Building a Pre-IPO Secondary Sleeve in a Family-Office Portfolio

How a family office sizes, structures and manages a disciplined allocation to late-stage private secondaries

Chennakeshav Adya
Independent Researcher
June 2026

ABSTRACT

A pre-IPO secondary holding is the purchase of shares in a late-stage private company from an existing holder, an early investor, a founder or an employee, before any public listing. For a Gulf family office, such positions have become easier to access and harder to ignore as companies stay private for longer and a deep secondary market has formed around their shares. The question this paper addresses is not whether to chase individual names but how to build a coherent sleeve, a sized, structured and governed allocation that sits inside the wider portfolio rather than alongside it as a series of opportunistic bets. The paper treats the sleeve as a design problem with four levers: how large it should be relative to the total portfolio, how it should be accessed and structured, how single positions should be concentrated and diligenced, and how the whole allocation should be governed and marked over a long and uncertain holding period. It sets out a transparent, stylised framework that links each lever to the risks that distinguish pre-IPO secondaries from public equity and from primary venture: price formation against an often stale last round, thin information, long and uncertain time to liquidity, single-name concentration, and the fee and carry layering of the access route. The figures are illustrative and are designed to make the mechanics legible rather than to forecast returns. The paper is written for the principal, the family-office chief investment officer and the adviser who must size and supervise such an allocation, and it is an analytical and educational document, not investment advice, an offer, a solicitation or a recommendation of any security, strategy or transaction.

JEL Classification: G11, G24, G32, G34, D14

Keywords: pre-IPO secondaries, family office, portfolio construction, private markets, liquidity, concentration, late-stage venture, allocation sizing

1. INTRODUCTION

A family office that has decided to take exposure to late-stage private companies faces a question that is easy to state and hard to answer well. It is not, in the first instance, which company to buy. It is how much of the portfolio such positions should occupy, how they should be bought and structured, and how the resulting holdings should be supervised over the years that pass before a listing or a sale turns paper into cash. The individual opportunity tends to arrive framed as a single, urgent decision: a chance to buy shares in a well-known private company at what is presented as an attractive entry. The discipline this paper argues for is to refuse to answer that question on its own terms and to answer a prior one instead. What would a sensible standing allocation to such positions look like, and does this particular opportunity belong inside it?

That standing allocation is what we call a sleeve. A sleeve is a sized, structured and governed pocket of the portfolio with its own mandate, its own limits and its own rules for entry, monitoring and exit. The difference between a sleeve and a collection of opportunistic positions is the difference between an allocation and an accident. A sleeve is decided in advance, in the calm of portfolio planning, rather than in the heat of a single deal. It has a target size and a ceiling. It has concentration limits that no single name may breach. It has a chosen route to access and a view on the fees that route imposes. And it has a governance routine that marks the holdings honestly and reviews them on a schedule rather than only when a crisis or an exit forces attention. The thesis of this paper is that the sleeve, not the individual position, is the right unit of decision for a family office, and that almost everything that goes wrong with private secondaries can be traced to a failure to think at the level of the sleeve.

The relevance of the question has grown because the underlying market has grown. Companies that two decades ago would have listed within a few years of founding now remain private for far longer, raising successive private rounds and reaching valuations that once belonged only to public markets. The consequence is that a large and increasing share of the value created by the most successful private companies accrues while they are still private, and that early holders, founders, employees and seed investors, accumulate stakes they may wish to sell before any listing. Out of that supply and the corresponding demand a secondary market has formed, with brokers, platforms, dedicated funds and structured vehicles that allow a later buyer to acquire shares from an earlier holder. For a family office, this market is both an opportunity and a trap. It is an opportunity because it offers entry to companies that were previously inaccessible. It is a trap because the ease of access disguises how different these positions are from the public equities they superficially resemble.

This paper is written for the Gulf family office in particular, because the calibration of the framework reflects conditions in the United Arab Emirates and Saudi Arabia in 2026: deep pools of patient capital, a strong appetite for direct and co-investment exposure, a preference for structures that are clean and Shariah-aware where required, and a governance culture in which the principal remains close to the decision. None of these features changes the underlying mechanics of a pre-IPO secondary, but each shapes how a Gulf family office should size and supervise the sleeve. The framework is general; the calibration is local.

