cycle had become an artefact of an earlier era, and the restructuring was designed to remove all artificial time horizons so that investments would no longer have expiration dates.[5]

StepStone Group — SPRING (2022). SPRING (StepStone Private Venture and Growth Fund) is a closed-end evergreen tender fund that provides broadly diversified exposure to venture and growth equity through an open-architecture approach, combining secondary purchases, direct investments, and primary fund commitments. The fund targets global access to experienced venture and growth managers for both individual and institutional investors.[6]

Hamilton Lane — Venture Capital and Growth Evergreen Fund (2025). In May 2025, Hamilton Lane launched the Hamilton Lane Venture Capital and Growth Fund (HLVCG), a continuously offered evergreen vehicle registered under the Investment Company Act of 1940, available to accredited US investors and their advisors. In October 2025 a global version followed (HLGVG), extending availability to parts of Europe, Asia, Latin America and the Middle East. The fund invests in disruptive technologies and innovative companies, leveraging the firm's expertise in co-investments and secondaries. Hamilton Lane's stated rationale is that growth companies are staying private longer, and the fund seeks to address common barriers to entry into this space by offering immediate capital deployment, flexible liquidity options, and curated exposure.[7]

[5] Sequoia Capital, *The Sequoia Capital Fund: Patient Capital for Building Enduring Companies*, October 2021, https://sequoiacap.com/article/the-sequoia-fund-patient-capital-for-building-enduring-companies/.

[6] StepStone Group, *SPRING — StepStone Private Venture and Growth Fund*, fund overview, https://www.stepstonegroup.com/strategies/spring/.

[7] Hamilton Lane, *Venture Capital and Growth Evergreen Fund*, product page, https://www.hamiltonlane.com/en-us/strategies/evergreen/us-private-wealth/venture-growth-fund; press release, May 6, 2025.

Redalpine — Summit Fund (2020). Founded in 2006, Redalpine is a Swiss venture capital firm with nearly two decades of European VC experience. Its Summit Fund, which the firm describes as the first European evergreen venture capital fund designed to offer semi-liquidity, immediate exposure, and compounded returns, has no fixed maturity and reinvests realised proceeds. It provides access to Redalpine's portfolio of high-performing companies from seed to pre-IPO, targeting annual returns above 15%. The fund's stated motivation is structural: it allows the manager to remain patient while founders navigate technology adoption cycles and markets mature, avoiding the forced exits that a ten-year fund imposes precisely when the best companies are still building their most durable value.[8]

We will arguin the implications of these vehicles move on in the analysis.

2.5 The Permanent Capital

The evergreen funds discussed so far are not an isolated recent phenomenon. They represent the latest manifestation of a broader and long-standing concept in private markets: permanent capital. Situating evergreen vehicles within this category is useful because it makes clear that the defining feature is not the semi-liquidity offered to investors, but the absence of a predetermined liquidation horizon for the vehicle itself.

The law firm Troutman Pepper Locke offers a concise and widely referenced definition: permanent capital is an investment for an indefinite period of time in an underlying vehicle. The vehicle may take any form, a corporation, a trust, or a partnership, and the investing entity may be publicly listed or privately held.[9] The most recognised, and arguably most successful, example of permanent capital is Berkshire Hathaway. The same source describes the typical architecture: permanent capital vehicles (PCVs) borrow the structural principles of traditional fund formats, often a limited partnership with a general partner, a special limited partner holding the carry, and a separate management company, but layer on the valuation and redemption mechanics familiar from hedge funds, together with the incentive arrangements more commonly found in a perpetual holding company. A structurally important feature,

[8] Redalpine, *Summit Fund*, fund page, https://www.redalpine.com/summitfund.

[9] Troutman Pepper Locke, *Permanent Capital Vehicles*, practice overview, https://www.troutman.com/insights/permanent-capital-the-essentials/.

from the manager's perspective, is that PCVs have access to capital that is neither subject to shareholder redemption nor to any obligation to distribute proceeds as income is generated, unlike most venture capital and private equity funds, where the contractual lifecycle forces both.

It should be noted that this description reflects US legal practice and partnership conventions. The concept of permanent capital, however, is jurisdiction-neutral. What matters for the purposes of this thesis is the economic architecture, not the regulatory envelope.

2.6 Research intent and methodology

The growing presence of permanent capital vehicles in private markets raises a natural question: why are these structures emerging, and what problem do they solve? Whether these vehicles represent a genuine structural improvement over the traditional drawdown fund or simply a complementary offering aimed at different investor segments is the central question this thesis sets out to investigate. The analytical starting point is the traditional ten-year drawdown fund organised as a limited partnership. Before evaluating alternatives, one must understand why this model became dominant in the first place, what incentives and institutional logic sustain it, and then identify where it generates inefficiencies or misalignments, particularly from the limited partner's perspective. Only against that baseline does it become meaningful to ask whether permanent capital vehicles resolve those problems, merely shift them, or introduce new ones.

To structure this analysis, the thesis adopts the venture capital cycle framework introduced by Gompers and Lerner [1], who decompose the life of a VC fund into three phases: fundraising, investing (and monitoring), and exit. Each phase is analysed along the same three dimensions: first, the structural logic and incentive architecture of the drawdown model; second, the inefficiencies or tensions that the academic and practitioner literature has identified within that phase; and third, the response offered by permanent capital alternatives, whether they mitigate, reframe, or merely displace the underlying problem. Conclusions are then drawn by bringing together the findings across all three phases.

2.6.1 Scope: vehicles considered

Within permanent capital, this thesis focuses on two broad categories that are most directly comparable to VC drawdown funds. The first is the **evergreen or semi-liquid fund**. The second category comprises **investment holdings**, both listed and non-listed, that is, corporate entities holding a portfolio of venture or technology investments on a permanent basis, financed through their own balance sheet rather than through LP commitments.

Qualitative methodology: interviews. The qualitative component of this thesis draws on semi-structured interviews with practitioners across the VC ecosystem. The interview framework is organised by theme and aligned with the three phases of the venture capital cycle, so that each conversation yields evidence that maps directly onto the analytical structure of the thesis. The intent is to capture multiple perspectives those of general partners, on how structural features of different fund models affect incentives, decision-making, and outcomes.

On the traditional drawdown side, the thesis includes interviews with representatives of United Ventures and Panakes Partners, two Italian VC funds operating under the conventional limited partnership model. On the permanent capital side, the primary case is LIFTT S.p.A., a non-listed investment holding focused on deep-tech and technology transfer in Italy. Reaching evergreen VC funds proved considerably more difficult: the market is still small, access is limited, and the most relevant European case the Summit Fund managed by RedAlpine could not be reached for an interview. A particularly instructive case is that of AENU, a Berlin-based fund that launched as an evergreen vehicle before converting back to a traditional drawdown structure, reportedly due to insufficient institutional LP appetite for the innovative format. Direct contact with the fund was not possible; the analysis of this case therefore relies on published interviews and press coverage.

Quantitative methodology: simulation. The analytical chapters are supplemented by a quantitative simulation that models the LP's experience across different fund structures. Building on the framework developed by Brown and Volckmann (2025), who compare the cash-flow dynamics, portfolio risk, and expected returns of an evergreen fund against a staggered series of drawdown funds, the simulation extends the analysis to include investment holdings and introduces additional layers—fee

structures, borrowing costs and liquidity stress—that the original model treats as constant or omits.[10] The simulation is calibrated with parameters drawn from the academic literature and, where possible, informed by the interview evidence. Its purpose is not to rank specific vehicles but to isolate the structural mechanisms through which fund design affects LP outcomes—an exercise that complements and quantifies the qualitative discussion. The methodology and results of the simulation are presented in detail in Chapter 3, where they directly inform the analysis of fundraising-phase inefficiencies.

[10] G. Brown and W. Volckmann, "Evergreen vs. Drawdown Funds: Risk, Returns and Cash Flows," Institute for Private Capital, June 2025. The base model and its equations are described in §2 and Appendix A of the paper.

Chapter 3

Fundraising

3.1 Introduction

This chapter turns to the first phase of the venture capital cycle outlined in the introduction. Building on the architectural overview of the different vehicles—from traditional limited partnerships structured as drawdown funds, to the more recent evergreen and semi-liquid models, to permanent-capital investment holdings—the analysis that follows focuses exclusively on the fundraising phase. To preserve analytical rigour, the chapter confines itself to issues that arise during capital formation. Compensation mechanisms such as management fees and carried interest, the time horizon of investments, the technical-expertise gap between managers and investors, the mechanics of capital deployment, and exit-related dynamics are all deferred to the next chapter. This separation is deliberate: it isolates the frictions that general partners face specifically when seeking new resources, before those resources are put to work. The chapter proceeds in four steps. A first section explains the functional and historical reasons why the drawdown model became dominant, drawing on both classic and recent academic literature. A second section examines the tensions inherent in the traditional fundraising process. A third section then asks how permanent-capital alternatives position themselves in relation to those tensions, and what further challenges they may introduce. Throughout these three sections, the qualitative analysis is supported by evidence drawn from semi-structured interviews with investors and operators active in the industry. A fourth section presents the results of a Monte Carlo simulation designed to evaluate the returns and risk profile

of several investment strategies from the standpoint of a limited partner choosing among them. The central argument of the chapter is that the way capital is raised and structured is not a mere administrative detail but a foundational choice whose effects propagate directly to the outcome experienced by the limited partner—shaping returns, risk exposure, cash-flow predictability, and effective access to the best opportunities the market has to offer.

3.2 The dominant structure: motivations and incentives

Capital is not raised and deployed all at once but progressively, in tranches, over the investment period. The first problem that the fundraising of a venture capital fund must solve is, in fact, one of timing. Estimates from a broad sample of funds show that a VC fund with a five-year investment period deploys its capital in a staggered fashion over the early years of its life, rather than in a single initial outlay [6]. Asking the investor to transfer the full amount at subscription would immobilise capital inside the fund while it waits to be put to work, imposing an opportunity cost that erodes fund returns before the first investment is even made—capital that the LP could otherwise allocate at the risk-free rate or in public markets. The main problem here is that fund's return would suffer from the resulting idle liquidity, a phenomenon known as cash drag on the IRR.

The solution adopted by the drawdown model decouples the commitment from the actual transfer of funds. The investor subscribes a binding commitment but contributes the resources in successive tranches: a portion at closing and the remainder called at later dates, with the general partner retaining substantial control over the timing of each call [20]. The commitment is made enforceable by severe penalties for default, which can extend to the forfeiture of up to half of the limited partner's capital account [20]. Capital is thus mobilised as it is needed, while the uncalled portion remains at the investor's disposal. It is in this sense that the drawdown, on the fundraising side, represents an efficient solution: the way capital is formed and made available minimises idle liquidity and aligns the investor's outflows with the fund's deployment schedule. This efficiency, however, is not costless. Because the general partner controls the timing of calls, the burden of standing ready to honour requests

that are unpredictable in both timing and size falls on the investor. Moreover, the management fee is traditionally calculated on committed capital—not only on the amount actually invested—at least during the investment period [20] [6] [21]. The investor therefore begins to bear costs on resources that are not yet at work. It is precisely these tensions that offset the efficiency just described and that anticipate several of the issues discussed in the subsequent analysis of the model's frictions.

For a fund to deploy capital in illiquid assets intended to be held for years, it must have the assurance that the capital will not be withdrawn in the meantime. The limited partnership obtains this assurance through a set of structural conditions: a dissolution date agreed in advance, restrictions on the transfer of interests, and, above all, a prohibition on withdrawal before term [20]. In Sahlman's sample, 72% of funds have an economic life fixed at ten years—the standard form of the closed-end, finite-life vehicle, common to both venture capital and buyout [22]. Sahlman frames these conditions primarily as requirements for the partnership's tax treatment, but an economic reading points to a different implication: the resulting lock-up of investor capital is what makes the fund viable as a fundraising proposition. These features are not defects; they are the conditions that render the vehicle financeable. The prohibition on withdrawal guarantees the manager a stable capital base on which to build a multi-year, illiquid strategy, neutralising the risk of unpredictable redemptions that would otherwise force asset sales at the worst possible times. This is precisely the structural difference with respect to open-ended vehicles, which must hold liquidity reserves and may have to suspend or ration redemptions under stress.

The flip side of finite life is that the management company raises capital in distinct, successive funds: each vehicle is born, invests, liquidates, and closes, and the manager is compelled to return to the market periodically to raise the next fund before the current one is exhausted [23]. This gives rise to what can be described as a *fundraising cycle*: capital formation is not a one-off event but a recurring appearance before investors.

The lock-up of capital in illiquid assets and the mechanics of the fundraising cycle restrict the pool of investors able to subscribe such commitments. Access to venture capital funds has historically been closed to the vast majority of individual investors, and not only for the economic reasons just discussed but also for regulatory ones. On the regulatory side, funds are offered exclusively to investors who meet sophistication tests based on net-worth and income thresholds well above those of the average

household. Managers accordingly set minimum investments on the order of millions or tens of millions of dollars. To ease the burden of repeatedly sourcing new capital every few years—as the fundraising cycle demands—managers typically turn to the same investors from prior vintages rather than admitting new ones: demand from existing LPs with participation rights in successor funds is often sufficient to fill the vehicle [24]. This practice also mitigates reinvestment risk for LPs, who can re-commit to managers whose approach they already know. These investors are, almost by definition, institutional: pension funds, university endowments, sovereign wealth funds, insurers, and funds of funds, all of which possess time horizons that match the fund's ten-year life, the resources to manage uncalled commitments and absorb cash-flow variability, and portfolios broad enough to diversify across vintages and managers, thereby dampening the idiosyncratic risk of any single investment.

Yet the institutional base is far from homogeneous, and the internal dispersion is as striking as the gap between institutions and retail. The private equity market is characterised by particularly severe information asymmetries, which amplify the differences among institutions in terms of evaluation capacity. Lerner, Josh and Schoar, Antoinette and Wongsunwai, Wan [25] document large performance differences between endowments, which consistently invest in top-performing funds and leverage private information acquired as insiders when making reinvestment decisions, and the average institutional LP; banks, by contrast, rank persistently at the bottom (Table III). This finding matters because it shows that regulation is not the only barrier to entry: the ability to extract value from the asset class is itself unevenly distributed, even among investors who nominally satisfy the fund's requirements. In practice, it is an investment for the few. The exclusion of retail is not a defect but a consequence of a structure that demands sophistication to be managed effectively—a lesson illustrated as early as the very first venture capital vehicle, ARD, whose public shareholders struggled with the partnership's illiquidity and valuation opacity.[1] What remains is the fact that a very large share of the potential capital market is excluded—precisely the space that permanent-capital alternatives seek to tap.

Once established, the model consolidated itself through mechanisms that make it difficult to displace. The first is standardisation. Limited partnership agreements

[1] American Research and Development Corporation, founded in 1946, is widely regarded as the first institutional venture capital fund. Its experience with retail shareholders is discussed in Lerner (2010).

are structured in remarkably similar fashion from one fund to the next: closed-end, finite life, management fees on committed capital, carried interest typically set at 20%, and GP activity constrained by negative covenants [20] [22]. Compensation is defined at inception, rarely renegotiated, and tied to observable, verifiable returns [26]. This homogeneity lowers contracting and evaluation costs and produces a broad, comparable universe of funds across which LPs can allocate and diversify.

The second mechanism is the track record. Successive vintages generate comparable performance histories, and past performance is a strong predictor of the ability to raise a follow-on fund and of its size; moreover, outperformance tends to persist across consecutive funds managed by the same partnership [27]. The cycle is thus self-reinforcing: good results feed the next raise.

The third is a form of implicit reinvestment contract. The explicit contract requires liquidation at term, but it is complemented by a tacit understanding that capital providers will reinvest in future funds sponsored by successful managers. It is precisely this expectation of reinvestment, hand in hand with a proven track record, that makes it rational for the manager to build infrastructure and expertise that outlive any single fund [28]. At the industry level, the equilibrium reinforces itself even through market cycles: entry and performance are procyclical, yet established managers prove less sensitive to downturns than new entrants [27].

3.3 The drawdown frictions

To sum up, the drawdown structure originates from an efficient, self-reinforcing equilibrium: the commitment-plus- capital-calls mechanism mobilises capital only when it is needed, minimising idle liquidity [20] and, with it, the problem of cash drag—that is, uninvested liquidity depressing the illiquid return. Whether cash drag is a problem at all depends on whether the liquidity is a structural feature of the vehicle or an unintended residual: in an evergreen fund, for instance, a cash buffer exists to service redemption requests and is therefore integral to the design. The preceding sections have also shown how lock-up and finite life make the vehicle financeable by guaranteeing the manager a stable capital base for a multi-year, illiquid strategy [20], and how the standardisation of terms, the portability of the track record, and the implicit expectation of reinvestment in successor funds have consolidated the model's dominance [22] [27] [27].

Since alternatives to this model do exist, however, it is worth asking whether these same features, functional as they are, also entail a cost during the phase in which capital is raised and made available. The right question is not whether the model has defects, but whether the very choices that made it dominant generate genuine inefficiencies for the investor, or rather expected and partly compensated costs that only appear to be inefficiencies. The literature offers concrete reasons to investigate, but does not deliver a uniform verdict—and it is precisely this heterogeneity that justifies a careful, case-by-case analysis.

Five candidate frictions emerge: fundraising volatility, cash drag, the J-curve, distortion of performance metrics, and closed access to retail investors. They do not, however, carry equal weight. For performance metrics the inefficiency is clear-cut and well documented; for cash drag borne by the LP it is equally documented and verifiable in principle through simulation; for fundraising volatility and retail access, by contrast, it is debatable whether they constitute genuine inefficiencies at all. The J-curve sits in between: it is more a measurement artefact than a structural defect. The pages that follow address each point in turn, drawing on the academic literature, the interviews, and a quantitative simulation.

3.3.1 Fundraising volatility

The dependence of fundraising on market conditions is well documented: venture capital activity is volatile and tracks public markets [29], new fund entry and performance are procyclical [27], and capital inflows during periods of strong demand tend to inflate valuations [30].[2]

This does not, however, automatically qualify as an inefficiency—or at least not entirely. While younger managers do exhibit procyclical behaviour, experienced managers increase their investment activity in response to favourable signals without this penalising their success rates [29]; and much of the underperformance of funds raised in hot markets attenuates once compared with public equities over the same period—suggesting that the LP did not necessarily make a worse choice, because the alternative, at that point, also performed poorly [31]. By the same logic, being bound by the fundraising cycle and having to call capital during adverse market windows is

[2]The finite-life, closed-end nature of the fund, which triggers a continuous cycle of discrete fundraises, makes venture capital inherently more sensitive to market conditions, which in turn shape both performance expectations and timing.

not inherently harmful: the evidence shows that calling capital in bad times yields higher returns, both in absolute and relative terms, because a liquidity premium compensates the investor [31]. In venture capital this link is even stronger: cash flows and performance are considerably more cyclical than in buyout, and the association between countercyclical capital calls and superior returns is correspondingly larger [31].

Practitioner experience corroborates the pattern while adding an important nuance. Diana Saraceni, co-founder of Panakès Partners, confirms that market timing weighs heavily on fundraising: the first close of the firm's second fund, the Purple Fund, took place in 2021, when the pandemic had created a tailwind for life sciences investment. Yet the fund reached its hard cap of EUR 175 million only in February 2023, in a markedly less favourable environment—evidence that a strong track record from the previous fund and a loyal anchor-investor base can partially offset adverse conditions.[3] Saraceni also observed that even sophisticated institutional LPs are reluctant to commit during periods of market stress, despite the potentially superior vintage-year returns that such periods offer—a behavioural pattern consistent with the procyclicality documented in the academic literature and with the finding in Lerner, Schoar, and Wongsunwai [25] that the ability to time commitments is itself unevenly distributed across LP classes. Massimiliano Magrini, co-founder and Managing Partner of United Ventures, adds a converging observation from a different empirical vantage point. Asked how much market timing had weighed on the fundraising of United Ventures' third fund—closed, per the interviewer's framing, in 2023 amid high interest rates and a difficult market—Magrini answered in general terms: capital, he explained, is typically raised from investors who had already committed to the manager's previous fund.[4] This is, in his account, the ordinary mechanics of the drawdown model: the same closed, repeat-player circle of institutional investors as the channel through which track record and the implicit reinvestment contract sustain fund-to-fund continuity.

3.3.2 Cash drag

The drawdown mechanism was designed, as noted, precisely to avoid cash drag on the fund's return: capital is mobilised only when needed, so as not to leave idle

[3] Interview with Diana Saraceni, Co-founder and Managing Partner of Panakès Partners

[4] Interview with Massimiliano Magrini, Co-founder and Managing Partner of United Ventures.

liquidity depressing the illiquid return [20]. Why, then, list it among the candidate frictions? The reason is that the same label covers three distinct phenomena, only one of which directly affects the fund's IRR.

The first is the cash held inside the fund: capital already called but not yet invested, which sits in liquid form during the interval between the call and the closing of a deal. This is cash drag in the strict sense—the drag that lowers the IRR—and it is also the one that the commitment-and-call mechanism minimises by timing calls close to deal closings. Practitioner experience confirms this assessment: Saraceni notes that at Panakès Partners capital calls occur approximately every four years in concentrated tranches, and characterises the resulting cash buffer inside the fund as minimal and not materially harmful to overall illiquid returns.[5] The second is the effect of management fees calculated on committed rather than invested capital: in the early years the investor pays on resources that are not yet at work, eroding net returns [20] [6]. This component anticipates the J-curve phenomenon discussed in the following subsection.

The third channel is arguably the most consequential, because it constitutes a genuine inefficiency in a way that the first two do not: the first is minimised by design, and the second is a conscious, structurally functional cost justified by expected returns. Limited partners see their committed but as-yet-uncalled capital earning less than it would if it were fully deployed in illiquid assets. It is held largely in low-yielding liquid instruments so as to be available when attractive deals arise, rather than sitting idle in the portfolio. And because the timing and size of capital calls can be unpredictable, LPs are all the more compelled to keep a substantial share of their commitment in liquid, low-return form, depressing their total return on the allocation. Critically, this cost does not enter the fund's IRR—uncalled capital is not a cash flow of the fund—and therefore remains an inefficiency clearly experienced by the LP but invisible to the metric most commonly used to assess performance, a point to which we return when discussing performance metrics below. Brown and Volckmann illustrate the magnitude for private equity buyout: in their simulation a single drawdown fund has on average only about 28.7% of committed capital invested at any given time, with the remainder earning the risk-free rate (Brown & Volckmann, 2025, Table 1, p. 7). Although their calibration targets buyout rather than venture capital, the mechanism is the same and, if anything, is likely to be

[5] Interview with Diana Saraceni, Co-founder and Managing Partner of Panakès Partners

more pronounced in VC, where deployment is slower and distributions arrive later. The aggregate signal of this committed-but-idle capital is dry powder, which by early 2020 had reached a record level of approximately USD 276 billion in venture capital alone [11]. This inefficiency will be analysed in detail through the simulation presented later in the chapter.

3.3.3 The J-curve

The second type of cash drag leads directly to the next theme, because the management fee calculated on committed capital is also one of the three drivers of the phenomenon known as the J-curve, together with the fund's formation costs and the investments made in portfolio companies at a stage when exits have not yet materialised [20]. In the early years the fund exhibits negative cash flows: setup costs and fees are paid immediately, capital is deployed, yet returns remain deferred because exits during the fund's first years are very difficult—not impossible, but rare. Reported performance therefore starts below zero and recovers only later, tracing the characteristic J shape; accordingly, standard performance measures such as IRR and multiples are typically depressed at the beginning of a fund's life and improve gradually [23].

This may look like an inefficiency, but it is better understood as a structural property of the vehicle, managed knowingly by fund managers and amply justified by the returns that follow. Saraceni confirms that in life sciences the J-curve is structurally deep, given the length of R&D and clinical trial timelines, but notes that at Panakès Partners it was "well managed" through what she describes as an excellent ongoing dialogue with investors on expectations.[6] It is, at bottom, a measurement phenomenon—all the more so because the interim NAVs through which performance is communicated to investors during this phase prove unreliable: they tend to overstate the value of the more mature, less active funds in particular [32]. The J-curve, in other words, matters less in itself than for what it signals: the fragility of the metrics used to gauge performance in the fund's early years—a point taken up directly in the next subsection.

[6] Interview with Diana Saraceni, Co-founder and Managing Partner of Panakès Partners

3.3.4 Performance metrics

The distortion of performance metrics is a genuine, well documented inefficiency, and one that is more acute in venture capital than in the rest of private equity, because VC carries more unrealised value, positions that are harder to appraise, and returns that are more heavily skewed. The starting point is that the metric the industry uses to communicate returns—the since-inception IRR—does not represent the rate at which the investor's wealth actually grows. It is driven, and effectively locked in, by the earliest large cash flows, and is better read as an internal discount rate that renders the stream of cash flows internally consistent than as a measure of wealth accumulation [33]. The IRR is simply the rate that equates inflows and outflows, but it implicitly assumes that every euro distributed is reinvested at the same IRR until the end of the fund's life—an assumption that no investor can fulfil in practice.

Its manipulability follows directly. Because early cash flows dominate the calculation and because the IRR is a time-sensitive measure, it rewards managers who exit winners quickly, who use fund-level credit facilities to delay capital calls,—though the prevalence of this practice may vary by segment: Saraceni notes that in mid-market European VC obtaining fund-level credit facilities is far from straightforward, as anchor investors impose conditions that constrain their use.[7]and who enjoy strong early performance [33]. In venture capital, where a single early home run can set the IRR of an entire fund, the effect is amplified. The IRR also ignores committed but uncalled capital—since it is not a cash flow of the fund—and therefore fails to capture the under-investment discussed under cash drag, returning a more favourable picture than the LP actually experiences.

On the valuation side, interim NAVs are subjective and tend to be biased upward, particularly for more mature, less active funds [32]. When rigorous corrections are applied—rewriting unrealised NAVs and adopting more conservative weightings—reported performance falls markedly: in the sample of Phalippou, Ludovic and Gottschalg, Oliver[34] the average Profitability Index drops from 1.01 to 0.88, with an average alpha of roughly -3% per year, and the underperformance is present in both buyout and venture capital but more pronounced in the latter.[8]

[7] Interview with Diana Saraceni, Co-founder and Managing Partner of Panakès Partners.

[8] Phalippou and Gottschalg's sample covers 852 funds (616 VC, 236 buyout) with vintage years 1980–1993. The corrected PI of 0.88 corresponds to an annual underperformance of approximately 3% relative to the S&P 500 (Table 2, Panel B, p. 39 of SSRN version). Verify against the published version in *Review of Financial Studies*, 22(4), 2009, pp. 1747–1776.

That this distortion is specific to the fundraising phase is no accident: NAVs are managed upward precisely around the time of fundraising, especially by managers with little reputational capital and few realised exits to show, only to be marked down more frequently after the new fund closes [23]. The metric, in short, is bent strategically at the moment it is needed most to convince.

The most telling consequence is that, for venture capital, the basic question of whether net returns exceed public markets remains unresolved—and the reason for this irresolution is itself instructive. On one side, estimates that write down unrealised NAVs and impose a high beta on VC find strong and statistically significant net underperformance [32, 34]; on the other, estimates that trust reported NAVs and assume a unit beta find VC above the market [27], with a capital-weighted PME of 1.21. The disagreement arises not from the raw data but from how valuations and risk are treated—it is, in other words, a direct manifestation of the very metric unreliability the section describes. The literature has accordingly developed alternative measures such as the PME, which, however, retains a limitation on the point that matters most for VC: if the fund's beta exceeds one, the PME overstates risk-adjusted performance [27]. The inflated perception of performance, built on numbers that look like returns but are not, has plausibly contributed to the rising flow of capital into these funds and to their penetration of retail portfolios [33].[9]

3.3.5 Closed access to retail investors

The drawdown model is, by construction, closed to the vast majority of individual investors. Accredited- or qualified-investor requirements, minimum commitments on the order of millions of dollars, and the participation rights enjoyed by existing LPs in successor funds drastically narrow the pool of potential subscribers [24]. As discussed above, this restriction is not accidental but a direct consequence of the vehicle's design: the product demands long horizons, the capacity to absorb cash-flow variability, and the scale to diversify across vintages and managers.

The literature paints an ambivalent picture. On one hand, excluding retail reduces the pool of capital available to venture and limits the diversification options of the individual investor—a real cost if that investor could benefit from exposure to illiquid

[9]The problem is compounded in evergreen vehicles, whose performance is by construction anchored to NAV rather than to realised cash flows—a point taken up in the section on permanent-capital alternatives.

returns. On the other, the evidence suggests that the ability to extract value from the asset class is unevenly distributed even among institutional investors themselves: the performance dispersion across LP classes is large, with university endowments consistently investing in top-performing funds and banks ranking persistently at the bottom [25]. If within the institutional base itself—that is, among investors who by definition satisfy every access requirement—the disparity in skill is this pronounced, opening the door to retail investors does not guarantee better outcomes for the new entrants.

There is, moreover, a structural argument. The illiquidity and lock-up inherent in the drawdown are not merely a cost borne by the investor; they can function as a screening device, attracting LPs with long horizons and a low probability of liquidity shocks while deterring those who would need to exit early [35]. Under this reading, the lock-up is a form of investor selection that lowers the manager's cost of capital in future raises and protects the remaining investors from forced liquidations. The closure to retail is therefore not simply a barrier: it is part of the mechanism that sustains the equilibrium. Saraceni's observation that even sophisticated institutional LPs are reluctant to commit during downturns and resist longer lock-ups—despite the potentially superior vintage-year returns—reinforces this point: if institutions themselves strain against the model's demands, extending access to less experienced investors is unlikely to be frictionless.[10]

The fact remains, however, that the growth of permanent-capital and semi-liquid vehicles means the market is now experimenting with precisely the democratisation that the traditional model forecloses. Whether this experiment works—and for whom—is an open question. The most recent evidence on interval funds shows that retail investors benefit from access to illiquid, information-insensitive assets such as private credit, but underperform in information-sensitive strategies such as equity and private equity [36]; and their returns improve only when they co-invest alongside sophisticated investors, whose presence sustains the manager's incentive structure.

The question that arises at this point is whether permanent-capital vehicles—evergreen and semi-liquid funds and investment holdings—represent a structural response to these tensions, and at what price. The analysis that follows examines how each alternative positions itself relative to the frictions identified in the drawdown model, asking whether it resolves, transforms, or replaces them with new ones. Academic

[10] Interview with Diana Saraceni, Co-founder and Managing Partner of Panakès Partners.

sources on this topic are scarcer: the peer-reviewed literature on evergreen funds in PE and especially in VC is still limited, the vehicles are few and often recent, and empirical evidence on returns and performance is thin. A significant share of the available material comes from industry sources or recent working papers—and it is precisely this gap that the interviews and the simulation are designed to fill.

3.4 How permanent-capital vehicles raise capital

In the drawdown model the manager returns to the market every three to five years, closes the fund, and from that point the capital is locked. In permanent-capital vehicles the process works differently—as the introduction already anticipated for some of them—and the consequences for the fundraising phase are equally distinct. The change is not merely operational; it redefines the relationship between the investor's capital and the structure of the vehicle.

3.4.1 Evergreen funds

In an evergreen fund the investor subscribes at net asset value and the capital is deployed almost immediately. There are no binding commitments and no capital calls: the capital is at work from the day it enters. The management fee is calculated on NAV—that is, on capital actually deployed—rather than on committed capital [2, 37], which stabilises the ratio between fees paid and capital at work. The immediate deployment eliminates the initial period in which fees run on capital not yet invested, the principal driver of the J-curve in its classical form. The J-curve is attenuated but does not disappear: think of a fund invests the bulk of its capital from day one in early-stage companies that will not generate visible value immediately, and valuations will remain conservative for a time. NAV may even decline in the first quarters if accounting policies recognise losses before gains. Overall exposure is nonetheless high from the outset, and the fee component is rendered more proportionate—a point whose implications for manager incentives will be examined in the chapter on investing and monitoring.

Capital is raised through continuous subscriptions, which should in principle eliminate vintage-year risk: the LP need not time the fundraising window, and Redalpine presents this as "cycle-agnostic access" that shields investors from vintage disper-

sion.[11] But the dependence on market conditions does not vanish; it transforms. If fundraising slows because the market turns cold or the manager fails to convince, the fund loses the ability to refinance its liquidity sleeve and to make follow-on investments in existing portfolio companies. Fundraising thus becomes a permanent operational requirement rather than a discrete event—and if capital raised outpaces deal flow, the manager faces an incentive to invest in mediocre opportunities so as not to accumulate cash that depresses reported returns.[12] Managers who come from the closed-end world may, moreover, find themselves ill equipped to run this continuous machine: as the authors of "Private Equity for All" observe, GPs from the world of private, closed-end funds "are inexperienced with managing liquidity in their funds" (38).

Immediate deployment also eliminates the cash drag from uncommitted capital that weighs on the LP's portfolio in the traditional model. Redalpine, presenting its Summit Fund—the first evergreen venture capital fund in Europe, launched in 2020 with a target of one billion Swiss francs—summarises four advantages from the investor's perspective: liquidity and control (scheduled redemptions without forgoing exposure), reduced operational burden (no capital calls to manage), accessibility (lower minimum tickets than in drawdown funds), and return dynamics (continuous compounding and J-curve mitigation). UBS (2023), cited by Redalpine itself, concludes that no structure is inherently superior: perpetual funds are a compelling entry point for new investors and a complementary sleeve for mature programmes—not a substitute.[13]

This narrative is consistent with the academic literature. Gourier, Elise and Phalippou, Ludovic and Westerfield, Mark M. [39] show that commitment risk—the cost of being bound to an illiquid obligation without knowing when the capital will be called—is a real cost for the LP, one that does not diversify away even across many funds, although diversification itself introduces a further complexity: the management of cash flows across overlapping vintages (see also [2]). The evergreen eliminates this risk, and it is arguably its most robust structural advantage. KKR's position (May 2026) is equally explicit: "both open-end and closed-end funds can play

[11] Redalpine, "Beyond closed-end: a guide to semi-liquid evergreen funds in venture" (blog, November 2025).

[12] Olin Business School / Washington University, "The Future of Evergreen Investment" (April 2026).

[13] UBS, "Accessing Private Equity with Perpetual Capital Funds" (November 2023), cited in Redalpine (2025).

an important role for private markets investors," framing the two as complements.[14] Hamilton Lane positions evergreen funds as an entry point for retail investors who lack the scale to manage diversified drawdown programmes, noting that institutions "typically opt for drawdown vehicles" because they have more tools to handle capital calls.[15]

What Redalpine's presentation does not mention, however, is the price of this efficiency: the liquidity buffer. A share of the capital—typically between 10 and 20 per cent—is held in liquid instruments to honour redemption requests [2, Section 2]. In venture capital this sleeve is particularly costly to maintain. Unlike private credit, where coupon flows replenish the buffer naturally, VC returns arrive in concentrated, unpredictable bursts—a significant exit may not occur for years—and in the meantime the sleeve must be refinanced through new subscriptions or secondary-market sales. It is a source of cash drag inside the fund that, in a drawdown, is marginal thanks to the commitment-and-call mechanism but that in an evergreen may be material, effectively replacing the external drag on the LP's portfolio even as it resolves it. The simulation presented later in the chapter will measure the relative magnitude of the two.

Under normal conditions a well-maintained buffer does not materially harm final returns—the problem, in short, does not arise. It does, however, make the IRR a less appropriate tool for communicating performance, because the fund permanently holds a low-yielding liquid pocket. The NAV accordingly becomes the instrument through which interim performance is reported and the yardstick for continuous subscriptions. The real issue with the buffer may lie elsewhere. The promise of liquidity is conditional: redemptions are typically capped at 5% of NAV per quarter, with notice periods of 30 to 90 days.[16] If requests exceed the cap they are fulfilled pro rata, with the surplus rolled to the next period. Should a sufficient number of investors seek to exit at the same time, the mechanism breaks down. Phalippou, Ludovic [37] warns that the semi-liquid structure is exposed to run-like dynamics when applied to illiquid assets—and venture capital is among the most illiquid asset classes of all. Clayton, William W. and de Fontenay, Elisabeth [38] describe the dynamic explicitly as "similar to bank runs, where investors are incentivized to beat

[14] KKR, "Evergreen Fund: What You Need to Know" (May 2026).

[15] Long Angle, "Private Equity Evergreen Funds: The Complete Guide" (November 2025), citing Hamilton Lane.

[16] Moonfare, "Evergreen Fund Structures in Private Equity" (April 2026); IQ-EQ, "Private Capital Perspectives" (December 2025).

each other to the exits". The case of BREIT, although it concerns real estate rather than VC, provides the most instructive stress test: between late 2022 and 2023 the fund satisfied only a fraction of redemption requests, raising doubts about whether reported NAVs reflected realisable prices [38]; [37]. In venture capital, where assets are even less liquid and cash-flow generation is virtually nil, the risk is potentially more severe.

The evergreen ecosystem in venture capital is, it should be said, still small and taking shape. Redalpine cites roughly two hundred evergreen VC vehicles worldwide, but a number of those listed—Novartis Venture Fund ($750 million), Samsung Catalyst Fund—are corporate venture capital operations funded from the parent company's balance sheet, with no external LPs holding redemption rights: structurally closer to a non-listed holding than to a semi-liquid fund. Others, such as Swisscom Ventures (evergreen since 2007, with external LPs from 2018), operate under a hybrid model whose liquidity mechanics are not fully comparable to an evergreen with NAV-based subscription and redemption. Among the better-known evergreen VC vehicles with external LPs and a semi-liquid mechanism are the Summit Fund itself, Sequoia (which restructured into an open-ended format in 2021), Hamilton Lane VC Evergreen, and StepStone SPRING—but the landscape is evolving and no exhaustive list is possible. Redalpine itself continues to manage traditional closed-end funds in parallel: its most recent, RAC VII, closed at $200 million in August 2024 in oversubscription.[17] The coexistence of the two models is not coincidental and underscores the point made above: the evergreen is intended as a complement to the drawdown, not as a substitute.

A particularly instructive case is that of AENU, the Berlin-based climate-tech fund founded by the Heilemann brothers in 2022. AENU was launched explicitly as an evergreen, with the declared objectives of contributing to the democratisation of the asset class, creating long-term alignment with founders, and offering innovative liquidity options to LPs. The first close, drawn from family, friends, and high-net-worth individuals—including Skype founder Niklas Zennström and several German unicorn CEOs—was 60% oversubscribed.[18] But in 2023, when AENU sought to open the raise to institutional LPs such as the European Investment Fund and KfW, the evergreen structure proved, in Heilemann's words, "not palpable to the

[17] Redalpine fund page and PitchBook.

[18] Sifted, 2 November 2022; TechFundingNews, 5 July 2022.

institutional LP market."[19] The fund was forced to convert to a traditional drawdown, ultimately closing at EUR 170 million in 2024—above its EUR 140 million target but under a radically different structure.[20] Heilemann, while still a believer in the merits of the evergreen, acknowledges that the format requires reputational capital that an emerging manager in an adverse market simply does not possess: "As an emerging manager in a market crisis, there's no way you can push through," he observes, with an explicit reference to Sequoia as a "very established" manager capable of "almost enforcing that upon their LPs."[21] His stated hope is to return to the evergreen format "in two or three fund generations' time." The lesson is clear: the evergreen in VC is conditional on the manager's bargaining power, the maturity of the LP base, and the prevailing market environment.

3.4.2 Non-listed investment holdings

The non-listed investment holding operates on a different logic altogether. LIFTT, in the Italian context, is a *Società per Azioni* with authorised capital of approximately EUR 100 million and 215 investors including banking foundations (Compagnia di San Paolo), universities (Politecnico di Torino), and private investors, founded in 2019 to facilitate technology transfer from research to market.[22] There is no subscription at NAV, no commitment in the drawdown sense, and no periodic redemption mechanism. Capital is raised through negotiated equity subscriptions in successive rounds attracting new shareholders progressively as the portfolio is built and the track record takes shape.

As Giuseppe Merlino, LIFTT's CFO, explains, the holding form was chosen precisely to avoid the closed-end fund structure: LIFTT replicates the logic of an evergreen— deploying capital as opportunities arise, without commitment-and-call mechanics—but without offering periodic liquidity to investors.[23] Unlike an evergreen, however, LIFTT does not raise capital for each individual transaction—which Merlino describes as "totally inefficient" given a pace of 15–20 investments per year. Instead, fundraising is planned around the deployment schedule: the team estimates how much capital can be deployed in the coming years based on the size of the target

[19] AENU press release, 11 September 2024.
[20] Sifted, 12 September 2024.
[21] Sifted, 12 September 2024.
[22] LIFTT S.p.A., institutional website (liftt.com); EIT Manufacturing.
[23] Interview with Giuseppe Merlino, CFO, LIFTT S.p.A.

market and the expected pipeline, and raises new rounds before the previous capital is exhausted, so that scouting and deal execution can proceed without interruption.[24] Cristina Odasso, Head of Business Analysis at LIFTT, confirms the same logic: if EUR 100 million are to be invested over three years, the company raises in tranches—say EUR 30 million at a time—building the pipeline so that capital can be allocated rapidly and is not left idle.[25]

There is no J-curve in the classical sense—no fees on uncommitted capital— but they may suffer cash drag from a liquidity sleeve, not becasue they have a redemption mechanism to support but because a cash buffer is needed for internal operations. Capital compounds in full over time. The major cost is total illiquidity: the investor who wants to exit depends on selling the stake to another shareholder or on an eventual IPO of the holding, with timing and pricing outside the investor's control. There are no gates because there is nothing to gate—but there is no structured exit either. Merlino notes that LIFTT has always been clear with its investors: this is patient-capital venture, and the investor must share that patience. The medium- to long-term exit options are two: dividend distributions, once realised gains are large enough that not all proceeds can be reinvested, or an IPO—with the decision ultimately resting with the shareholders in assembly.[26] He adds that certain institutional investors require a written assurance of return by a fixed date—typically year twelve—and that raising capital from such investors is difficult under a holding structure.[27]

3.4.3 Listed investment holdings

In the case of a listed holding, the investor buys shares on the secondary market at the prevailing market price; new capital is raised through rights issues or placements whose pricing depends on the market price, not on the NAV of the underlying portfolio. Retail access is unrestricted and liquidity is that of the stock market. The trade-off, however, is well documented: listed closed-end vehicles holding illiquid assets typically trade at persistent discounts to NAV, reflecting valuation uncertainty and the illiquidity of the underlying assets (Cherkes, Sagi & Stanton, 2009, cited

[24] Interview with Giuseppe Merlino, CFO, LIFTT S.p.A.

[25] Interview with Cristina Odasso, Head of Business Analysis, LIFTT S.p.A.

[26] Interview with Giuseppe Merlino, CFO, LIFTT S.p.A.

[27] Interview with Giuseppe Merlino. This observation mirrors Saraceni's point that the LP participation constraint, not GP preference, is what sustains the ten-year model.

in Clayton, William W. and de Fontenay, Elisabeth [38]). In venture capital the discounts are particularly pronounced, with VC portfolios on the secondary market changing hands in 2024 at discounts of 40–60% relative to reported NAV.[28]

The investor's return therefore depends on the price, not on the NAV, and the price is driven by sentiment. When the stock trades at a steep discount, a capital increase is dilutive for existing shareholders and expensive for the holding, which may preclude fundraising precisely when fresh resources are most needed—a constraint that mirrors the fundraising volatility of the drawdown model but is driven by public-market dynamics rather than the institutional fundraising cycle. An analysis of a basket of listed shares of large PE firms shows that exposure via listed equity has outperformed the respective flagship private funds over both five- and ten-year horizons.[40]. The result should be read with caution, since that exposure is to the manager as a business rather than to the portfolio as a fund, but it does suggest that the discount penalty is not necessarily destructive over the long run.

3.4.4 About performance metrics

In drawdown structures, as discussed, the metrics problem centres on the IRR — a measure that fails to capture true economic returns and is susceptible to manipulation through subscription credit facilities and the timing of cash flows. In permanent capital vehicles, the problem does not disappear; rather, it shifts from the IRR to the NAV, and in certain respects becomes more acute. The IRR loses its analytical relevance in evergreen structures because there are no discrete cash flows comparable to those of a drawdown fund, nor is there a defined sleeve (avoided by the commitment-and-call mechanism). Capital enters and exits at NAV without a well-defined start or end date for each individual investor. Performance is therefore measured against NAV appreciation over time, and NAV simultaneously becomes the performance metric, the transaction price for subscriptions and redemptions, and the basis on which management fees are calculated. This dual role of NAV — simultaneously a performance signal and a transaction price — is acknowledged by practitioners. Giuseppe Merlino, noted that an evergreen vehicle *"must publish its NAV at every reporting date: in practice, the unit value is computed as little more than a multiplication, yet it is the parameter the fund must consistently monitor"*. Merlino further observed however,

[28] Chronograph, "2024 Secondary Market Trends" (December 2025); VC Corner, "VC Secondaries Pricing Report" (August 2025).

that IRR retains analytical value at the portfolio level — measuring the return on each euro from investment to exit — but that the cash buffer inherent to evergreen structures structurally depresses it, since idle liquidity dilutes overall performance even as it reduces risk. In a drawdown fund, NAV is a periodic accounting estimate with few immediate transactional consequences — no investor enters or exits at that price — and management fees are charged on committed capital. In an evergreen vehicle, NAV carries direct cash consequences [37]: an overstated NAV causes incoming investors to overpay (effecting a wealth transfer to existing investors), allows redeeming investors to exit at above-fundamental values, and inflates the fee base of the manager. Valuation errors that would remain inconsequential in a drawdown fund become, in the evergreen context, direct channels of economic transfer. Phalippou, Ludovic [37] cite the Hamilton Lane case: interests in PE funds were acquired at approximately 80% of reported NAV and marked up to 100% the day after subscription — an illustration of how valuation discretion can translate into an instantaneous gain at the expense of incoming investors [37].

In venture capital, the problem is further amplified. Portfolio companies are typically pre-revenue, with no frequent market transactions to anchor their valuation, and the NAV of recent VC funds has become a persistent and economically significant component of reported value at advanced fund ages — no longer as transitory as in older vintages [41]. If this persistent, difficult-to-verify NAV becomes the basis on which evergreen transactions occur, the metrics problem of the drawdown structure is not merely replicated but worsened. This opacity is not merely theoretical. Diana Saraceni, of Panakès Partners, noted in interview that in venture capital, formal portfolio revaluations occur only upon a new financing round. Between rounds, limited partners have no independent signal of their holdings' true value. If such episodic, GP-controlled revaluations become the basis for evergreen subscription and redemption prices, the wealth transfer risk identified by Phalippou, Ludovic [37] is not attenuated but amplified.

For the non-listed holding company, the problem takes a different form: there is no transactional NAV (no investor enters or exits at NAV), so the risk of inter-investor wealth transfer does not arise. However, the investor lacks a transparent performance benchmark: the track record belongs to the holding as a corporate entity rather than to a comparable vintage fund, making benchmarking against drawdown and evergreen structures inherently difficult. For the listed holding company, NAV is replaced by the market price, which incorporates an illiquidity discount and prevailing market

sentiment — the metric is transparent, but volatile and potentially disconnected from the fundamental value of the underlying portfolio.

3.4.5 About democratisation

The most frequently invoked argument in favor of permanent capital vehicles is the democratization of access. Evergreen vehicles in particular guarantee access to more people due to their lower minimum ticket. KKR (2024) observes that institutional investors have been able to take advantage of the illiquidity premium and diversification benefits of private equity for several decades, while broader adoption among individual investors is only in its early stages.[29] Evergreen vehicles, by contrast, are meant to simplify and broaden access for individual investors. Hamilton Lane estimates that evergreen funds today account for roughly 5% of private markets (approximately $700 billion) and could reach 20% within a decade.[30] iCapital reports that Partners Group holds 30% of its assets under management in evergreen structures and is anticipated to launch six new evergreen products for the wealth channel.[31]

Access, however, is not a guarantee of returns. Balloch, Mainardi, Oh and Vokata (2026) document intermediation frictions that are unique to private equity and correlated with investor wealth: the most affluent individual investors outperform the least affluent by 9 percentage points in risk-adjusted public market equivalent terms, a gap driven largely by differences in financial advisor quality rather than by fund selection itself [42]. Broadening access without addressing these intermediation frictions does not automatically translate into investor benefit. Pegoraro, Shive and Zambrana (2026) reach a complementary conclusion for semi-liquid vehicles more broadly: retail investors benefit from these structures only in asset classes with

[29] Alisa Amarosa Wood, Paula Campbell Roberts and Justin Park, "An Evergreen Vehicle Can Be an Important Tool in Private Equity Asset Allocation," KKR Insights, May 2024, https://www.kkr.com/insights/evergreen-vehicle.

[30] Hamilton Lane, *2025 Market Overview* (March 2025), https://shareholders.hamiltonlane.com/2025-03-12-2025-Hamilton-Lane-Market-Overview-Private-Markets-Reach-an-Inflection-Point,-though-Long-Term-Fundamentals-Remain-Strong; see also Long Angle, *Evergreen Fund Investment Guide 2025*, https://www.longangle.com/alts-education/evergreen-funds, reporting the same figures with Hamilton Lane as primary source.

[31] iCapital, "The Future is Evergreen: The Next Generation of Private Market Funds" (January 2024), https://icapital.com/insights/private-equity/the-future-is-evergreen-the-next-generation-of-private-market-funds/; the underlying AUM figure is dated to SEC filings as of December 31, 2022.

low information sensitivity, such as private credit, where outperformance is driven primarily by distribution income; in information-sensitive strategies such as listed equity, interval funds underperform, and retail investors achieve superior returns only when they co-invest alongside larger, more sophisticated investors [36].

Phalippou (2026) adds a further layer to this picture. Semi-liquid vehicles import the institutional features of private equity—discretionary valuation, opaque fees, weak governance—into retail channels where investors cannot negotiate terms. Fees stack across layers: management and performance fees at the fund level, plus distribution and platform costs at the retail interface, further eroding the margin available to investors [37]. The conclusion is that value creation is not automatic through broad access alone; it requires transparency, governance, and discipline.

Lerner and Schoar (2004), already discussed above, argue that illiquidity is not merely a cost borne by investors but a screening device that selects LPs with a low probability of needing liquidity, thereby lowering the cost of capital in future fundraising rounds—even though, as shown earlier, some fundraising frictions persist regardless of investor sophistication [35]. Illiquidity is amplified through these new retail-facing vehicles, yet retail access sets aside the consequences of this screening mechanism and the protection that lock-ups previously offered to the stability of the vehicle and to other investors. The result is a fundraising process that remains complicated and a management task that is equally complex for managers who must confront new challenges they may not be prepared for (see the AENU case).

Democratization, in short, is real as a market phenomenon but conditional as a benefit: it works for certain assets and certain client bases, not universally, and its outcome for venture capital—where information sensitivity is at its highest—remains an open question.

3.5 Drawdown vs PCVs: A simulatated comparison

The preceding pages have analyzed how the drawdown model and permanent capital vehicles behave during the fundraising phase: the tensions inherent in drawdown, the responses offered by evergreen structures and holding companies, and the new problems these responses introduce. The limitation of this qualitative analysis, however, is that it identifies trade-offs without measuring them. How much is

the immediate-deployment advantage of evergreen or holding structures actually worth? How does the cash drag from an evergreen liquidity sleeve compare to that of uncalled committed capital in a drawdown fund? Answering these questions requires numbers.

The starting point is the highly relevant work of Brown and Volckmann (2025), the first study to rigorously compare, through a Monte Carlo simulation, the returns, risk, and cash flows of an investor choosing among three strategies: a single drawdown fund, an evergreen fund, and a staggered series of drawdown funds [2].

The setup is simple, and for that reason powerful. An LP holds $100 and seeks a constant 90% exposure to private assets and 10% in cash. All three vehicles invest in the same underlying assets, with identical market returns and identical independent alpha: the only difference is structural, not a matter of selection. Returns are calibrated on the net performance of buyout funds in the MSCI-Burgiss data, with a market beta of 1, average market return of 12% (standard deviation 18%), average alpha of 3% (standard deviation 5%), and a risk-free rate of 3%. Capital calls and distributions are modeled as random variables with ranges similar to those observed empirically. The simulation is run 10,000 times per strategy, and each strategy is liquidated at the end of year 12.

The single drawdown fund returns 6.5% annually on average, but is invested only 28.7% of committed capital on average, because for most of the fund's life capital sits in cash at the risk-free rate—"locked in" by the commitment but not yet called through unpredictable capital calls. The evergreen fund returns 13.5%, remaining 90% invested by construction. The drawdown series returns 12.7%, also reaching a 90% average invested share, but with a standard deviation of 19.5% (against 0% for the evergreen fund); the average fund beta is approximately 0.91 for both the evergreen fund and the series, closely tracking the average 90/10 allocation [2, Table 1]. The key result is that the evergreen fund outperforms the series even though both invest in the same underlying assets, because it keeps the 90/10 ratio constant. The mechanism is this: for a given average, a constant allocation produces higher compounded returns than a variable one. Counterintuitively, the evergreen fund shows greater dispersion in total return and in alpha but less in beta, because beta is proportional to the invested share, which for the evergreen fund is fixed; in the series, by contrast, an LP forced to hold an unexpected distribution in cash while awaiting a new capital call temporarily lowers portfolio beta, while an LP forced to borrow

to meet an unexpected capital call temporarily raises it—in fact, the series carries a negative cash position (i.e., is borrowing) in about 10% of simulated years. The series' cash flows are complex (reaching up to 146% of committed capital in a single period), cyclical, and unpredictable; the evergreen fund's cash flows are trivial—one outflow at time 0, one inflow at time 12.

The authors themselves acknowledge four limitations. First, the perfectly maintained 10% cash buffer is unrealistic, since in practice the manager must actively manage the liquidity sleeve. Second, fees are not modeled. Third, funding costs are asymmetric—borrowing costs typically run 1–3 percentage points above lending rates—and this asymmetry is not included, even though it would weigh on both drawdown-series performance (when the LP must borrow to meet a capital call) and evergreen performance (when the fund must borrow to meet redemptions). Fourth, the valuation pressure on evergreen NAV is not modeled: because new investors enter and existing investors exit at reported NAV, any systematic bias in that NAV—in either direction—creates a wealth transfer between old and new investors, an incentive for gaming that the simulation does not capture. The calibration is based entirely on buyout data, with a beta of 1 and a symmetric, normally distributed alpha, assumptions that do not hold for venture capital [2, §4].

The simulation's central mechanism finds independent support in industry practice, even outside the rigor of a Monte Carlo framework. Long Angle notes that, on a strict return-on-deployed- capital basis, drawdown funds tend to outperform evergreen vehicles marginally: once capital is called, a drawdown fund is fully invested in private assets, whereas an evergreen fund must permanently hold a liquidity sleeve to fund periodic redemptions.[32] This effect is present in B&V's own results as well: the evergreen fund's return (13.5%) remains below the return implied by a fully invested underlying asset (15%, given by a 12% market return plus a 3% alpha) precisely because of the liquidity sleeve [2, Table 1].

This drag, however, is second-order relative to the dominant effect: that of capital not yet deployed. A drawdown fund takes years to call the full commitment, remaining invested for only a minority of that time—28.7% on average for a single fund— while an evergreen fund allocates 90% of committed capital from day one [2, Table 1]. It is this second channel, not the first, that explains the wide gap between the single

[32]Long Angle, *Evergreen Fund Investment Guide 2025* (October 2025), https://www.longangle.com/alts-education/evergreen-funds, Table 1, row "Total Expected Return."

drawdown return (6.5%) and the evergreen return (13.5%). KKR quantifies the same mechanism in terms of the net IRR required to reach an equivalent net multiple on invested capital (MOIC) over a ten-year horizon (Table 3.1), and Long Angle summarizes it through a comparable Hamilton Lane analysis over an eight-year horizon: an evergreen vehicle fully deployed from inception needs only a 10% annualized return to reach a 2.1x MOIC, against 16% for a closed-end fund with dispersed capital calls.[33]

Table 3.1 Net IRR Required to Reach a Given Net MOIC over a Ten-Year Horizon

	2.5x	3.0x	3.5x
Evergreen vehicle, net IRR	9.6%	11.6%	13.3%
Drawdown fund, net IRR	13.5%	18.4%	22.5%

Source: KKR, "An Evergreen Vehicle Can Be an Important Tool in Private Equity Asset Allocation" (May 2024), https://www.kkr.com/insights/evergreen-vehicle, Exhibit 2b. Note: illustrative deterministic example, not a stochastic simulation; assumptions differ from those of Brown and Volckmann (2025) and from the Hamilton Lane example cited above.

This validation, however, applies only to B&V's original calibration, which is based entirely on buyout data. Venture capital presents structurally different characteristics—a heavily back-loaded distribution profile (the well-known VC J-curve) and a higher market beta—that require the model to be recalibrated before any conclusions can be drawn for the object of this thesis.

The simulation extends Brown and Volckmann (2025) in several directions. First, the model is recalibrated for venture capital. VC parameters differ substantially from those of buyout, and each parameter choice is discussed in turn below.

3.5.1 Parametres

Beta. Brown, Lundblad, and Volckmann (2025) report fund-level beta estimates for VC funds obtained through three independent methods applied to the MSCI-Burgiss database: the Dimson (1979) lagged market-model regression yields a beta of 1.73, the Korteweg and Nagel GPME model yields 2.33, and the Brown, Gredil, and Ghysels nowcasting approach yields 1.45 [43]. The Dimson estimate is adopted for two reasons. Methodologically, the Dimson approach accounts for the NAV

[33] Long Angle (2025), citing Hamilton Lane, *Evergreen Funds and Private Wealth*, 2025 Market Review, https://explore.hamiltonlane.com/2025-market-overview/evergreen-funds.

smoothing implicit in quarterly fund reporting, which is precisely the phenomenon that makes standard OLS estimates on private fund data unreliable. Moreover, beta depends on the level at which it is measured: using deal-level betas (2.33 from Korteweg-Nagel, or higher) would implicitly assume that the LP invests in a single company with no diversification, which is not the case. A lower, more conservative beta is therefore used, consistent with the diversification achieved by a fund holding multiple startups.

Alpha. The same authors report a net Direct Alpha of just 0.2% annually for U.S. VC funds, after risk-adjustment at $\beta = 1.73$, compared with 4.4% for buyout funds [43]. This near-zero value suggests that the idiosyncratic component of VC fund performance—the manager's skill premium net of market exposure—is itself close to zero before fees, in clear contrast to the $\mu_\alpha = 3\%$ assumed by Brown and Volckmann for buyout, which corresponds to their own reported net Direct Alpha of 4.4%. On this basis, the simulation adopts $\mu_\alpha \approx 0\%$ for VC.

Return dispersion. VC returns are characterized by extreme positive skewness: a small number of exceptional investments dominate aggregate returns, while the majority of portfolio companies return less than invested capital. Cochrane (2005) documents a standard deviation of arithmetic (log) returns of 89% at the individual-deal level, across a large sample of venture-backed firms [44, p. 18]. At the fund level this dispersion compresses considerably through diversification across 20–50 companies, though it remains materially wider than for buyout funds. Korteweg and Sorensen (2015) report a standard deviation of the ten-year log return of 2.623 for VC funds against 1.552 for buyout funds—a ratio of approximately 1.7 to 1—while the standard deviation of IRR shows an even wider gap, 54.8% for VC against 18.6% for buyout [45, Table I]. Guided by these figures, the simulation adopts $a = 10\%$ as the fund-level dispersion parameter for VC, roughly double a typical buyout fund-level figure of 5%, a simplification within the empirically documented range rather than an exact replication of either ratio.

Market volatility. VC funds invest predominantly in early-stage technology and life-sciences companies, sectors with above-average covariance relative to growth equity indices. Korteweg and Nagel (2024) calibrate their own simulation exercise

with a market factor volatility of 15% and an idiosyncratic volatility of 25% per year for VC-like funds [46, Appendix C]. Building on this evidence of the two-component risk structure, and consistent with the higher-beta, higher-covariance profile of early-stage technology and life-sciences investments, market return volatility is increased from the 18% assumed by Brown and Volckmann for buyout to 22% for VC in the present simulation—a moderate upward adjustment relative to the buyout baseline, in the same direction as, though not numerically identical to, the calibration used by Korteweg and Nagel.

3.5.2 The simulation

Before extending the Brown and Volckmann (2025) framework to venture capital, the buyout baseline itself was independently rebuilt from the description in the paper's Appendix A—which reports the simulation's parameters, capital call schedule, and distribution schedule—and tested against the authors' own published results in Table 1 (single drawdown fund and evergreen fund columns; the series of drawdown funds is treated separately and is not part of this validation, see the discussion in Subsection: Series discussion. Running 10,000 independent simulations, the rebuilt model reproduces the order of magnitude and the direction of every metric reported in Table 1, with deviations generally within one to two percentage points: the single drawdown fund returns an annualized 6.3% against B&V's reported 6.5%, while the evergreen fund returns 12.6% against their reported 13.5%; the average percentage invested for the single drawdown fund is 26.9% against a reported 28.7%. These deviations were verified to be stable rather than the product of sampling noise: increasing the number of simulations by an order of magnitude did not narrow the gap, which indicates a systematic rather than a random source of divergence.

The most plausible explanation is the absence of the authors' original random seed and source code, neither of which is disclosed in the paper. Because the evergreen fund's return compounds a single random draw over twelve consecutive periods, even small differences in the underlying pseudo-random stream are amplified multiplicatively by the time the simulation reaches year twelve—consistent with the pattern observed here, where the deviation is smallest for quantities with little compounding (the single drawdown fund's alpha deviates by only 0.05 percentage points) and largest for quantities that compound over the full horizon (the evergreen fund's annualized return and alpha deviate by roughly 0.9 percentage points each).

Unpublished implementation choices not fully specified in the text—for instance, the exact order in which annual market returns, idiosyncratic alpha, capital calls, and distributions are drawn—may also contribute marginally, though the underlying mechanics were verified line by line against the paper's own description and match it in structure.

Because an exact bit-for-bit replication is not achievable without the original seed, these residual deviations are treated here as a disclosed limitation rather than as evidence of a structural error. The rebuilt simulation is structurally and conceptually faithful to Brown and Volckmann's design, reproduces every qualitative result of Table 1 (the evergreen fund's higher average return, its full capital deployment relative to the drawdown fund's average 28.7% invested share, and comparable risk-adjusted performance measures), and is therefore treated as an equally valid basis for the extension to venture capital that follows, on the understanding that the absolute figures carry a margin of one to two percentage points relative to the original paper.

3.5.3 Series discussion

The series of overlapping drawdown funds—the third strategy in Brown and Volckmann's (2025) comparison—is discussed here only at the level of the mechanism the authors describe, not as an independently reproduced quantitative result. An attempt to rebuild the six-fund staggered structure of Appendix A.3 was made, but it proved unstable in a way that could not be resolved through recalibration: even in a minimal three-fund configuration (the two seasoned funds plus a single new commitment, without any of the dynamic resizing described for later vintages), the reconstructed series exhibited a negative cash position in 55–70% of simulated paths during the first two evaluation years, far in excess of the "about 10% of years" the authors themselves report for their own series simulation. Because B&V's own code and random seed are not published, it was not possible to determine whether this instability reflects a missing implementation detail or a feature of the mechanism itself under different draws. Rather than present series-level figures whose reliability could not be established, this thesis limits itself to a qualitative discussion—developed in the following section—of how the series' reliance on timely recycling between seasoned and newly committed funds should be expected to interact with the venture capital J-curve.

3.5.4 Extending the Evaluation Horizon for Venture Capital

Brown and Volckmann (2025) force full liquidation of all simulated funds at the end of year 12, a choice the authors themselves describe as a modeling convenience adopted to keep the simulated time frame manageable; they further note that extending the fund-life cap to 15 years would not meaningfully alter their main results [2]. That robustness claim, however, is made exclusively for their buyout calibration, and there is no reason to assume it extends to venture capital without independent verification.

Bartlett and Ramella (2026) provide exactly this verification, and their findings suggest that a twelve-year horizon is considerably more restrictive for private capital generally, and for venture capital in particular, than B&V's own robustness check implies. Using linked quarterly fund cash flow data for funds formed between 1995 and 2014, the authors document that a substantial share of both private equity and venture capital value is realized only after the traditional ten-year fund term. The pattern, however, is markedly asymmetric across asset classes. For buyout-style private equity, the shift is present but comparatively contained, with distribution profiles remaining broadly similar across vintage cohorts through year ten and only a modest, if economically meaningful, continuation of distributions thereafter [41]. For venture capital, by contrast, the shift is pronounced: for VC vintages formed after 2000, the authors report that median cumulative distributions continue to grow substantially beyond year ten, increasing by roughly 70 to 120 percent relative to their year-ten level by the end of the observation window [41]. This implies that, for the median fund, somewhere between 45% and 59% of eventual cumulative distributions had been received by year ten, with the remainder arriving over the following decade.

This asset-class asymmetry motivates treating the evaluation horizon itself as a parameter that should be recalibrated for venture capital, rather than inherited from B&V's buyout-tuned convenience choice. The simulation is accordingly extended to $T = 20$ for the VC calibration — corresponding to the maximum horizon over which Bartlett and Ramella track the most data-complete vintage cohort, those formed between 2000 and 2004 [41] — while the buyout calibration is left at B&V's original $T = 12$, preserving comparability with their published results.

It should be stated plainly that Bartlett and Ramella do not provide a year-by-year distribution schedule at the granularity the simulation requires; their reported evidence describes only the normalized cumulative distribution path relative to a year-ten baseline, not the underlying annual rates. The year-by-year schedule used for $T = 20$ is therefore the author's own construction, calibrated so that the simulated median share of cumulative distributions received by year ten falls within the 45–59% range implied by the paper — a target the simulation meets, with a median of 50.5% verified automatically at runtime. The specific annual rates should be understood as a stylized approximation consistent with this empirical anchor, not as a direct transcription from the source.

Figure 3.1 plots the resulting exposure profile for a single drawdown fund under each calibration: the buyout fund (B&V, $T = 12$) against the venture fund ($T = 20$). Two features stand out. First, the venture fund's invested share never approaches the near-complete drawdown the buyout fund exhibits by year 12; instead of a single fund life concentrated within a twelve-year window, the venture fund's capital remains partially deployed across a substantially longer horizon, consistent with the slower and more protracted realization pattern Bartlett and Ramella document. Second, extending the horizon lowers, rather than raises, the fund's time-averaged invested share (31.7% for VC across 20 years, against 39.7% under a earlier 12-year calibration): the additional years 13 through 20 are predominantly a late-stage wind-down phase, during which an increasing share of NAV has already been converted to cash awaiting final distribution, and their inclusion in the average pulls the overall figure down even though capital remains committed for longer in absolute terms. This is a mechanical consequence of averaging over a longer window that includes an extended tail of low-exposure years, not evidence that VC investors are, at any given point in time, less exposed to private assets than the shorter-horizon calibration suggested.

3.5.5 Results

The evergreen fund outperforms the single drawdown fund by 5.65 percentage points annualized. The gap is driven almost entirely by exposure: the evergreen fund is invested at a fixed 90%, against a time-averaged 31.72% for the drawdown fund, while the expected return on invested capital is high ($\beta\mu_m = 1.73 \times 12\% \approx 20.8\%$) relative to the 3% risk-free rate earned on uninvested cash. Evergreen alpha

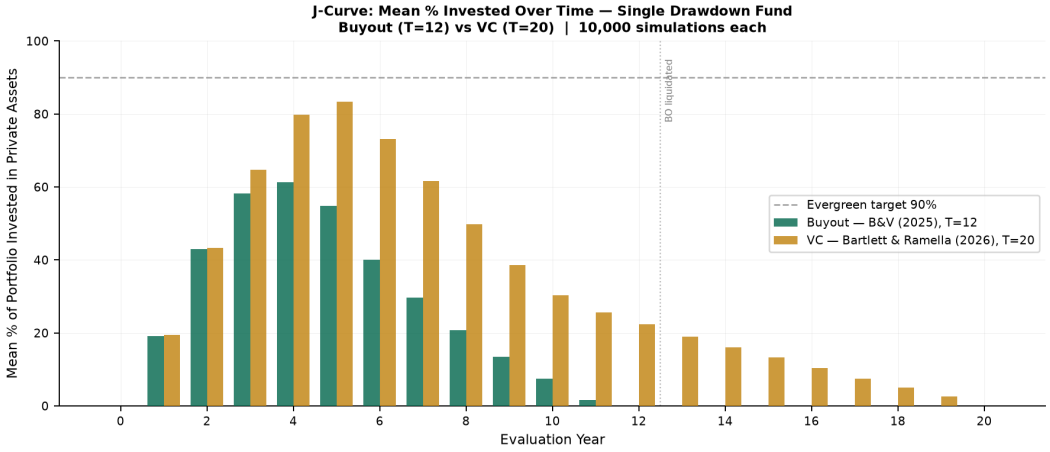

Fig. 3.1 Mean percentage of portfolio invested in private assets over time, single drawdown fund, buyout ($T = 12$, Brown and Volckmann 2025 calibration) versus venture capital ($T = 20$, calibration described in the text). Based on 10,000 Monte Carlo simulations per asset class. The venture capital distribution schedule is the author's own construction, calibrated to match the cumulative-distribution timing documented in Bartlett and Ramella (2026).

Table 3.2 Single Drawdown Fund vs. Evergreen Fund, Venture Capital Calibration ($T = 20$, $N = 10{,}000$ simulations each)

Metric	Single Drawdown	Evergreen
Annualized return (CAGR) — mean	7.57%	13.22%
Annualized return (CAGR) — sd	4.31%	8.64%
% invested — mean	31.72%	90.00%
% invested — sd	4.13%	0.00%
Alpha (Jensen) — mean	0.52%	1.99%
Beta — mean	0.573	1.557
Beta — sd	0.1704	0.0990

(1.99%) exceeds drawdown alpha (0.52%), the same direction B&V report for buyout, consistent with constant allocation capturing idiosyncratic performance more fully than a variable one. The evergreen beta (1.557) matches $0.9 \times 1.73 = 1.557$ exactly, confirming internal consistency with the fixed 90/10 construction; the drawdown fund's lower mean beta (0.573) reflects its lower average exposure, while its higher beta dispersion (sd 0.1704 vs. 0.0990) reflects that exposure varying substantially across simulations and over time. This same mechanism — a fixed, higher beta transmitting market volatility directly to returns — explains why the evergreen fund's CAGR is also more dispersed (sd 8.64% vs. 4.31%) despite its investment policy being, by construction, the more mechanically simple of the two.

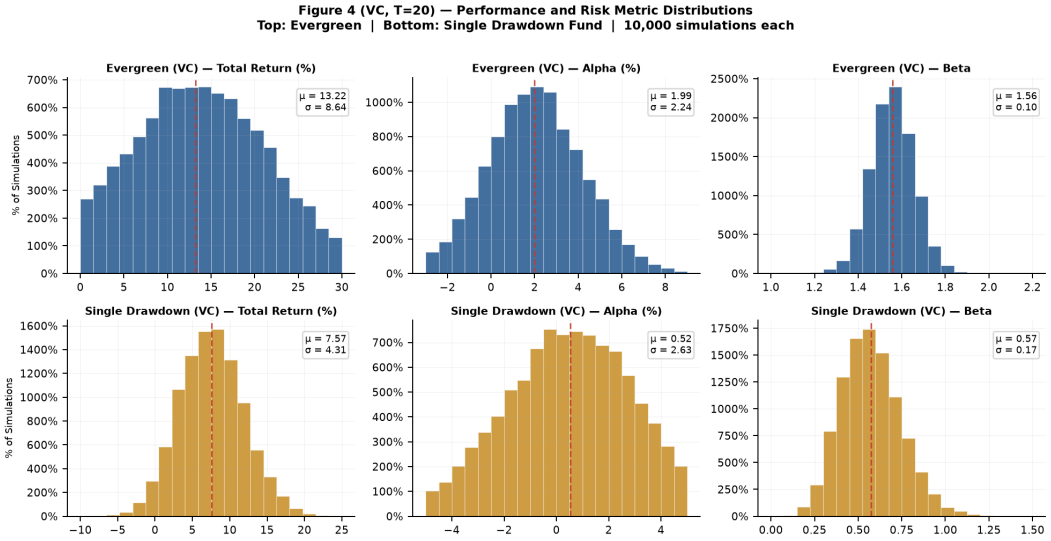

Fig. 3.2 Distributions of annualized return, alpha, and beta across 10,000 simulations, evergreen fund (top) and single drawdown fund (bottom), venture capital calibration, $T = 20$.

Figure 3.2 shows the evergreen fund's wider dispersion in total return and alpha alongside its narrower beta dispersion, the same qualitative pattern B&V report for buyout [2, Table 1]: a fixed 90/10 allocation transmits market variance directly into return and alpha variance, while the drawdown fund's fluctuating exposure instead concentrates the variance in beta itself.

Figure 3.3 contrasts the two cash flow profiles: the evergreen fund's is trivial by construction, a single outflow at $t = 0$ and a single inflow at $t = T$; the drawdown fund's shows sustained two-sided activity across the full horizon, with the extended $T = 20$ window making visible the prolonged distribution tail documented in Bartlett and Ramella (2026, Sec. 3.2) that a $T = 12$ window would have truncated.

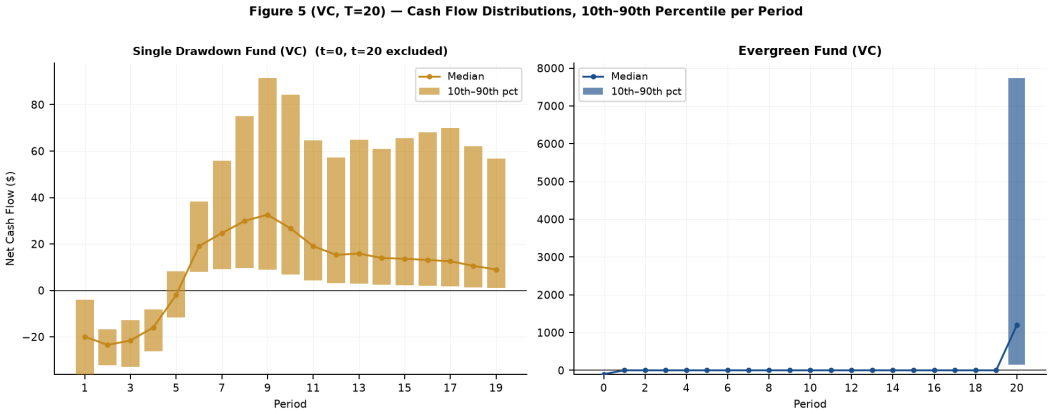

Fig. 3.3 Net LP cash flow, 10th–90th percentile per period, single drawdown fund (left, periods 1–19) and evergreen fund (right, all periods), venture capital calibration, $T = 20$.

About the qualitative discussion of the series of overlapping drawdown funds developed in subection: Series discussion, we can now extend in light of the extended-horizon evidence presented above. The premise that venture capital's slow realization profile complicates the cash-recycling mechanism underlying B&V's series construction was already well founded on the fund life-cycle described in the literature: capital is called rapidly — Bartlett and Ramella (2026) report that, for the fastest-calling cohort, median VC funds had drawn 100% of committed capital within 6.5 years of formation [41, Sec. 3.1] — while distributions follow a markedly slower and more protracted schedule.

The evidence assembled in this chapter sharpens this concern. Bartlett and Ramella's finding that a substantial share of VC value is realized only after year ten, with distributions continuing to accrue steadily through year twenty [41, Sec. 3.2], indicates that the horizon over which recycling would need to operate is considerably longer than B&V's own twelve-year buyout window. Layering a substantially slower and more extended venture capital distribution profile onto a mechanism that is already complex under its own native calibration gives no reason to expect the recycling chain would become more, rather than less, reliable.

Even setting this instability aside, a hypothetical series of VC funds attempting to replicate a 90/10 evergreen-like exposure over a twenty-year window would require substantially more active cash flow management than B&V's own buyout series, which they already describe as involving cash flows reaching up to 146% of the initial commitment in magnitude within a single period [2, Sec. 3]. A twenty-year VC

horizon would require a correspondingly larger number of rebalancing cycles across overlapping vintages, compounding rather than resolving the underlying mismatch between call speed and distribution speed. This thesis therefore treats the series of VC drawdown funds as an open question rather than a quantified result: the mechanism B&V describe for buyout does not straightforwardly transfer to an asset class whose realization horizon is both slower and, per the most recent evidence, still lengthening.

3.5.6 Net of fees

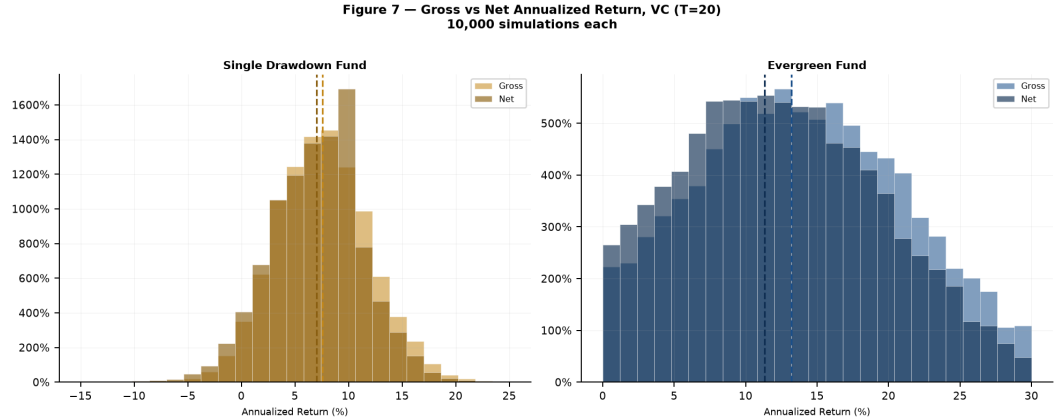

Fig. 3.4 Gross vs. net annualized return, single drawdown fund and evergreen fund, VC calibration, $T = 20$.

Table 3.3 Gross vs. Net Annualized Returns, Venture Capital Calibration ($T = 20$)

Metric	Single Drawdown	Evergreen
CAGR gross — mean	7.57%	13.22%
CAGR net — mean	7.03%	11.32%
Fee drag (gross − net)	0.54pp	1.90pp

Net of fees, the Evergreen–Drawdown gap narrows from 5.65 to 4.28 percentage points, driven by an asymmetric fee drag: the evergreen's recurring 1.68% management fee, compounded over twenty years, erodes final NAV by roughly 29%, while the drawdown fund's carried interest — averaging $49.67 per $100 committed, but collected only once at liquidation and only in the 32.1% of paths that clear the 8% hurdle — has a proportionally smaller effect on the annualized rate. A fee's

weight on CAGR depends on its timing and recurrence, not only its nominal size or its magnitude in dollars. This is consistent with, though not directly established by, Metrick and Yasuda's (2010) discussion of GP compensation, where management fees are characterized as a comparatively predictable revenue stream and carried interest as contingent and concentrated in a single realization event. The result is specific to this calibration (twenty-year horizon, 1.68% evergreen fee) and should not be generalized; it also represents an upper bound on the evergreen's net advantage, since no performance fee is applied to the evergreen fund here despite both the Merlino interview and Brown and Volckmann (2025, Sec. 1, Sec. 4) confirming that such a fee is typical in practice.

3.5.7 Introducing liquidity stress and gating

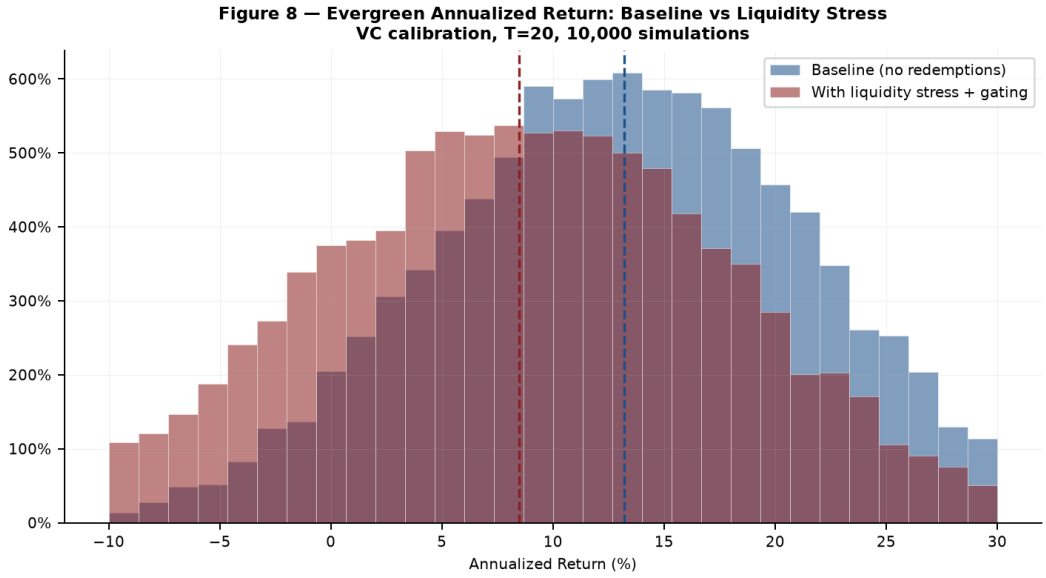

Fig. 3.5 Evergreen annualized return: baseline vs. with liquidity stress and gating, VC calibration, $T = 20$.

Individually, each friction erodes a meaningful share of the evergreen fund's gross advantage: the recurring management fee costs 1.90 percentage points, compounded over the twenty-year horizon; the liquidity stress and gating mechanism, calibrated to BREIT's actual 2022–2024 redemption episode, costs a further 4.73 points on its own. Combined, the two frictions are not merely additive: the evergreen fund's net return (6.61%, averaged across six independent simulation seeds, standard deviation

Table 3.4 Evergreen Fund: Return Drag from Fees, Liquidity Stress, and Their Combination (VC Calibration, $T = 20$)

Component	CAGR (mean)
Gross (Layer 1 baseline)	13.22%
Net of management fee only (Layer 2)	11.32%
Gross, with liquidity stress only (Layer 3)	8.49%
Net of fees *and* liquidity stress (combined)	6.61%
Drawdown fund, net of fees (Layer 2, for comparison)	7.03%

0.02 percentage points across seeds) falls *below* the drawdown fund's own net-of-fees return (7.03%), a gap of -0.43 percentage points. This reversal held in 100% of the seeds tested, indicating it is not an artifact of a particular random draw.

This result should not be read as a verdict that the evergreen structure is inferior to the drawdown fund in general. It is better read as quantitative support for the complementarity argument developed throughout this thesis: the evergreen fund genuinely resolves the exposure-timing inefficiency documented in the gross comparison (Chapter 3; Layer 1 above), but does so by taking on a different exposure — to liquidity risk in periods of redemption stress — of comparable magnitude. Neither structure is unconditionally superior once both frictions are priced in; each substitutes one source of fragility for another. The frequency and persistence of the gating episodes driving this result are examined below.

The frequency and persistence of gating in the simulation are worth reporting alongside the return figures. Gating binds — meaning requested redemptions, including any carried-over backlog, exceed the 20% annual cap — in 8.29% of all simulated fund-years. Over a full twenty-year path, this amounts to an average of 1.66 gating-bound years per simulation, and 77.7% of simulated paths experience at least one such year; only 22.3% complete the full horizon without ever binding against the cap.

Decomposing these gating years shows that 85.4% coincide with a market-triggered stress episode in that same year, while the remaining 14.6% occur in years with only normal (2%) redemption demand, where the cap nonetheless binds because of unresolved backlog carried over from an earlier stress year. This spillover requires closely spaced or repeated stress episodes: a single stress year leaves a backlog too small, on its own, to bind the cap again under normal demand. The pattern is

consistent, at least qualitatively, with BREIT's own experience, where redemption pressure persisted for sixteen consecutive months rather than resolving within the single quarter that triggered it.

The debt taken on to fund gated redemptions is similarly persistent. Across all simulations, the average debt episode — a contiguous run of years with a positive fund-level borrowing balance — lasts 3.99 years (median 3), with 21.9% of episodes extending five years or longer. At the twenty-year horizon, 26.7% of simulated funds still carry an outstanding balance that has not been fully repaid within the simulated period.

3.5.8 Non-listed/listed holdings considerations

Table 3.5 Evergreen (90/10) vs Investment Holding, Non-Listed (70/30), Gross Comparison (VC Calibration, $T = 20$)

Metric	Evergreen (90/10)	Holding (70/30)
CAGR — mean	13.22%	12.07%
CAGR — sd	8.64%	6.40%

The gap between the evergreen fund and the non-listed holding (1.15 percentage points) should not be read as evidence that the holding structure is inherently less exposed, or less performant, than the evergreen fund. The 70/30 split is an observation specific to LIFTT's current stage of portfolio construction, not a structural feature of the holding form itself; as Merlino notes, the holding and evergreen structures are, in his own account, substantively similar. Nothing in the holding's legal or operational structure precludes it from reaching a 90/10 allocation comparable to the evergreen fund's, and LIFTT's own trajectory (21.0% invested in 2019, rising non-monotonically to 69.6% by 2024, with a stated intention to expand the portfolio company count further) suggests this gap may close, or close further, as the vehicle matures. The comparison in Table 3.5 should therefore be read as illustrating a single observed case at a particular point in its life cycle, not as a claim about the relative merits of the two structures.

The listed-holding illustration in Table 3.6 and Figure 3.6 models a fixed 40% discount to NAV, calibrated on Molten Ventures plc (LSE: GROW), the closest available structural comparable to a listed venture capital holding company. The

Table 3.6 Listed Holding: NAV vs. Discounted Market Value by Year (40% Fixed Discount)

Year	Mean NAV	Mean Market Value	Discount ($)
0	100.00	60.00	40.00
5	201.66	120.99	80.67
10	416.54	249.92	166.62
15	852.22	511.33	340.89
20	1751.92	1051.15	700.77

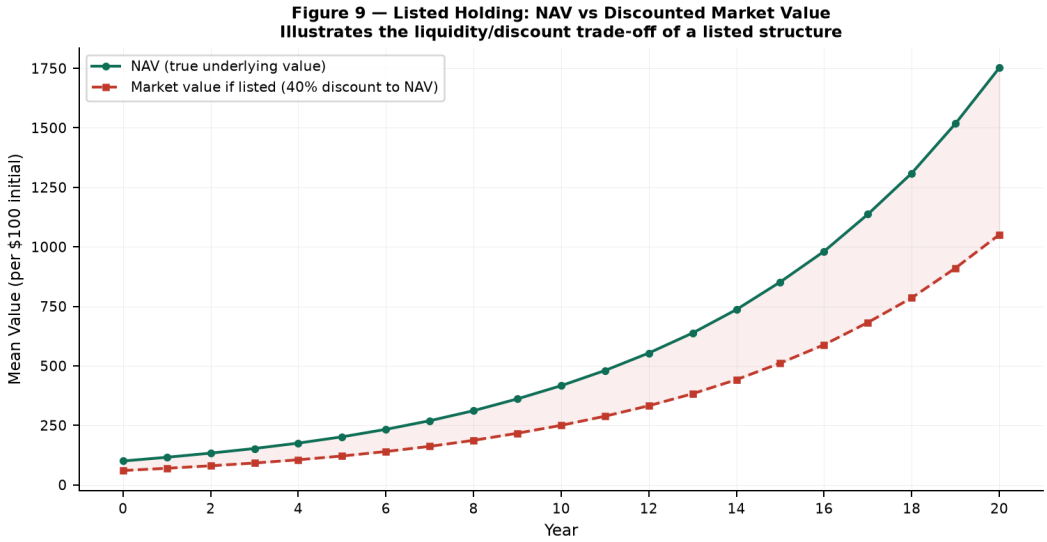

Fig. 3.6 Listed holding: mean NAV versus mean market value assuming a constant 40% discount to NAV, VC calibration, $T = 20$, 10,000 simulations. Discount magnitude calibrated on Molten Ventures plc (LSE: GROW).

observed discount has ranged from 53–66% between 2023 and mid-2025[34] to 29% as of the most recent observation.[35] 40% is used here as an approximate midpoint between the most recent reading and the more typical preceding range, not a formal statistical average. The model applies only a discount, not a premium, consistent with Molten's own observed history and with the broader tendency of listed vehicles holding illiquid, hard-to-value assets to trade below rather than above NAV, as the theoretical model in Cherkes, Sagi and Stanton (2009) predicts: their framework explains "why funds are issued at a premium to net asset value (NAV), and why they later usually trade at a discount" [47]; other listed PE/VC holding companies have, under different portfolio conditions, traded at a premium instead, so this pattern should be read as typical rather than universal.

This deviation between market price and NAV, however, affects only investors who transact in the secondary market at a given moment — it does not alter the holding's underlying NAV, and therefore does not affect the position of shareholders who simply hold their stake. A discount transfers value from the investor who sells to the investor who buys at that moment; a premium does the reverse. The cost illustrated in Table 3.6 is therefore specific to an investor who requires liquidity at that particular point in time, not a cost borne uniformly by all shareholders of the listed structure.

[34] Edison Group, "Molten Ventures — NAV down c 6% in the first half of FY24," November 2023, https://www.edisongroup.com/research/nav-down-c-6-in-the-first-half-of-fy24/32914/; Kepler Trust Intelligence, "Molten Ventures (GROW)," July 2025, https://www.trustintelligence.co.uk/investor/articles/fund-research-investor-molten-ventures-retail-jul-2025.

[35] Kepler Trust Intelligence, "Molten Ventures (GROW)," January 2026, https://www.trustintelligence.co.uk/articles/fund-profile-molten-ventures-jan-2026.

Chapter 4

Investing and Monitoring

4.1 Introduction

This chapter addresses the investing and monitoring phases of the venture capital cycle, which we treat jointly because the mechanisms deployed at investment stage as staging, contractual rights, board representation are designed precisely to shape the monitoring relationship that follows. In a traditional closed-end fund, capital is called from limited partners as the general partner identifies investment opportunities. Each year, managers screen hundreds of proposals before committing to a small number of companies. The selection problem is compounded by a fundamental misalignment of incentives: once equity is injected, the entrepreneur bears far less downside than the investor. Founders may pursue projects with high private returns but lower expected value for outside shareholders, or take on excessive risk knowing that success accrues disproportionately to them while losses fall largely on the fund. This is the classic agency problem in venture finance [20]. To manage this asymmetry, venture capitalists have developed a set of contractual and governance tools that structure both the investment and the ongoing relationship with the entrepreneurial team. The remainder of this chapter examines these mechanisms in detail, identifies structural tensions that arise from the ten-year fund model, and evaluates how permanent capital vehicles behave in these terms.

4.2 Staging the capital investment

The single most consequential tool available to venture capitalists is the ability to stage their investment over time rather than committing the full amount upfront [20]. The rationale is straightforward once one considers the statistical structure of venture returns. Data covering thirteen VC partnerships and 383 portfolio companies between 1969 and 1985 show that 34.5% of the total capital invested resulted in a partial or total loss, while a mere 6.8% of investments accounted for nearly half of the aggregate ending value [20]. More recent evidence confirms the persistence of this pattern: across hundreds of funds active between 1985 and 2014, roughly 6% of individual investments, representing 4.5% of dollars deployed, generated approximately 60% of total returns. [1] Given this extreme concentration, it would be reckless to deploy the full committed capital into a company at once. Instead, the venture capitalist structures the investment as a sequence of financing rounds, conventionally labelled in alphabetical order (Series A, Series B, and so on), each corresponding to a distinct stage of the company's development. The capital provided at each round is calibrated to be sufficient for the company to reach specific, pre-agreed results that mark progress from one stage to the next. If the results are met, the venture capitalist participates in the subsequent round; if they are not, the investor retains the option to withhold further funding and, in effect, abandon the project. This mechanism serves two purposes simultaneously. First, it protects the investor against downside: rather than discovering after ten years that a company has consumed all of its capital without producing results, the venture capitalist evaluates progress at each stage gate and cuts losses early on the ventures that are not working. Second, it disciplines the entrepreneur. A large upfront capital injection could distract the founding team from the primary objective of scaling the business, or worse, create incentives to extract personal compensation rather than reinvest in growth. By tying each new tranche to verifiable progress, staging forces the entrepreneur to demonstrate that the previous capital has been put to productive use before receiving more. Gompers, Paul A.[48] provided the first systematic empirical test of these intuitions. Analysing a sample of 794 venture-backed companies, he found that the intensity of staging varies predictably with the severity of information problems. Specifically, decreases in industry ratios

[1] Chris Dixon, "Performance Data and the 'Babe Ruth' Effect in Venture Capital," Andreessen Horowitz, 8 June 2015, https://a16z.com/performance-data-and-the-babe-ruth-effect-in-venture-capital/

of tangible assets to total assets, higher market-to-book ratios, and greater R&D intensities all lead to shorter funding duration and more frequent monitoring rounds. Early-stage companies, where uncertainty is highest and track record is thinnest, receive significantly less money per round than later-stage ones. The pattern is consistent with the theoretical prediction that staging functions as a monitoring device: the more difficult it is for the investor to observe the entrepreneur's actions and assess the venture's true prospects, the tighter the control over capital provision. From a theoretical standpoint, Hsu, Yao-Wen [49] complements the agency-cost analysis of Gompers, Paul A. [48] by formalising the staging decision as a compound option problem. Both analyses predict more intensive staging for younger ventures, though the underlying mechanisms differ: while Gompers emphasises the reduction of agency costs through frequent monitoring, Hsu highlights the value of preserving an abandonment option when the expected growth rate of the venture is low. Kaplan, Steven N. and Str"omberg, Per [50] extende this analysis by distinguishing between two forms of staging. The first, which they term *ex ante* staging, operates within a single financing round: part of the committed funding is made contingent on explicit milestones, so that the venture capitalist can withhold the remaining tranche if targets are not met. In their sample of 67 investments, the median deal committed $6.0 million, of which $4.8 million was disbursed at closing and the remainder tied to milestones. The second form, *ex post* staging, is the more familiar mechanism whereby the venture will require additional funding in a subsequent and separately negotiated round. The distinction matters because the two types respond to different risk drivers. Ex ante staging increases with internal risk that is, concerns about management quality, moral hazard, or the need for costly monitoring. Ex post staging, by contrast, tightens significantly with external risk, such as market uncertainty or competition. This suggests that the spacing between rounds is not purely an agency tool but also reflects the investor's need to preserve an abandonment option when the environment itself is uncertain. A further refinement concerns the choice between the two operational forms that staging can take. Bienz, Carsten and Hirsch, Julia [51] distinguished between *round financing*, in which each subsequent tranche is negotiated separately when the venture needs additional capital, and *milestone financing*, in which the contingencies that trigger disbursement are specified in the initial contract. The distinction is consequential because it determines the degree of contract completeness and, with it, the scope for renegotiation. Their empirical analysis of German venture capital deals confirms that the decision to stage *per se* is

driven by uncertainty and information asymmetry, consistent with Gompers, Paul A.[48]. However, the choice of *which form* of staging to adopt depends on a different set of factors: the expected distribution of bargaining power at the time of refinancing and the predictability of the development process. When development milestones can be specified credibly in advance, milestone financing protects the entrepreneur against opportunistic renegotiation by the investor. When they cannot, the parties default to round financing, which leaves the venture capitalist with greater discretion but also exposes the entrepreneur to a hold-up risk [51]. It is worth noting that the decision to discontinue a project is structurally assigned to the venture capitalist, not to the entrepreneur. Sahlman, William A. [20] observed that investors and founders "are likely to disagree on certain issues, including if and when to abandon a venture", and the entrepreneur's private benefits from running the firm create a systematic bias toward continuation. Gompers, Paul A.[48] confirms this empirically, concluding that the evidence indicates venture capitalists are concerned about "the lack of entrepreneurial incentive to terminate projects when it becomes clear that projects will fail". This is what the foundamental literature about Venture Capital tells for this mechansim. This important practice today has consolidated and it's used as a must also in PCV, showing that is a feature that goes beyond the model structure. Staging is what makes the investor's right to walk away credible and enforceable at regular intervals. However, the disciplining function of staging is not without costs. Tian, Xuan and Wang, Tracy Yue [52] show that venture capitalists who are more tolerant of early failure measured by their tendency to continue funding ventures that have not yet met milestones back firms that ultimately produce significantly more patents and patents of higher impact. Their finding suggests that staging, while effective as a monitoring device, can become counterproductive when applied too rigidly. In particular, capital constraints and career concerns can lead younger and less experienced venture capitalists to terminate promising projects prematurely, sacrificing long-run innovation for short-run portfolio discipline. As will be argued move on, the finite life of a drawdown fund may systematically contribute the cost identified by Tian, Xuan and Wang, Tracy Yue [52]: when the fund's remaining life shrinks, the pressure to terminate underperforming investments intensifies regardless of their long-run potential, turning a mechanism designed to discipline into one that penalises patience (is the case of deep tech investments).

4.3 Contractual rights and deal structure

Beyond staging, venture capitalists rely on a battery of contractual provisions to manage the agency relationship with the entrepreneur. These provisions allocate cash flow rights, voting control, and decision rights between the parties, and are designed to be highly state-contingent [53]. The result is a financing contract that effectively shifts a substantial portion of the risk in a new venture from the investor to the entrepreneur [20, p. 510].

4.3.1 The choice of security.

The dominant instrument in US venture capital is convertible preferred stock. In the sample of 213 financing rounds analysed by Kaplan and Strömberg, convertible preferred was used in 96% of cases and was the exclusive instrument in 80% [53]. The logic is intuitive: convertible preferred combines downside protection with upside participation. If the venture fails, the investor's liquidation preference ensures a senior claim on whatever residual value exists. If the venture succeeds, conversion into common stock allows the investor to share proportionally in the gains [20, 53]. This dual nature also serves a screening function: by accepting a security that subordinates their own claim in bad states, entrepreneurs signal confidence in their business plan [53]. A common variant is *participating* convertible preferred, which allows the holder to receive both the liquidation preference and a pro rata share of the remaining value, effectively "double-dipping" [**?**]. Kaplan and Strömberg report that participating preferred was used in 48% of their sample [53].

4.3.2 Cash flow protections.

Several provisions protect the investor's economic interest. *Anti-dilution clauses*, present in approximately 95% of financings, adjust the conversion price of existing preferred stock downward if a subsequent round is raised at a lower valuation. The most common form is the weighted average ratchet, used in 78% of cases, which recalculates the conversion price as a function of the number of new shares issued and their price [53]. *Liquidation preferences* ensure that the venture capitalist's claims in liquidation are at least as large as the original investment in over 98% of cases. In 44% of financings, cumulative preferred dividends further strengthen

these claims [53]. *Redemption rights*, present in 79% of financings with an average vesting period of five years, give the investor the right to sell the security back to the company, providing an exit mechanism even when no liquidity event has occurred [53].

4.3.3 Control and governance rights.

Venture capitalists insist on board representation as a condition of financing. In 25% of the Kaplan and Strömberg sample, the venture capitalist controlled a majority of board seats, and in 18% of cases board composition was state-contingent, meaning that VC control increased if the company failed to meet performance benchmarks [53]. Beyond the board, contracts typically include vesting schedules for founder shares, buy-back provisions that allow the company to repurchase equity from departing employees, and information rights obliging the company to provide regular financial statements and budgets [20]. *Automatic conversion* provisions, present in 95% of financings, force preferred stock to convert into common upon a qualifying IPO, typically above a minimum share price or market capitalisation threshold. This prevents individual preferred holders from blocking a value-creating public offering in order to preserve their liquidation rights [53, 54].

4.3.4 Preemptive rights and the option to follow success.

Just as staging preserves the right to abandon, preemptive rights and rights of first refusal preserve the option to invest more if the venture proves successful. These provisions give existing investors the right to participate in subsequent financing rounds in proportion to their holdings, ensuring that they are not diluted by new investors entering at a higher valuation [20, p. 505].

4.3.5 The evolution of contract terms.

While the architecture described above has remained broadly stable, the specific terms have shifted considerably over the past two decades. Gornall, Will and Strebulaev, Ilya A. [54] compile data from multiple sources to track how the frequency of individual provisions has changed since the original Kaplan and Strömberg sample.

The trend is toward less aggressive investor protections on the cash flow side: the fraction of deals with a liquidation multiple above one fell from 71% in the early 2000s to between 1% and 6% in more recent samples; participating preferred declined from 41% to between 4% and 11%; and redemption rights dropped from 79% to between 6% and 12% [54]. At the same time, automatic conversion provisions have become near-universal, rising from 95% to 100%. The pattern suggests a market that has moved toward simpler, more standardised contract structures, with investors relying less on complex cash flow protections and more on governance mechanisms such as board control and protective provisions.

4.3.6 International and experience-based variation.

The contract architecture described above reflects primarily US practice. European venture capital operates with greater variation. Cumming, Douglas [55] analyses 223 investments across 11 continental European countries and finds that stronger VC control rights, particularly drag-along rights and the right to replace the founding CEO, are associated with a roughly 30% higher probability that the company will exit through an acquisition rather than an IPO. Conversely, the use of common equity, which implies weaker investor control, is associated with a 12% higher probability of an IPO exit. These findings underscore that the allocation of control rights in the initial contract has material consequences for how the investment is eventually realised. Contract design also varies with the experience and capabilities of the venture capitalist. Bengtsson, Ola and Sensoy, Berk A. [56] show that more experienced VCs, who possess superior monitoring abilities and more frequently join the boards of their portfolio companies, obtain weaker downside-protecting cash flow rights than less experienced investors. The authors interpret this as evidence that high-quality VCs substitute governance involvement for contractual protection: rather than loading the contract with liquidation preferences and anti-dilution clauses, they negotiate board seats and rely on their ability to intervene directly when needed [56].

4.4 Syndication

Venture capitalists rarely invest alone. Instead of committing to a deal as sole financier, they routinely form syndicates in which two or more investors jointly

take equity stakes in the same company. The practice is pervasive: depending on the market and the period, somewhere between 40% and 80% of all venture capital investments involve more than one investor [57]. Syndication has attracted considerable scholarly attention, with nearly sixty published studies examining its causes, mechanisms, and consequences [57]. The literature identifies three broad families of motivations for syndication, operating at different levels of analysis [57]. The first is financial: by sharing the equity commitment with other investors, a venture capitalist can limit the exposure of any single fund to one deal, thereby improving the diversification of its portfolio [57]. The second is resource-based: different investors bring different skills, industry knowledge, and networks, and pooling these resources can improve both the initial selection of the investment and the subsequent support provided to the company [57]. The third is strategic: by inviting another firm into a deal today, a venture capitalist creates an expectation of reciprocity, gaining access to that firm's proprietary deal flow in the future [57]. The relative importance of these motivations is not uniform. Manigart [58] survey 317 venture capital managers across six European countries and find that financial considerations dominate: on a five-point scale, respondents rate financial motives at 3.67 on average, compared to 2.47 for resource-based motives and 2.71 for deal flow considerations (Table 6). This contrasts with findings from North America, where value-adding and selection rationales tend to be more prominent [58]. A particularly interesting result emerges when the analysis distinguishes between lead and non-lead investors. Lead investors, who initiate the syndicate and typically commit the largest amount of capital, syndicate primarily to manage financial risk: they actively seek partners for deals that exceed their individual risk appetite. Non-lead investors, by contrast, join syndicates primarily for resource-based reasons. Partnering with a respected lead investor enhances their legitimacy, exposes their managers to new sectors and practices, and provides access to deals they would not have sourced on their own [58]. The distinction matters because it reveals that syndication is not a single phenomenon with a single rationale but a practice that serves different purposes for different participants within the same deal. Massimiliano Magrini, of United Ventures, was asked directly which of these motives predominates in his firm's syndication decisions. His answer touched on the full set identified in the literature— financial risk management, access to complementary skills, and reciprocity in future deal flow—without singling one out as dominant.[2] He added

[2]Interview with Massimiliano Magrini, Co-founder and Managing Partner, United Ventures.

an important qualification, however: United Ventures operates principally as a lead investor by strategy, and does not need to syndicate in order to share financial risk or co-invest; the firm is equally comfortable investing alone. This nuances Manigart's finding that lead investors syndicate chiefly for financial risk management: for at least one experienced lead investor, syndication appears to be a matter of choice rather than necessity, driven more by the other rationales than by an inability to bear the full risk of a deal alone.

4.4.1 Syndication networks and performance.

The repeated interaction of venture capitalists through co-investment has given rise to extensive networks whose structure has measurable consequences for performance [59]. Using data on 3,469 US venture capital funds started between 1980 and 1999, Hochberg, Yael V. and Ljungqvist, Alexander and Lu, Yang [59] construct syndication networks from the co-investment histories of VC firms and measure each firm's position using five centrality metrics. The results are striking. A one-standard-deviation increase in a lead investor's *eigenvector* centrality, which captures how well connected the investor is to other well-connected investors, raises the portfolio company's probability of surviving to the next funding round from 66.8% to 72.4% in the first round, from 77.7% to 82.8% in the second, and from 79.2% to 86.6% in the third (Section IV.A, p. 279). At the fund level, a one-standard-deviation increase in network centrality is associated with approximately 2.5 percentage points of additional IRR from a sample average of 15% [59, p. 253]. Crucially, the authors show that this effect is not merely a reflection of privileged access to better deals. Well-networked venture capitalists appear to perform better because they are able to draw on their network to provide superior value-added services to the companies they finance [59, Section IV.A]. Tian, Xuan[60] provides complementary evidence on the value-creation channel: startups financed by syndicates, compared to those backed by a single investor, produce significantly more patents and patents of higher impact. The mechanism is the combination of heterogeneous resources that individual investors could not provide in isolation. [61] confirms that the network effects documented by Hochberg and colleagues translate directly into investment activity: using panel data on 3,797 venture capital firms and 1,290 business angels, he finds that investors with more direct contacts exhibit higher deal flow, with a one-standard-deviation increase in strong ties raising the probability of investing by 9.4% [61,

Section 5]. The dominant channel is access to complementary resources rather than deal sourcing or reciprocity [61].

Taken together, these findings establish syndication not as a peripheral feature of venture capital but as a core mechanism through which the industry generates value. The ability to form, maintain, and leverage co-investment relationships is a determinant of both the quality of investment decisions and the level of support available to portfolio companies after the investment is made.

4.5 Monitoring and post-investment involvement

The venture capitalist's role does not end once the investment is made and the contract is signed. On the contrary, a substantial body of evidence indicates that the ongoing relationship between the investor and the entrepreneurial team is one of the most consequential features of venture capital as an asset class. Unlike a passive shareholder, the venture capitalist takes an active part in the governance and development of the portfolio company, dedicating considerable time and resources to what the literature broadly terms "monitoring and value-added" activities [62, 63].

4.5.1 The scope of involvement.

Sapienza, Harry J. and Manigart, Sophie and Vermeir, Wim [62] provide the most detailed cross-country picture of how venture capitalists interact with the companies they finance. Surveying VC firms and their portfolio company CEOs in the United States, the United Kingdom, France, and the Netherlands, they document three categories of roles. The first and most valued is *strategic*: the venture capitalist serves as a sounding board for management decisions, provides financial and business advice, and reviews plans and targets [62, Table 3]. The second is *interpersonal*: the investor acts as a mentor, coach, and confidant to the CEO, particularly in the early stages of the company's development when the founding team faces decisions it has never encountered before. The third is *networking*: the venture capitalist facilitates introductions to potential customers, suppliers, strategic partners, and professional service providers, and assists in recruiting key members of the management team. Across all four countries, VCs rate their strategic involvement as the most important contribution they make, with interpersonal roles coming second and networking

assistance last [62, Table 3]. The intensity of this involvement varies considerably across markets. American and British venture capitalists devote roughly 194 and 154 hours per year, respectively, to each portfolio company, meeting with the CEO slightly more than once a month on average. Their French and Dutch counterparts are substantially less involved, dedicating 79 and 75 hours per year respectively [62, Table 3]. The finding that effort and perceived value-added tend to move together suggests that active involvement is not merely a cost of doing business but a channel through which the investor generates returns.

4.5.2 Professionalization.

One of the most tangible forms of value-added is the venture capitalist's role in accelerating the professionalisation of the startup's internal organisation. Hellmann, Thomas and Puri, Manju [64] study a sample of Silicon Valley startups and find that companies backed by venture capital adopt stock option plans more than twice as fast as comparable firms without VC backing, and hire their first vice president of sales and marketing significantly earlier. More broadly, entrepreneurs report that their venture capital investors are influential in shaping human resource policies, with a 48% higher likelihood of the investor playing a formative role in this area [64, Tables 3–5]. These results are consistent with the view that venture capitalists do not simply monitor from a distance but actively intervene in the organisation of the firm, pushing it toward the structures and processes it will need to scale.

4.5.3 Value-added through the network.

Beyond direct involvement in governance and operations, venture capitalists create value by connecting portfolio companies to their broader network of relationships. As documented in the previous section, syndication generates co-investment ties that span the industry. Lindsey, Laura [65] shows that these ties have a concrete economic payoff: companies that share a common venture capital investor are significantly more likely to form strategic alliances with one another than companies that do not. The effect is strongest for alliances involving the transfer of technology, suggesting that venture capitalists facilitate connections that would be difficult for the companies to establish on their own.

4.5.4 The distinction between advice and conflict management.

Cumming, Douglas and Johan, Sofia A.[63] draw a useful analytical distinction between two dimensions of the post-investment relationship. The first is *advice*: the proactive provision of strategic guidance, industry knowledge, and operational support. The second is *conflict management*: the reactive exercise of control rights when the interests of the investor and the entrepreneur diverge. Analysing a sample of European VC funds, they find that both formal governance mechanisms (contractual cash-flow and control rights) and informal ones (trust, reputation, and management structures) significantly facilitate the provision of advice. Conflicts, by contrast, are more closely tied to the legal environment: higher-quality legal systems in the entrepreneur's country of residence mitigate the frequency and severity of VC-entrepreneur disputes. The finding underscores that monitoring in venture capital is not a single activity but a composite of support and oversight, each governed by partly different institutional factors [63, Section 1].

4.5.5 Peer monitoring within syndicates.

Finally, the monitoring function extends beyond the bilateral VC-entrepreneur relationship. Bayar, Onur and Chemmanur, Thomas J. and Tian, Xuan [66] show that venture capitalists who co-invest in the same company monitor not only the startup but also each other. Syndicates composed of investors with comparable levels of experience produce better outcomes, because each member is credibly able to evaluate the behaviour and contribution of the others. This peer monitoring mechanism reinforces the governance structure established through staging and contractual rights, and provides a further rationale for why syndication is so pervasive in the industry.

4.6 The rise of deep tech

For much of the past two decades, venture capital portfolios were heavily concentrated in digital businesses with comparatively short development cycles: business-to-business software, SaaS platforms, and digital marketplaces. The Covid-19 pandemic marked an inflection point. Ambrois etAl [67] document, through a global survey of

VC firms, how the crisis led general partners to reassess their investment practices, with unobservable dimensions such as screening criteria and sector focus proving more susceptible to change than structural, LP-visible commitments. In the post-pandemic landscape, investment attention has shifted toward frontier sectors such as synthetic biology, advanced materials, and quantum computing.[3] Deep tech investments quadrupled globally between 2016 and 2020, from approximately $15 billion to more than $60 billion, yet roughly two thirds went to artificial intelligence and synthetic biology alone, leaving the vast majority of deep tech sectors significantly underfunded.[4] More recently, McKinsey estimates that European deep tech funds have delivered a weighted average net IRR of 16%, compared to 10% for traditional technology funds, based on a sample of 115 deep tech and 1,572 traditional funds with vintages between 2003 and 2020.[5]

4.6.1 What defines deep tech.

A recent systematic review by Kortsch etAl [68] proposes six defining attributes of deep tech innovation at the invention level: it originates from scientific breakthroughs, is embodied in tangible physical systems, has enabling platform potential across multiple application domains, requires extended development timelines, demands high capital intensity, and faces pervasive, multi-layered uncertainty. Development horizons from initial scientific discovery to market readiness typically range from five to fifteen years, with long pre-revenue phases during which the venture consumes capital without generating returns [68]. Founders interviewed in qualitative studies describe a "temporal mismatch" between these cycles and the expectations of venture capital investors, whose fund structures typically require portfolio liquidity within a much shorter timeframe [68]. Romasanta etAl [69] complement this picture by observing that traditional VC models rely on short investment cycles and have historically shown limited interest in deep tech companies, in part because they "lack technical knowledge, which hinders their ability to evaluate either the technology or the company" [69, p. 21].

[3] See United Ventures, *6 Trend Tecnologici per il Post-Pandemia*, March 2021, https://unitedventures.com/2021/03/19/6-trend-tecnologici-per-il-post-pandemia/.

[4] BCG and Hello Tomorrow, *The Deep Tech Investment Paradox* (2021), Exhibits 1 and 4.

[5] McKinsey, *European Deep Tech: What Investors and Corporations Need to Know*, December 2024, Exhibit 2. Data from Preqin.

4.6.2 Implications for the analysis.

The core investment mechanisms described in the preceding sections of this chapter (staging, contractual rights, syndication, monitoring) remain applicable to deep tech ventures. What has changed is the context in which they operate: development cycles that extend well beyond typical fund horizons, risk that is scientific rather than primarily commercial, and a need for technical competence that goes beyond traditional financial and managerial expertise. Understanding this shift is necessary to appreciate both the structural gap that separates scientific research from industrial application and the different ways in which fund structures can respond to it.

4.6.3 The technology transfer gap

The gap between scientific research and industrial application is not, strictly speaking, a drawdown fund inefficiency. It is a broader structural problem in the innovation ecosystem. Research-based inventions (RBIs) developed within universities frequently remain trapped in what the literature calls the "Valley of Death": the transition period between original research and commercial application where conventional funding sources are unwilling to bear the risk [70]. Ughetto, Elisa and Battaglia, Daniele and Micol, Federico [70] study 15 research projects with commercialisation potential at Politecnico di Torino and identify three classes of inhibitors: institutional (lack of resources, misaligned academic career incentives), interpersonal (disconnection between researchers and potential industrial partners), and cultural (a focus on theoretical content over applied development). Over 75% of university patents are never commercialised, according to the estimates they cite [70].

One institutional response has been the creation of Proof-of-Concept programmes within universities. Ughetto, Elisa and Battaglia, Daniele and Micol, Federico [70] analyse 12 such programmes at the same university and find that they generate three classes of interconnected enablers (relational, structural, and cultural) that facilitate the commercialisation of embryonic research-based inventions.

At the market level, new investor types have emerged to operate at the boundary between academic research and venture finance. Ughetto, Elisa and Battaglia, Daniele and Micol, Federico [71] study these "IP Venturing Funds" through interviews with seven such funds operating across Europe and the United States. They find that IP

Venturing Funds differ from classic venture capital investors in two respects: they invest small amounts of capital in startups with technologies that are still unproven, effectively providing proof-of-concept financing "with strings attached," and they are willing to invest in projects where the entrepreneurial drive typical of a founding team may not yet be present [71, Abstract]. All seven funds in their sample share a strongly "hands-on" approach, though they adopt different operational models and decision-making processes to reduce the several sources of uncertainty (legal, technological, managerial, market, and value-chain) that characterise the commercialisation of research-based inventions [71]

4.6.4 Who does not contribute to bridging the gap.

Traditional drawdown funds face structural limitations on all three dimensions that deep tech requires. On duration, the ten-year fund life (with typical extensions of one to two years) imposes a deadline that may arrive well before the technology has reached commercial maturity. On competence, venture capital teams are typically composed of professionals with financial and commercial backgrounds; as BCG and Hello Tomorrow put it, VC funds are "structurally unfit (lifetime, size, incentives) to invest in deep tech" and often "lack the expertise needed to understand advanced science, engineering risks and to support ventures."[6] McKinsey corroborates this from the founder side: "deep tech founders can be frustrated with the lack of detailed expertise that generalist VCs bring to investment and development conversations."[7] On incentives, the management fee structure, typically calculated on committed capital, may not encourage the patient engagement that deep tech ventures require; this dimension will be examined in greater detail in the sections that follow, with evidence from practitioner interviews.

4.6.5 Who can contribute.

Permanent capital vehicles, by contrast, are not subject to the ten-year fund life constraint. Their capital base allows for longer holding periods, and their governance structures can accommodate the multi-disciplinary teams that deep tech investment

[6] BCG and Hello Tomorrow (2021), Executive Summary, p. 4.

[7] McKinsey, *European Deep Tech: What Investors and Corporations Need to Know*, December 2024, p. 11.

demands. LIFTT S.p.A. represents a concrete instance of this model. LIFTT describes itself as "a not-only-profit investment holding company that aims to offer an innovative and dynamic vision of venture capital applied to technology transfer."[8] LIFTT invests in innovation from the earliest stages, participating in pre-seed and seed rounds, and it will always be active investors, participating in the corporate life until the exit, for which it does not impose time obligations.[9] As of today, LIFTT counted over 215 private and institutional investors, including Azimut, LINKS Foundation, Compagnia di San Paolo, and CRT Foundation, had raised €104.7 million, and had completed 51 investments in innovative startups and SMEs from a pipeline of over 3,000 companies across all market sectors, with a focus on deep tech.

4.7 The demand side: are investors ready for patient capital?

The evidence reviewed so far explains why deep tech ventures need longer holding periods and more technical scrutiny than a traditional ten-year fund can offer. Diana Saraceni confirms this diagnosis from the supply side and adds a critical qualification on the demand side.[10]. She confirmed that deep technical competence is fundamental in life sciences venture capital and accelerates portfolio company development, consistent with Panakès's own team composition, which includes engineering backgrounds among the founders and dedicated biotech specialists. The most substantive point to emerge from the interview, however, concerns duration. Saraceni stated that she would want five additional years of fund life, specifically for deep tech and life sciences investments, where development cycles can extend to twelve or fifteen years. Crucially, she located the constraint not on the general partner's side but on the limited partner's side: even sophisticated institutional investors are not willing to remain illiquid for that long. What they need is exit events and periodic liquidity, regardless of whether a longer holding period would ultimately produce a better outcome for the underlying investment. This demand-side constraint is compounded

[8]European Investment Bank, "Italian Business Innovation Bolstered with €100 Million Initiative by EIB and Venture Capital Firm LIFTT," press release, 9 September 2024, https://www.eib.org/en/press/all/2024-323-italian-business-innovation-bolstered-with-eur100-million-initiative-by-eib-and-venture-capita

[9]LIFTT S.p.A, https://www.liftt.com/en/lift-en/

[10]Interview with Diana Saraceni, Co-founder and Managing Partner of Panakès Partners

by a second problem: the opacity of interim valuations. In life sciences, formal revaluations of a portfolio company occur only when a new financing round takes place, and in biotech this can happen every two or three years. Between rounds, the limited partner has limited visibility into what their holding is actually worth; the only reference point is the valuation set at the last round, which is by construction a stale and approximate estimate. Extending the fund's life by five years would therefore not simply extend the illiquidity period. It would extend, in parallel, the period during which investors have little more than an approximate, infrequently updated figure to gauge the value of their commitment. Taken together, these two elements help explain why the shift toward permanent capital vehicles has not been as swift or as widespread as the supply-side case built in the preceding sections might suggest. It is not only that general partners have been slow to recognise the structural mismatch between fund life and deep tech development cycles. It is that the investor base best positioned to fund such vehicles, namely professional and institutional investors, still largely prices liquidity very highly and is reluctant to accept the valuation opacity that comes with longer, less frequently marked holding periods. When asked directly whether she would change the ten-year model, Saraceni confirmed that duration was the element she would change first, but she was careful to frame the current structure as a binding constraint imposed by investor preferences rather than a design choice that general partners consider optimal. Massimiliano Magrini, of United Ventures, converges with Saraceni's diagnosis while widening its scope beyond life sciences and deep tech in general. Asked whether the case for a longer holding period was specific to deep tech, Magrini partially disagreed: in his view the need is not exclusive to deep tech but applies, more or less, to every sector, since companies now stay private considerably longer than before.[11] Yet he located the binding constraint in the same place Saraceni did: on the LP side. LPs, he noted, want liquidity sooner rather than later, but are not yet ready to buy into an evergreen fund. He traced this reluctance to the composition of the investor base itself: capital is typically raised from the same institutional investors who backed the manager's previous fund, and an investor base assembled around a 5+5 structure is difficult to redirect toward an evergreen vehicle without effectively resetting the relationship—something he judged not easy to do, since some institutional investors simply have no appetite for an evergreen fund. He extended the point to United Ventures' own growth fund, noting that it would not have been possible to propose

[11] Interview with Massimiliano Magrini, Co-founder and Managing Partner, United Ventures.

an evergreen structure to the investors who had backed that vehicle, since they would not have accepted it. By this perspective, the ten-year fund, in other words, persists not because it is a first-best structure but because it is the structure that satisfies the participation constraint of an investor base that is not yet ready for something more patient.

4.7.1 The response from permanent capital practitioners: aligned expectations and technical diligence

The picture that emerges from Saraceni's interview, namely an investor base not yet ready for extended illiquidity, was put directly to LIFTT's CFO.[12] Merlino's response reframes the problem rather than contradicting it. He stated that anyone investing in LIFTT has always been made aware, from the outset, that an investment in the company is not one from which a short or medium-term return should be expected: LIFTT has consistently presented itself as patient-capital venture investing, and an investor considering LIFTT has to accept that condition or choose not to invest at all. In his account, the two natural exit routes for LIFTT shareholders, namely dividend distributions or an eventual listing, are both medium-to-long-term propositions that were communicated as such from the start. This does not resolve the tension identified by Saraceni so much as relocate it: rather than a structural mismatch between what deep tech requires and what any investor will accept, Merlino describes it as a matching problem between the vehicle and the specific investor base it attracts. He noted explicitly that institutional investors who require a contractual, written commitment that capital will be returned by a specific year are difficult to raise from if one is running a holding company rather than a closed-end fund. LIFTT's capital base is consequently self-selected: the 107 shareholders recorded by the end of the first fundraising cycle in February 2021 include institutional entities, venture capitalists, entrepreneurs, managers, and professionals,[13] a composition that suggests LIFTT has built its investor base around parties willing to accept the patient-capital proposition rather than attempting to convert a broader, liquidity-sensitive pool of institutional capital. A related point was raised independently in

[12] Interview with Giuseppe Merlino, CFO, LIFTT S.p.A.

[13] LIFTT S.p.A., *Annual Report 2020*, p. 66 (table of subscriptions by investor category as of 28 February 2021).

the interview with LIFTT's investment team.[14] Asked how the model compares to the evergreen structure adopted by some peers, Luca Salerno, business analyst at LIFTT, argued that the absence of quarterly redemption windows, while it removes a liquidity feature that evergreen funds can offer, also removes a source of potential misalignment between what is promised to investors and what the vehicle can actually deliver. In this account, having a structure like LIFTT's "puts things in the clear from the start" and tends to align investor expectations more closely with reality than a structure that promises periodic liquidity but may struggle to honour it under stress.

4.8 Technical competence in practice.

Magrini, of United Ventures, offers a complementary perspective on how technical competence is actually built into a fund's practice. Asked whether deep tech investing requires dedicated technical expertise inside the fund itself, he described the question as one without a single answer: in his account, each manager works out its own solution to the problem, with some funds carrying technological competence in-house and others bringing it in on a deal-by-deal basis as needed.[15] What matters, in his view, is less the specific organisational choice than two underlying requirements: a genuine vision of the technology in question, and investors capable of doing their job well enough to support the entrepreneurs they back. Beyond these two conditions, he was explicit that practice varies considerably from firm to firm, each having arrived at its own approach in a different way. LIFTT represents one particular resolution of this heterogeneity: The Annual Report 2020 provides independent, document-based confirmation of the technical diligence that Saraceni identified as a prerequisite for deep tech investing. LIFTT's internal valuation process is explicitly described as extending well beyond financial indicators: the periodic review of each holding is required to give "extreme importance" to information concerning the technical and scientific content of the underlying project, alongside its technical, commercial, and financial development plan.[16] At the governance level, every investment proposal must pass through an Investment Committee composed exclusively of members external to the operating team, selected specifically for

[14] Interview with Luca Salerno, Business Analyst Team, LIFTT S.p.A.

[15] Interview with Massimiliano Magrini, Co-founder and Managing Partner, United Ventures.

[16] LIFTT S.p.A., *Annual Report 2020*, p. 144 (section on valuation criteria for balance sheet and management purposes).

their background in entrepreneurship, innovation, and technological and scientific development, before it can reach the Board of Directors for final approval.[17] This structure operationalises, at the level of internal governance, the point that Saraceni made about the necessity of technical expertise in evaluating long-cycle, science-based ventures.

4.8.1 From screening to monitoring: evidence from LIFTT's investment process

The interview with Luca Salerno of LIFTT's investment team offers a detailed, phase-by-phase account of how the absence of a fixed fund life reshapes each of the mechanisms examined in the preceding sections of this chapter.[18]

Screening. The interviewee argued that the principal advantage of LIFTT's structure at the screening stage is the ability to accept projects with substantially longer development horizons than a closed-end fund could accommodate, precisely because there is no risk of failing to return capital to investors within a fixed window. A further point, not previously identified in the academic literature reviewed for this chapter, concerns the stability of selection criteria over time. In a traditional ten-year fund, which typically has only four or five years in which to deploy capital, selection criteria tend to shift as the investment period progresses: early years permit longer-cycle, higher risk bets, while the final years of the investment period require assets with much shorter maturation windows, generally no more than five or six years, in order to allow enough time for an exit before the fund must wind down. LIFTT, lacking this rigid temporal structure, applies the same early-stage selection criteria regardless of when in its operating history an opportunity is evaluated. The internal process works as this: proposals move through a five-stage funnel (Lead, Qualified, Due Diligence, Negotiation, Portfolio) with an average total duration of two months, and every qualified proposal must clear both an internal Strategic Committee and an Investment Committee composed of six external experts before reaching the Board for final approval.[19]

[17] LIFTT S.p.A., *Annual Report 2020*, p. 170 (section on operational risk mitigation).
[18] Interview with Luca Salerno, Business Analyst Team, LIFTT S.p.A.
[19] LIFTT S.p.A., *Annual Report 2020*, pp. 46–48.

Staging. On the question of whether investment tranches are tied to milestones, the interviewee confirmed that LIFTT defines a development roadmap with specific milestones at the time of each investment, though whether these milestones are formally linked to the release of further tranches varies from deal to deal. He was explicit, however, that this practice is not a structural consequence of LIFTT's permanent capital model but rather an operating convention the firm has chosen to follow: staging, in his words, is largely "invariant" across fund structures and depends more on the type of investment and sector than on whether the vehicle has a fixed life. What does change under a permanent capital structure is the substance of the milestones themselves. Because LIFTT is not bound to a defined investment period, it can define success milestones oriented toward well-specified inflection points in the venture's development, rather than milestones whose primary function is simply to confirm that the company has survived to the next round. On exit discipline, the interviewee described minimum performance thresholds that a portfolio company is expected to approach before LIFTT will consider an exit, except in genuinely opportunistic cases where an exit represents the only viable path forward. Capital for follow-on rounds is concentrated on the assets that are performing best relative to the milestones agreed at the time of the original investment; whether underperformance relative to plan justifies discontinuing support is assessed deal by deal, since it matters whether the shortfall stems from poor execution or from external factors outside the company's control.

Contractual instruments. Asked directly about the type of security used, the interviewee stated that LIFTT relies exclusively on direct equity rather than convertible preferred stock, describing it as the financial instrument the firm prefers. Governance rights, such as board seats, vesting, and anti-dilution provisions, are described as highly deal-specific: LIFTT seeks governance rights particularly at the early stage, both to add value and to verify that the agreed development plan is being pursued, but beyond this there is no standard template that LIFTT applies uniformly across deals.

Syndication. It was identified a specific source of friction that arises from LIFTT's longer investment horizon when forming a syndicate. Co-investors who are approaching the end of their own fund's investment period may hold a different view than LIFTT on the appropriate timing of an exit for a shared asset, since LIFTT has

the option of waiting for further value creation while a fund nearing the end of its life may not. The interviewee was careful to frame this primarily as a problem for the co-investor rather than for LIFTT, and noted that its practical significance depends heavily on which party leads the round: when another investor with a shorter horizon is the lead, the potential misalignment can matter for the whole syndicate; when LIFTT itself is the lead investor, it is LIFTT that sets the strategic direction of the round, and the burden of adjusting shifts to the other participants.

Monitoring. When I proposed that LIFTT's post-investment support looked more like technical de-risking through technology transfer than commercially oriented support aimed at preparing an exit, the interviewee agreed with this framing and elaborated on it: the firm's focus, in his account, is on value creation rather than on the urgency of reaching an exit, precisely because LIFTT is not forced by a fund termination date to exit before the company has expressed its maximum value. This is corroborated by the operational description in the Annual Report, where portfolio support is explicitly organised around product, business, and human-resources roadmaps tied to pre-agreed milestones and KPIs, with the intensity of LIFTT's project management involvement calibrated to the maturity of each portfolio company.[20] Support of this kind can also be remunerated through a work-for-equity arrangement when it reaches a substantial scale, rather than being priced as a cash fee.[21]

4.9 Compensation, fund size, and the "too much money" problem

The literature reviewed earlier in this chapter established the mechanics of general partner compensation: a management fee, typically calculated on committed capital, and carried interest on realised profits. Litvak, Kate [21] provides the most detailed empirical account of how this structure actually operates, based on a hand-collected dataset of sixty-eight venture partnership agreements. She shows that VC compensation in fact has three components rather than two: alongside the management fee and carried interest, the rules governing the timing of distributions during the fund's

[20] LIFTT S.p.A., *Annual Report 2020*, p. 49.
[21] LIFTT S.p.A., *Annual Report 2020*, p. 171.

life constitute a third element, since they frequently generate what amounts to an interest-free loan from limited partners to general partners [21, Section II]. Critically, compensation is less performance-based than commonly assumed: for funds of vintage years between 1986 and 1997, roughly half of total GP compensation derives from the nonrisky management fee rather than from carried interest [21, p. 161]. The management fee itself is far from standardised. Litvak identifies nine distinct formulas across her sample, all combinations of two choices: the base (committed capital or managed capital, meaning the cost basis of invested and undistributed assets) and whether the applicable percentage is fixed or varies over the fund's life [21, Section II.A]. The net present value of the management fee, computed over the fund's typical ten to eleven year life as a percentage of committed capital, ranges from 3.66% to 20.15% across her sample, with a mean of 14.14% [21, Table 1], indicating that the textbook "2% and 20%" rule of thumb substantially understates the true magnitude of fixed compensation in many funds. The choice of base carries real incentive consequences. When the management fee is calculated on managed rather than committed capital, the general partner has a direct incentive to accelerate investment and delay distributions, since doing so keeps the fee base larger for longer [21, Section II.A]. More broadly, because the fee is pegged to committed capital rather than to the fair value of the portfolio, and because it constitutes the majority of GP pay, the structure rewards the act of raising capital more directly than it rewards the creation of value in the underlying companies.

4.9.1 The consequences of raising too much.

An incentive to raise large amounts of capital does not necessarily translate into a benign outcome for the fund. Gompers, Paul and Lerner, Josh [30] show that inflows of capital into venture funds are associated with higher valuations on the funds' new investments, an effect that is most pronounced in states with the highest concentration of venture capital activity, consistent with competition for a limited number of attractive deals [30]. Crucially, they find that the success rate of investments made during periods of high inflows and elevated valuations, measured through IPO completion or acquisition at favourable terms, is not significantly different from that of investments made in colder markets [30, Section 5.3]. The implication is that abundant capital inflates prices more than it improves the underlying quality of the projects financed. A related mechanism operates on the investment side of

the same fund: Gompers, Paul A. [48] shows that a one standard deviation increase in a fund’s capital commitments is associated with a two-month reduction in the average time between successive financing rounds, suggesting that larger pools of committed capital push general partners toward faster, and potentially less disciplined, deployment.

4.9.2 A counterpoint from practice.

When this dynamic was put directly to Diana Saraceni, in the context of Panakès’s own scaling from a EUR 75 million first fund to a EUR 175 million second fund, she did not describe fund growth as unconstrained. Target fund sizes are set jointly with anchor investors, and the conditions attached to that relationship prevent disproportionate scaling; in her account, track record and investor discipline regulate how much a fund is allowed to grow between vintages.[22] This suggests that the “too much money chasing too few deals” dynamic documented by Gompers and Lerner is not left entirely to the general partner’s discretion: the hard cap negotiated with anchor investors, and the reputational stakes tied to track record, act as an external check on the incentive to over-raise. This does not contradict the academic finding so much as identify a governance mechanism that can attenuate it in practice, at least for funds with a sufficiently engaged and disciplined investor base.

4.9.3 Alternatives to traditional compensation

Absent a fund structure, LIFTT cannot replicate a traditional carried interest arrangement. According to Merlino, the firm’s employees are instead incentivised through a periodic stock option plan, calibrated to role and seniority, which converts into equity in the company over time. He described this as serving both a retention function and, more importantly, an incentive to maximise the value of the company itself, since the options are only valuable if LIFTT’s equity appreciates.[23] Odasso confirmed the same underlying logic from the business analysis side: compensation at LIFTT consists of salary, bonus, and participation in the holding structure through a partner arrangement, and no individual has a personal fee-based incentive to push a given investment through the approval process, since everyone is aligned to the same

[22] Interview with Diana Saraceni, Co-founder and Managing Partner of Panakès Partners
[23] Interview with Giuseppe Merlino, CFO, LIFTT S.p.A.

overall corporate objectives.[24] Both accounts describe a compensation structure tied to the value of the company as a whole rather than to the volume of capital raised or deployed, which is structurally closer to a NAV-based logic (what happen in an evergreen) than to a committed-capital management fee. It is worth being explicit that this comparison should not be read as establishing that a NAV-based or equity-based compensation structure is unambiguously superior. Carried interest, whatever its shortcomings, does tie a portion of GP compensation to a realised outcome, the exit, which creates a form of discipline that a pure NAV fee does not automatically replicate. Odasso acknowledged this tension directly when the point was raised, agreeing that the two logics differ but maintaining that both are ultimately meant to align compensation with asset performance, whether through exit-contingent carry or through the value implicit in unrealised NAV.[25] One could further argue that because carried interest is realised only at exit, it may create an incentive to optimise for the specific financial metrics that are visible at that moment, such as the reported multiple or IRR, rather than for the durable value of the underlying business, in a way that a fee tied to an independently assessed NAV would not.

4.10 The zombie holding question

The absence of a fixed termination date raises an obvious concern: without the discipline of a fund life that forces liquidation, could a permanent capital vehicle end up holding underperforming positions indefinitely, simply because there is no external deadline compelling a decision?

4.10.1 The practitioner response.

Odasso framed the issue as one of internal fund discipline rather than of the absence of a termination date. In her account, capital allocation at LIFTT is governed by expected return: when an asset's expected return deteriorates, the corresponding capital is redirected toward other projects, and continuing to allocate capital to an underperforming position would be, in her words, an illogical decision that internal

[24] Interview with Cristina Odasso, Head of Business Analysis, LIFTT S.p.A.
[25] Interview with Cristina Odasso, Head of Business Analysis, LIFTT S.p.A.

decision-making bodies and process filters are designed to prevent.[26] She noted that both new investment proposals and follow-on proposals for existing portfolio companies must pass through multiple committees, which she described as ensuring that positions are considered from a portfolio-wide perspective, with several decision filters intended to guarantee impartiality and equity across positions.[27] This is consistent with the exit discipline described independently by LIFTT's investment team in the context of staging (see above), where minimum performance thresholds and milestone-based criteria determine whether capital continues to flow to a given asset, and where capital is described as being concentrated on the assets performing best relative to plan. It is worth being explicit about the limits of this evidence. Both accounts are self-reported descriptions of internal governance by the organisation under study, not independently verified outcomes. The existence of committees and stated allocation rules demonstrates that LIFTT has designed mechanisms intended to prevent indefinite holding, but it does not, on its own, demonstrate that these mechanisms are effective in practice, nor does it rule out the behavioural pull toward inertia that the question was originally meant to probe. A closed-end fund's discipline on this point is structural and externally enforced by the looming termination date; LIFTT's discipline, by contrast, rests on internal governance and self- assessed expected returns, which is a materially different kind of safeguard.

4.10.2 A complication from the drawdown side.

The broader academic and industry evidence, however, complicates the framing of zombie holdings as a risk specific to, or newly introduced by, permanent capital vehicles. If anything, the phenomenon is currently most visible on the drawdown side of the industry. Facing an environment of higher interest rates since 2022 and a corresponding collapse in exit opportunities, traditional private equity sponsors have increasingly found themselves unable to sell portfolio companies at acceptable valuations, leading to a documented rise in so-called "zombie companies" trapped inside closed-end structures that were explicitly designed to force liquidation.[28] The industry's principal response has been the continuation fund, a mechanism

[26] Interview with Cristina Odasso, Head of Business Analysis, LIFTT S.p.A.

[27] Interview with Cristina Odasso, Head of Business Analysis, LIFTT S.p.A.

[28] See the discussion and press citations collected in Clayton, William W. and de Fontenay, Elisabeth [38], Section II.A.2, including the observation that these pressures have "intensified since 2022, when central banks began rapidly raising interest rates."

through which a sponsor transfers portfolio companies from an existing, expiring fund into a new vehicle that the same sponsor also manages, effectively extending the holding period while offering only partial liquidity to investors.[29] By 2025, roughly one fifth of all private equity sales involved exactly this kind of related- party transaction.[30] Because the sponsor acts simultaneously as seller, buyer, and manager, and sets the timing and valuation of the transaction without an external market check, these structures raise governance concerns considerably more serious than the ones typically attributed to permanent capital vehicles, precisely because they attempt to engineer continued holding while formally preserving the appearance of a fund termination.[31] The comparison suggests that the relevant question is not whether a vehicle has a fixed termination date, but what governs the decision to continue holding a position when that date approaches or when no such date exists at all. A closed-end fund without a genuine exit option does not automatically avoid indefinite holding; it merely relocates the decision to a related-party transaction that carries its own, arguably more severe, conflicts of interest. A permanent capital vehicle without a termination date does not automatically fall into indefinite holding either, provided its internal governance genuinely ties capital allocation to updated expected returns, as LIFTT describes. Both outcomes depend on the quality of governance rather than on the presence or absence of a calendar deadline as such.

4.11 Chapter conclusion

This chapter examined how venture capital finds, structures, and supports its investments, and how that process changes once the vehicle is no longer bound to a ten-year, closed-end structure.

The traditional model rests on a coherent set of mechanisms designed to manage the agency problem between investor and entrepreneur: staged capital infusion, contractual rights that allocate risk and control asymmetrically, syndication that pools risk and resources across investors, and active monitoring that extends well beyond the provision of capital. These mechanisms are well documented and, on the whole, robust: the evidence reviewed in this chapter, from Sahlman's founding account to

[29] Clayton, William W. and de Fontenay, Elisabeth[38], Section II.A.2.
[30] Clayton, William W. and de Fontenay, Elisabeth[38], Section II.A.2, citing industry data.
[31] Clayton, William W. and de Fontenay, Elisabeth[38], Section II.B.1.

Gompers's empirical test of staging, to Kaplan and Strömberg's contractual data, to the more recent work on syndication networks and post-investment value-added, does not suggest that the mechanisms themselves are broken. What the chapter's later sections suggest is that these mechanisms operate within a temporal constraint, the fund's finite life, that was not itself designed for every type of venture the industry increasingly needs to finance.

The deep tech evidence made this constraint concrete. Development cycles that extend to ten or fifteen years, the technical competence required to evaluate scientific rather than purely commercial risk, and the persistent gap between university research and commercial application (the Valley of Death) together describe a category of investment opportunity that the traditional fund structure handles poorly, not because its governance tools fail, but because its clock runs out first. The interviews conducted for this thesis, particularly with Diana Saraceni, confirmed this diagnosis from inside the industry, while adding a qualification the literature alone would not have surfaced: the constraint binds primarily on the demand side. It is not that general partners fail to see the mismatch; it is that the investor base able to fund a longer- duration vehicle has not yet shown itself willing to accept the accompanying illiquidity and valuation opacity. The rise of permanent capital vehicles is therefore better understood as the gradual formation of a self-selected investor base around patience, as Merlino described it for LIFTT, than as a wholesale shift in what all investors are prepared to accept.

The evidence from LIFTT showed that the practical consequences of removing the fund-life constraint are real but narrower than a simple duration argument would predict. Selection criteria remain stable over time rather than shifting as an investment period nears its end; milestones are designed around genuine value-creation inflection points rather than mere survival; exits can be timed around peak value rather than forced liquidation; and post-investment support can prioritise value creation over the urgency of reaching an exit. At the same time, LIFTT's own investment team was explicit that not every operational feature observed at the firm follows mechanically from its permanent capital structure: staging practice itself, for instance, was described as a matter of firm convention rather than a structural necessity. This distinction matters for the argument of this thesis as a whole: permanent capital changes what a vehicle *can* do, not everything that it *does*.

On compensation, the chapter found a genuine misalignment in the traditional model, documented by Litvak and consistent with the pattern of capital inflows chasing a limited number of deals documented by Gompers and Lerner, but also found that this misalignment is not left entirely unchecked in practice. Anchor investor discipline and hard cap negotiations, as described by Saraceni, and the equity-based incentive structures used by LIFTT, as described by both Merlino and Odasso, offer partial, governance-based correctives that operate independently of whether a vehicle has a fixed fund life. The comparison between carried interest and NAV-based or equity-based compensation does not yield an unambiguous winner: each ties compensation to a different, imperfect proxy for value creation, and the hypothesis that carry might additionally bias General Partners toward exit-metric optimisation remains, on the evidence gathered here, a plausible but undocumented conjecture rather than an established finding.

Finally, the zombie holding question illustrated a broader methodological point that will recur in the remainder of this thesis. A risk that appears, on first inspection, to be a distinctive vulnerability of permanent capital vehicles turns out, on closer examination of the traditional model's own recent behaviour under stress, to be at least as serious on the drawdown side, where sponsors facing blocked exits have turned to continuation fund transactions that concentrate rather than resolve the underlying conflict of interest. The lesson generalises beyond this specific question: the absence of a fund termination date is not, by itself, the source of either the traditional model's inefficiencies or the permanent capital model's new challenges. What matters in both cases is the quality of the governance that sits behind the investment decision, a theme the next chapter will return to in the context of exit.

Chapter 5

Exit

5.1 How exit works in the drawdown model

Exit is the phase that closes the venture capital cycle opened by fundraising and continued through investing and monitoring. It is also the point at which the carried interest is finally crystallised: management fees are earned throughout the fund's life, but carry is only realised once investments are exited and proceeds are distributed to investors through the waterfall already discussed above. This distribution mechanic is the same mechanism already quantified in the simulation (Section 3.5). The requirement to pay out realised proceeds rather than reinvest them is precisely what keeps the drawdown fund's average exposure to private assets well below the evergreen fund's constant, fully-deployed allocation, and it accounts for the entire gross annualised performance gap documented there between the two structures. The compounding forgone at exit has already been measured.

5.1.1 Exit channels: composition and a caveat on comparability.

The Gompers–Lerner framework identifies three main exit routes — IPO, trade sale/M&A, and secondary sale — alongside write-off in the case of failure. Two sources in the library offer a picture of how these channels are actually used, though they are not directly comparable and should not be merged into a single figure. Bartlett, Robert and Ramella, Paolo [41], using PitchBook data on US funds raised 1995–2014 (Table 3 and the discussion in Section 5.5, p. 28), report that M&A

dominates exit volume for both PE and VC funds, that IPOs account for 18% of VC exits in the pooled sample (against 12% for PE), and that secondary transactions represent a smaller but non-negligible share. The IPO share has fallen markedly across vintages, from 26% of VC exits in the 1995–1999 cohort to 17% in the 2010–2014 cohort (Figure 10), while M&A — and in particular M&A by "serial" acquirers, defined as buyers completing at least fifteen acquisitions of PE- or VC-backed companies within a four-year window — has absorbed the difference and, for VC funds, provides liquidity one to one-and-a-half years faster than ordinary acquisitions (Figure 11). A separate European-focused survey in the library reports, for a sample of European VC respondents, that trade sales account for 40% of exits, IPOs for only 7%, and outright failure for 22%, with 28% of investments exceeding their original objectives (Table 3.1).[72] The two samples are not directly comparable — one is a US, fund-level PitchBook sample; the other a European, GP-level survey — so they should not be averaged or presented as a single statistic. But they agree on the substantive point: trade sale/M&A is the dominant exit channel, and the IPO, while the most visible in the popular imagination of venture capital, is in both samples a minority outcome.

Bartlett, Robert and Ramella, Paolo [41] note that the ten-year fund "remains a contractual benchmark but no longer corresponds to the economic horizon over which value is realised and returned" (p. 4): even for recent vintages, distributions relative to reported NAV have stayed broadly stable, but actual fund life has lengthened well past the stated ten years as GPs work through the unwinding of residual, illiquid positions. The formal duration of the drawdown vehicle increasingly diverges from the real time needed for value to mature, echoing — from the standpoint of how long funds actually live rather than how long GPs would like them to live — the same duration mismatch that Saraceni describes from the LP side.

5.1.2 Exit pressure

The need to raise a new fund every three to four years, combined with the current fund's fixed life, turns exit decisions into reputation-and-carry-driven choices rather than value-driven ones. In its original formulation this was framed as pushing GPs toward premature M&A *or IPO*.

Masulis, Ronald W. and Nahata, Rajarishi [73] explain the pressure directly: a VC fund has a fixed capital base and a fixed maturity date, and — in order to retain favourable tax treatment — must pay out realised proceeds rather than reinvest them in the remaining portfolio. As a fund nears maturity it therefore has progressively less capital available to support its remaining companies, while facing growing pressure to exit them; this is reinforced by the GP's incentive to liquidate earlier so as to lower the cost of capital and raise IRR, the metric on which performance is judged. The authors state plainly that "VCs in more mature funds nearing liquidation can experience greater incentives to complete acquisitions to eliminate any further funding needs of these portfolio firms… This liquidity pressure can lead VCs to pressure target managers to sell their firms more quickly, even at the cost of a lower purchase price". The paper's central empirical result, is that as funds approach maturity, acquirer announcement returns rise and target takeover premia fall significantly — the standard signature, in the finance literature, of a seller accepting worse terms under time pressure. However, it is worth being precise about scope: the paper's sample consists exclusively of acquisitions of VC-backed private firms. It says nothing, one way or the other, about IPO timing — the question simply falls outside what the paper measures. Yao, Ting and O'Neill, Hugh M. [74] test the M&A/IPO distinction directly, on a sample of 219 US venture-backed surgical-device ventures founded between 1989 and 2014. The result, stated in the abstract and repeated in the discussion, is clear: VCs' exit pressure significantly increases the hazard of high-value M&A, low-value M&A, and liquidation, but has no effect on the hazard of IPO. The authors offer a specific reason for this: promising ventures on an IPO trajectory tend to attract new investors willing to buy out the time-pressured VC, effectively decoupling those companies from the fund's own deadline. About this type of problem, asked explicitly whether moving from a ten-year structure to a more patient-capital one would change "the pressure on follow-on or exit decisions, or the entrepreneurial discipline" felt by founders, Massimiliano Magrini — who himself manages drawdown funds on the standard five-plus-five schedule — agreed with that.[1] The mechanism also surfaces in interview. Discussing whether LIFTT's longer horizon creates friction inside investment syndicates, Luca Salerno (LIFTT) confirms that co-investing funds approaching the end of their own investment period create exactly this kind of misalignment: LIFTT can wait for an asset to mature,

[1] Personal interview with Massimiliano Magrini, Co-founder and Managing Partner, United Ventures.

while a fund near the end of its life cannot, purely because of its own calendar.[2] This does not add new evidence for the mechanism itself, but it confirms that the dynamic documented in the academic literature is recognised in practice by a permanent-capital investor observing it in other funds' behaviour. On the LP in interview, Saraceni expressed a general preference for holding on to positions in successful portfolio companies rather than being forced to distribute, and said she would have wanted to keep several of her own positions open longer than the fund structure allowed.

5.1.3 The permanent-capital response

The permanent-capital response is, in its most direct form, the removal of that term itself: no obligation to liquidate by a preset date, and therefore no need to sell a winning position simply because the fund's calendar requires it. On its own website, Sequoia explicitly describes an "open-ended liquid portfolio made up of public positions in a selection of our enduring companies". After an IPO, rather than distributing the resulting shares to LPs as a drawdown fund would, Sequoia retains them within the master fund, in a continuous reinvestment loop. Redalpine makes a similar argument more directly on its company blog: by avoiding forced exits, evergreen vehicles can capture outsized gains that occur beyond the ten-to-twelve-year mark, and the firm's own Summit Fund reinvests partial proceeds and maintains exposure to top-performing companies while still offering orderly redemptions to investors. In both cases the substantive point is the partial or total reinvestment of proceeds in place of distribution, not the fund structure as such, which is not repeated here. Basnet, Anup and Blomkvist, Magnus and Cumming, Douglas [75] address a limitation in the earlier literature, which measured VC-backed IPO performance over a fixed post-IPO window without distinguishing whether the VC was still present, and consequently found no clear outperformance. By hand-collecting actual VC exit dates for a sample of US VC-backed IPOs, they show that while the lead VC remains present, the stock earns a significant positive abnormal return; once the VC exits, this outperformance holds up only for highly-reputable VCs with large syndicates and long holding periods — for lower-reputation VCs it disappears. The mechanism they propose is active monitoring, not simply the selection of better companies at the outset. This is close to the exact profile — high reputation, broad syndicates, long

[2] Interview with Luca Salerno, Business Analyst Team, LIFTT S.p.A.

holding periods — on which Sequoia builds its own case: holding a position after an IPO has, on this evidence, a real empirical basis for firms of that calibre, though not as a rule that generalises to venture capital as a whole. A more informative account, precisely because it does not come from an evergreen vehicle describing itself, is Magrini's. Asked what makes Sequoia's model the "textbook" version of evergreen venture capital, he explains that Sequoia placed an evergreen vehicle on top of its closed-end funds specifically to avoid a situation in which someone is forced to liquidate a stake prematurely when holding it, even once listed, could deliver equal or superior returns.[3] He describes the mechanism as operating on two levels: the LP invests only in the master fund, where LP inflows and outflows occur, while the underlying funds invest directly in portfolio companies, so that liquidity reaches the LP only after it has already been realised at the level of the underlying position. When I asked directly whether this allows the vehicle to benefit from longer holding periods on companies that, even after an IPO or an exit, still carry value the fund could capture, his answer agreed with that. LIFTT offers a partial parallel, though it is not, technically, an evergreen structure. Odasso explains that whether proceeds from an exit are reinvested into the remaining portfolio, as in an evergreen fund, depends on a shareholder resolution: the holding can choose to retain proceeds for reinvestment in higher-potential assets, or distribute them pro rata to shareholders, whereas a "pure" evergreen fund offers redemption windows at fixed intervals, priced off the fund's reported value.[4] Merlino, for his part, distinguishes the two exit routes available to a LIFTT shareholder — a listing or dividends — from what he calls "a third way": the genuine evergreen structure, which allows redemptions even on a three-to-five-year basis, but never on the full stake at once.[5]

5.1.4 Considerations

Redalpine uses an evergreen structure, the Summit Fund, but not across its whole portfolio: traditional vintage funds cover the early stage, while the evergreen vehicle sits alongside specifically for the growth part. This makes it natural to put the same question directly to Magrini, who runs a traditional growth fund: would an evergreen structure have made sense for United Ventures' own growth fund? The answer is no.

[3] Personal interview with Massimiliano Magrini, co-founder and Managing Partner, United Ventures.

[4] Personal interview with Cristina Odasso, Head of Business Analysis, LIFTT S.p.A.

[5] Personal interview with Giuseppe Merlino, CFO, LIFTT S.p.A.

It would not have been possible to propose it to investors who had already committed to the existing growth fund, because, in his words, they "would absolutely not have accepted it."[6] This is the same LP path-dependency constraint already seen for venture funds, confirmed here for growth as well: the constraint does not depend on the investment stage, but on the kind of relationship already established with existing investors. This same LP-side constraint is the central theme of the Saraceni interview: even sophisticated institutional investors are not willing to remain illiquid for the twelve to fifteen years that deep tech often requires, and they need liquidity events.[7] Evergreen funds respond to this with redemption windows at pre-set intervals. LIFTT has none: as Merlino puts it, "essendo una società privata non è liquida, non puoi uscire quando vuoi, néanche parzialmente, cosa che invece è permesso a finestre nei tipici fondi Evergreen."[8] LIFTT's investors, however, know this from the start: the only two exit routes, a listing or dividends, have always been communicated as such. Magrini adds an important point. Redemption windows do not solve the illiquidity of the underlying investment; they are only a mechanism that lets the LP enter or exit the fund. They do not speed up the actual liquidation of the portfolio companies. Questioning directly whether the constraint was on the LP side, he answers: "The LP wants shorter-term liquidity, but they are not yet ready to buy into an evergreen fund. Because in the end it is not really about having an evergreen fund or not — what matters, in the end, is how much you can realise the value of the underlying portfolio. The evergreen fund only gives the option, eventually, to let LPs enter and exit when they want to enter and exit, but having an evergreen fund does not guarantee you a faster turnover of the portfolio — it simply means having windows for exit and entry." Portfolio companies remain private investments, with the same maturation timelines, regardless of the vehicle that holds them. The growth-fund case shows that when an existing LP base has to be converted, the constraint is binding: investors simply do not accept the change. Yet Magrini himself sees the coexistence of the two structures favourably, as a market option rather than a substitution: "This would be good for the market. I think it would be right for the market to have a combination of closed-end funds and permanent [funds]. I think that would be entirely desirable." This is exactly what Redalpine does in practice: vintage funds and the evergreen

[6]Personal interview with Massimiliano Magrini, co-founder and Managing Partner, United Ventures.

[7]Interview with Diana Saraceni, Panakès Partners; see Interview_Summary_Saraceni_Panakes, Section 3.6.

[8]Personal interview with Giuseppe Merlino, CFO, LIFTT S.p.A.

Summit Fund side by side, not one replacing the other — closing the block on the same example it opened with.

Chapter 6

Conclusion

This thesis set out to answer a single question: whether permanent capital vehicles represent a genuine structural improvement over the traditional ten-year drawdown fund, or merely a complementary offering aimed at a different kind of investor. To address it, the analysis adopted the Gompers–Lerner cycle as its scaffold, examining each of the three phases of a venture capital fund's life—fundraising, investing and monitoring, and exit—along the same three dimensions: the structural logic of the drawdown model, the inefficiencies the literature and practitioners attribute to it, and the response that permanent capital alternatives offer. What follows draws those findings together.

What the drawdown model does well, and where it strains

The recurring finding across all three phases is that the ten-year closed-end structure is not a flawed design but a coherent equilibrium—one that solves a real set of problems and, precisely because it solves them, has proven durable. The commitment-and-call mechanism minimises idle capital inside the fund; the finite life and lock-up make a multi-year illiquid strategy financeable; the standardisation of terms, the portability of the track record, and the implicit expectation of reinvestment in successor funds sustain the model's dominance across cycles. Several of the "inefficiencies" commonly levelled at the model turn out, on inspection, to be functional costs rather than

genuine defects. Fundraising procyclicality matters less than it first appears, because established managers with strong track records and loyal investors are relatively insulated from it, and because capital called in adverse windows tends to earn a liquidity premium. Cash drag in the strict sense is minimised by design. The J-curve is largely a measurement artefact of fees charged on committed rather than invested capital, not a structural loss.

Where the model genuinely strains is more specific, and the thesis located it precisely. The most consequential fundraising-phase inefficiency is the drag borne by the limited partner on committed-but-uncalled capital, held in low-yielding liquid form and invisible to the fund's IRR because it never enters the fund's cash flows. The simulation quantified the magnitude of the underlying exposure gap: a single drawdown fund holds, on average, only about a quarter of committed capital invested in private assets at any given time, with the remainder earning the risk-free rate—an exposure penalty that the evergreen structure, fully deployed from day one, does not incur. In the investing and monitoring phase, the binding constraint is temporal: the fund's finite life is not calibrated to the development horizons of the ventures the industry increasingly needs to finance. Deep tech, with cycles that can extend to twelve or fifteen years, exposes this mismatch most sharply. And at exit, the requirement to liquidate by a preset date can turns exit timing into a reputation-and-carry-driven decision rather than a purely value-driven one—a pressure the empirical literature associates with a higher hazard of trade sale and liquidation, though not of IPO.

What permanent capital resolves, and what it merely relocates

Permanent capital vehicles address the two inefficiencies that are structural rather than merely apparent. Immediate deployment removes the exposure-timing penalty; the absence of a termination date removes the forced-exit pressure and allows winning positions to compound. But the central lesson of the thesis is that these vehicles do not eliminate the underlying tensions so much as relocate them. The problem of performance measurement does not disappear—it migrates from the IRR to the NAV, which in an evergreen fund becomes simultaneously the performance signal, the transaction price for subscriptions and redemptions, and the fee base,

and in a listed holding is replaced by a market price that trades at persistent, often steep, discounts to the value of the underlying portfolio—40 to 60 per cent on the secondary market for venture portfolios in recent evidence. Liquidity, likewise, is not created but repackaged: as Magrini put it, redemption windows let the limited partner enter and exit the fund, but they do nothing to accelerate the maturation of the underlying companies, which remain private investments on their own timeline. The simulation made this trade-off concrete. Once the recurring management fee and a realistically calibrated liquidity-stress-and-gating mechanism are both priced in, the evergreen fund's net return falls slightly *below* the drawdown fund's own net-of-fees return—a reversal that held across every simulation seed tested. This is not a verdict that one structure is superior; it is quantitative support for the thesis's central claim. Each structure substitutes one source of fragility for another: the drawdown fund carries an exposure-timing penalty, the evergreen fund a liquidity-risk exposure of comparable magnitude.

The demand-side constraint

If the supply-side case for permanent capital is real but qualified, the decisive factor turns out to lie on the demand side—and this is perhaps the qualification the literature alone would not have surfaced. Across three independent interviews, the same constraint emerged. Saraceni would extend fund life for deep tech and life sciences, but located the binding constraint not on the general partner's side but on the limited partner's: even sophisticated institutional investors are unwilling to remain illiquid for that long, and they price the accompanying valuation opacity heavily. Magrini converged on the same point from a different vantage, and widened it beyond any single sector: companies now stay private far longer, so the need for patience is general, yet an investor base assembled around the 5+5 structure cannot easily be redirected toward an open-ended vehicle without effectively resetting the relationship—a constraint he confirmed was binding even for his own growth fund. Merlino reframed rather than contradicted this: LIFTT does not attempt to convert a liquidity-sensitive institutional pool but builds its capital base around investors who accept the patient-capital proposition from the outset. The rise of permanent capital is therefore better understood as the gradual formation of a self-selected investor

base around patience than as a wholesale shift in what all investors are prepared to accept.

Complementarity, not substitution

Taken together, these findings point to a single conclusion. Permanent capital vehicles and traditional drawdown funds are not substitutes competing to occupy the same space, but complementary structures whose suitability depends on the investor, the sector, and the tolerance for illiquidity. The simulation reinforced this from the investor's standpoint: the operational complexity of the drawdown model—capital calls, treasury management, the discipline required to manage uncalled commitments—favours large institutional investors equipped to handle it, while the operational simplicity of permanent capital vehicles suits those seeking sustained exposure without active treasury management, which is also what raises the broader, still-open question of democratisation and wider access to private markets. Yet access, as the evidence made clear, is not the same as benefit: broadening entry without addressing the intermediation frictions and governance weaknesses that accompany these structures does not automatically translate into investor return, and the value of semi-liquid formats appears to hold in information-insensitive assets far more readily than in equity venture. Practitioners themselves see the two models coexisting rather than one displacing the other. Redalpine's practice embodies it—vintage funds for the early stage, an evergreen vehicle alongside for growth—and Magrini stated the principle directly: a market combining closed-end and permanent structures would be, in his view, entirely desirable.

The broader methodological lesson that recurs throughout this thesis is that the presence or absence of a fund termination date is not, by itself, the source of either the traditional model's inefficiencies or the permanent capital model's new challenges. A closed-end fund without a genuine exit route does not avoid indefinite holding; it merely relocates the decision into a related-party continuation transaction with its own conflicts. A permanent capital vehicle without a deadline does not automatically drift into indefinite holding either, provided its governance genuinely ties capital allocation to updated expected returns. In both cases, what determines the outcome for the investor is the quality of the governance behind the investment decision, not the calendar. The ten-year fund persists not because it is a first-best design, but

because it satisfies the participation constraint of an investor base that is not yet ready for something more patient. Permanent capital changes what a vehicle *can* do; whether it does so to the investor's benefit remains, as ever, a question of governance.

Limitations and directions for further research

Several limitations qualify these conclusions and mark out where further work is needed. The qualitative evidence, though triangulated across general partners and a permanent-capital operator, rests on a small number of interviews, all drawn from the Italian and broader European market; no evergreen VC fund could be reached directly, and no limited partner or portfolio-company founder was interviewed, so the demand-side constraint that proved central is documented here through the general partners who observe it rather than through the investors who impose it. The simulation, for its part, isolates structural mechanisms under deliberately stylised assumptions: its purpose is to illuminate how fund design shapes limited-partner outcomes, not to forecast the return of any specific vehicle, and its calibration— particularly the venture capital distribution schedule and the liquidity-stress parameters—necessarily rests on choices that future work could refine against richer data. Three directions in particular would extend the analysis: direct evidence from limited partners on how they actually price illiquidity and valuation opacity; a larger sample of evergreen VC vehicles as the market matures and access improves; and closer study of the governance mechanisms that, on the argument advanced here, ultimately determine whether either structure serves the investor well.

References

[1] Paul Gompers and Josh Lerner. The venture capital revolution. *Journal of Economic Perspectives*, 15(2):145–168, 2001.

[2] Gregory W. Brown and William Volckmann. Evergreen vs. drawdown funds: Risk, returns and cash flows. *Institute for Private Capital Working Paper*, 2025. June 11, 2025.

[3] Luciano E. «Applied Corporate Finance and Private Capital — Lesson 1: Introduction to the course». Lecture slides, M.Sc. course, Politecnico di Torino, 2025. Academic Year 2024–2025.

[4] Luciano E. «Valuation». Lecture slides, Applied Corporate Finance and Private Capital (Lesson 7), M.Sc. course, Politecnico di Torino, 2025. Academic Year 2024–2025.

[5] Kaplan S. N. and Strömberg P. «Leveraged Buyouts and Private Equity». *Journal of Economic Perspectives*, 23(1):121–146, 2009.

[6] Metrick A. and Yasuda A. «The Economics of Private Equity Funds». *The Review of Financial Studies*, 23(6):2303–2341, 2010.

[7] Jensen M. C. «Eclipse of the Public Corporation». *Harvard Business Review*, 67(5):61–74, 1989.

[8] Doidge C., Kahle K. M., Karolyi G. A., and Stulz R. M. «Eclipse of the Public Corporation or Eclipse of the Public Markets?». *Journal of Applied Corporate Finance*, 30(1):8–16, 2018.

[9] Gözlügöl A. A., Greth J., and Tröger T. H. «The Oscillating Domains of Public and Private Markets». Law Working Paper 689/2023, European Corporate Governance Institute (ECGI), 2023. Also SAFE Working Paper No. 384. Available at SSRN: https://ssrn.com/abstract=4372011.

[10] Paul A. Gompers, Will Gornall, Steven N. Kaplan, and Ilya A. Strebulaev. Venture capitalists and COVID-19. *Journal of Venture Capital*, 23(2):83–104, 2021. NBER Working Paper No. 27824.

[11] Sabrina T. Howell, Josh Lerner, Ramana Nanda, and Richard R. Townsend. Financial Distancing: How Venture Capital Follows the Economy Down and Curtails Innovation. Working Paper 27150, National Bureau of Economic Research, 2020.

[12] KPMG. Venture pulse q4 2021: Global analysis of venture funding, 2021.

[13] Andrea Bellucci, Alexander Borisov, Gianluca Gucciardi, and Alberto Zazzaro. The reallocation effects of COVID-19: Evidence from venture capital investments around the world. *Journal of Banking and Finance*, 147:106443, 2023.

[14] Enrico Luciano. «Introduction to Private Capital», 2025. Lesson 15, Applied Corporate Finance and Private Capital, Politecnico di Torino, April 14, 2025.

[15] McKinsey & Company. Global private markets report 2026: Private equity—clearer view, tougher terrain, 2026.

[16] Alisa Amarosa Wood, Paula Campbell Roberts, and Justin Park. An evergreen vehicle can be an important tool in private equity asset allocation. Technical report, KKR Wealth Insights, 2024.

[17] Alexander Blostein, Michael Vinci, Luke Bianculli, Anthony Corbin, Aditya Sharma, and Mohak Batra. Americas asset managers: State of retail alternatives. Technical report, Goldman Sachs Equity Research, 2024.

[18] Daniel Murphy and Juliana Hadas. Choose your vehicle: A closer look at private market fund structures. Technical report, Goldman Sachs Featured Insights, 2024.

[19] Bill Myers. SEC committee flips narrative on retail investors in private funds. *Private Funds CFO*, 2024.

[20] William A. Sahlman. «the structure and governance of venture-capital organizations». *Journal of Financial Economics*, 27(2):473–521, 1990.

[21] Kate Litvak. Venture Capital Limited Partnership Agreements: Understanding Compensation Arrangements. *University of Chicago Law Review*, 76(1):161–218, 2009.

[22] Andrew Metrick and Ayako Yasuda. Venture Capital and Other Private Equity: A Survey. *European Financial Management*, 17(4):619–654, 2011.

[23] Brad M. Barber and Ayako Yasuda. Interim Fund Performance and Fundraising in Private Equity. *Journal of Financial Economics*, 124(1):172–194, 2017.

[24] Josh Lerner. The Future of Public Efforts to Boost Entrepreneurship and Venture Capital. *Journal of Economic Perspectives*, 24(3):153–167, 2010.

[25] Josh Lerner, Antoinette Schoar, and Wan Wongsunwai. Smart Institutions, Foolish Choices: The Limited Partner Performance Puzzle. *The Journal of Finance*, 62(2):731–764, 2007.

[26] Paul Gompers and Josh Lerner. An Analysis of Compensation in the U.S. Venture Capital Partnership. *Journal of Financial Economics*, 51(1):3–44, 1999.

[27] Steven N. Kaplan and Antoinette Schoar. Private Equity Performance: Returns, Persistence, and Capital Flows. *The Journal of Finance*, 60(4):1791–1823, 2005.

[28] Bernard S. Black and Ronald J. Gilson. Venture Capital and the Structure of Capital Markets: Banks versus Stock Markets. *Journal of Financial Economics*, 47(3):243–277, 1998.

[29] Paul Gompers, Anna Kovner, Josh Lerner, and David Scharfstein. Venture Capital Investment Cycles: The Impact of Public Markets. *Journal of Financial Economics*, 87(1):1–23, 2008.

[30] Paul Gompers and Josh Lerner. Money Chasing Deals? The Impact of Fund Inflows on Private Equity Valuations. *Journal of Financial Economics*, 55(2–3):281–325, 2000.

[31] David T. Robinson and Berk A. Sensoy. Cyclicality, Performance Measurement, and Cash Flow Liquidity in Private Equity. *Journal of Financial Economics*, 122(3):521–543, 2016.

[32] Joost Driessen, Tse-Chun Lin, and Ludovic Phalippou. A New Method to Estimate Risk and Return of Nontraded Assets from Cash Flows: The Case of Private Equity Funds. *Journal of Financial and Quantitative Analysis*, 47(3):511–535, 2012.

[33] Ludovic Phalippou. The Tyranny of the Internal Rate of Return. SSRN Working Paper No. 5042563, 2025.

[34] Ludovic Phalippou and Oliver Gottschalg. The Performance of Private Equity Funds. *Review of Financial Studies*, 22(4):1747–1776, 2009.

[35] Josh Lerner and Antoinette Schoar. The Illiquidity Puzzle: Theory and Evidence from Private Equity. *Journal of Financial Economics*, 72(1):3–40, 2004.

[36] Stefano Pegoraro, Sophie Shive, and Rafael Zambrana. Democratizing Illiquid Assets: Liquidity Transformation and Performance in Interval Funds. SSRN Working Paper No. 5365061, 2026.

[37] Ludovic Phalippou. Private Equity's Retail Turn: Anatomy of a Mis-selling Risk. SSRN Working Paper No. 6080266, University of Oxford, Saïd Business School, 2026.

[38] William W. Clayton and Elisabeth de Fontenay. Private Equity for All: The Paradoxical Push to Democratize Private Markets. SSRN Working Paper No. 6157128, 2026.

[39] Elise Gourier, Ludovic Phalippou, and Mark M. Westerfield. Capital Commitment. SSRN Working Paper No. 2490430, 2022.

[40] Nori Gerardo Lietz. Should Mom Have Private Equity in Her 401(k)? SSRN Working Paper No. 6304280, Harvard Business School, 2025.

[41] Robert Bartlett and Paolo Ramella. The disappearance of the ten-year fund. Technical Report Working Paper No. 616, Stanford Law School, John M. Olin Program in Law and Economics, April 2026. SSRN 6664338.

[42] Cynthia Balloch, Federico Mainardi, Sangmin S. Oh, and Petra Vokata. «democratizing private markets? private equity performance of individual investors». Technical report, London School of Economics, Columbia Business School and Ohio State University, March 2026.

[43] Gregory W. Brown, Christian Lundblad, and William Volckmann. «risk-adjusted performance of private funds: What do we know?». Technical report, Institute for Private Capital, Kenan-Flagler Business School, University of North Carolina, 2025.

[44] John H. Cochrane. «the risk and return of venture capital». *Journal of Financial Economics*, 75:3–52, 2005.

[45] Arthur Korteweg and Morten Sorensen. «skill and luck in private equity performance». *Journal of Financial Economics*, 2015. Working paper, October 2015; published version 2017, vol. 124, pp. 535–562.

[46] Arthur Korteweg and Stefan Nagel. «risk-adjusted returns of private equity funds: A new approach». *The Review of Financial Studies*, 2024.

[47] Martin Cherkes, Jacob S. Sagi, and Richard H. Stanton. «a liquidity-based theory of closed-end funds». *The Review of Financial Studies*, 22(1):257–297, 2009.

[48] Paul A. Gompers. «optimal investment, monitoring, and the staging of venture capital». *Journal of Finance*, 50(5):1461–1489, 1995.

[49] Yao-Wen Hsu. «staging of venture capital investment: A real options analysis». Working paper, University of Cambridge, 2002.

[50] Steven N. Kaplan and Per Strömberg. «characteristics, contracts, and actions: Evidence from venture capitalist analyses». *Journal of Finance*, 59(5):2177–2210, 2004.

[51] Carsten Bienz and Julia Hirsch. «the dynamics of venture capital contracts». Working Paper 929491, SSRN, 2005.

[52] Xuan Tian and Tracy Yue Wang. «tolerance for failure and corporate innovation». *Review of Financial Studies*, 27(1):211–255, 2014.

[53] David J. Denis. «entrepreneurial finance: An overview of the issues and evidence». *Journal of Corporate Finance*, 10(2):301–326, 2004.

[54] Will Gornall and Ilya A. Strebulaev. «venture capital contracts». *Journal of Financial Economics*, 143(1):131–158, 2022.

[55] Douglas Cumming. «contracts and exits in venture capital finance». *Review of Financial Studies*, 21(5):1947–1982, 2008.

[56] Ola Bengtsson and Berk A. Sensoy. «investor abilities and financial contracting: Evidence from venture capital». *Journal of Financial Intermediation*, 20(4):477–502, 2011.

[57] Mikko Jääskeläinen. «venture capital syndication: Synthesis and future directions». Working paper, Helsinki University of Technology, 2010.

[58] Sophie Manigart, Andy Lockett, Miguel Meuleman, Mike Wright, Hans Landström, Hans Bruining, Philippe Desbrières, and Ulrich Hommel. «why do european venture capital companies syndicate?». 2002.

[59] Yael V. Hochberg, Alexander Ljungqvist, and Yang Lu. «whom you know matters: Venture capital networks and investment performance». *Journal of Finance*, 62(1):251–301, 2007.

[60] Xuan Tian. «the role of venture capital syndication in value creation for entrepreneurial firms». *Review of Finance*, 16(1):245–283, 2012.

[61] Jochen Christian Werth. «syndication networks and investment activity in the venture capital industry». Working paper, Goethe University Frankfurt, 2014.

[62] Harry J. Sapienza, Sophie Manigart, and Wim Vermeir. «venture capitalist governance and value added in four countries». *Journal of Business Venturing*, 11(6):439–469, 1996.

[63] Douglas Cumming and Sofia A. Johan. «advice and monitoring in venture finance». *Financial Markets and Portfolio Management*, 21(1):3–43, 2007.

[64] Thomas Hellmann and Manju Puri. «venture capital and the professionalization of start-up firms: Empirical evidence». *Journal of Finance*, 57(1):169–197, 2002.

[65] Laura Lindsey. «blurring firm boundaries: The role of venture capital in strategic alliances». *Journal of Finance*, 63(3):1137–1168, 2008.

[66] Onur Bayar, Thomas J. Chemmanur, and Xuan Tian. «peer monitoring, syndication, and the dynamics of venture capitalist interactions». *Journal of Financial and Quantitative Analysis*, 2018.

[67] Matteo Ambrois, Vincenzo Butticè, Annalisa Croce, Luca Grilli, and Elisa Ughetto. «the times they are a-changin': How venture capital firms change their investment practices under the covid-19 pandemic». *Small Business Economics*, 65:871–893, 2025.

[68] Johann Kortsch, Stefan Raff-Heinen, David Bendig, Martin Murmann, Colin Schulz, and Fiona Murray. «deep-tech innovation: A multi-method study toward a conceptual framework and research agenda». Working paper, 2022.

[69] Angelo Romasanta, Gozal Ahmadova, Jonathan Wareham, and Laia Pujol Priego. «deep tech: Unveiling the foundations». Working Paper 276, ESADE Business School, 2022.

[70] Daniele Battaglia, Emilio Paolucci, and Elisa Ughetto. «hurdles in university-industry technology transfer: Why research-based inventions are not transferred to the market?». *IEEE Transactions on Engineering Management*, 71:13415–13427, 2024.

[71] Elisa Ughetto, Daniele Battaglia, and Federico Micol. «betting on research-based inventions: An investigation of pre-seed deep-tech venture capital decision-making processes and activities». 2024. Available at https://iris.polito.it/%handle/11583/2992481.

[72] Anna Söderblom, Benjamin Le Pendeven, and Jeroen Verbouw. «Practices of European Venture Capitalists». Technical report, Alliance of European academic institutions, April 2023. Version 2023-04-11.

[73] Ronald W. Masulis and Rajarishi Nahata. «Venture Capital Conflicts of Interest: Evidence from Acquisitions of Venture-Backed Firms». *Journal of Financial and Quantitative Analysis*, 46(2):395–430, 2011.

[74] Ting Yao and Hugh M. O'Neill. «Venture Capital Exit Pressure and Venture Exit: A Board Perspective». *Strategic Management Journal*, 43(13):2829–2848, 2022.

[75] Anup Basnet, Magnus Blomkvist, and Douglas Cumming. «Long-Run IPO Performance and the Role of Venture Capital». Technical report, SSRN, November 2023. Available at SSRN: https://ssrn.com/abstract=4404904.