Three caveats frame the analysis. First, this is an educational and analytical document. It is not investment advice, an offer, a solicitation or a recommendation of any security, fund, strategy or transaction, and nothing in it should be read as such. Second, every figure and table is illustrative and stylised. The numbers are chosen to make relationships legible, not to predict outcomes, and they are not sourced data, forecasts or the performance of any actual position. Third, pre-IPO secondaries are a complex and risky activity that depends on specific facts, on legal and tax advice and on diligence that no general framework can replace. The paper offers a way of thinking, not a substitute for professional counsel.

The paper makes four contributions. First, it reframes the decision from the individual position to the sleeve, and sets out the four levers, size, access, concentration and governance, through which a sleeve is designed. Second, it provides a transparent sizing method that links the sleeve weight to the family office's risk tolerance, liquidity needs and information advantage. Third, it sets out an access spectrum that ranks the routes into pre-IPO secondaries by cost, control and diligence burden, and matches them to the size and ambition of the sleeve. Fourth, it offers a governance routine, including honest marking, concentration discipline and exit planning, that addresses the risks specific to holding illiquid private shares over a long horizon.

The remainder of the paper is organised as follows. Section 2 reviews the relevant literature on portfolio construction, illiquidity, private-market secondaries and family-office behaviour, and draws from it five propositions. Section 3 sets out the data and methodology, makes the illustrative nature of the framework explicit, and records the base-case assumptions. Section 4 presents the framework in the order of the propositions, working through sizing, access, concentration, liquidity and governance with the support of seven illustrative figures. Section 5 sets out a practical, staged process for building and running the sleeve. Section 6 is candid about the risks, the caveats and the limits of the framework. Section 7 concludes. Three appendices record the full assumptions, a sleeve-construction checklist and a glossary.

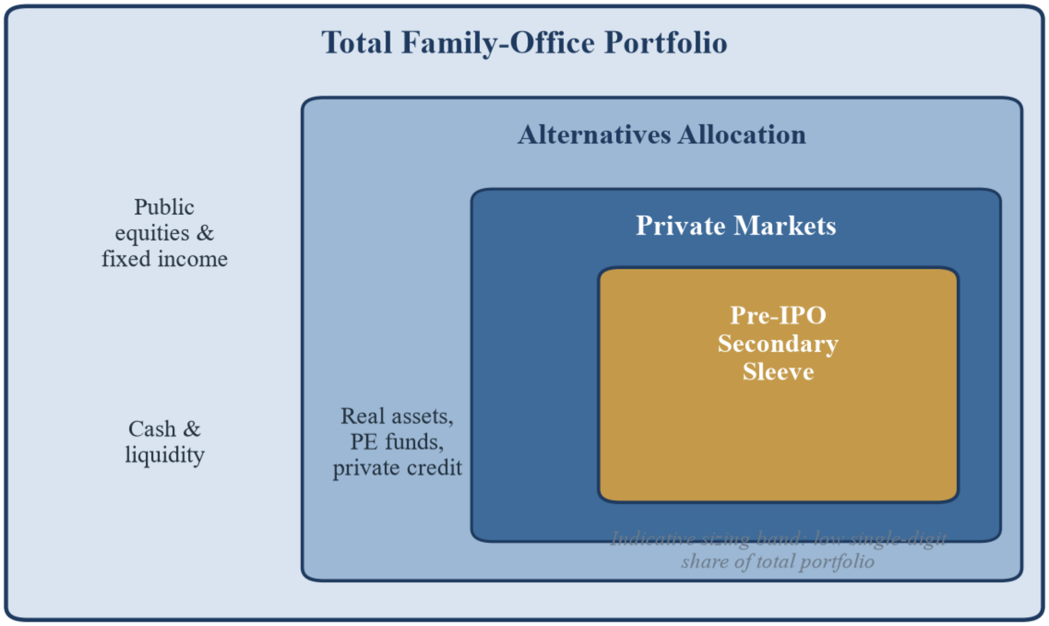

Figure 1. The Pre-IPO Secondary Sleeve Within the Family-Office Portfolio

Illustrative schematic. The sleeve sits inside the alternatives and private-markets allocation as a small, deliberately bounded pocket of the total portfolio, not as a free-standing set of bets.

2. LITERATURE REVIEW AND PROPOSITIONS

The framework draws on four bodies of work: the literature on portfolio construction and position sizing; the literature on illiquidity and its pricing; the literature on private-market secondaries and late-stage venture; and the literature on family-office investing and behavioural discipline. Each contributes a part of the case for treating the sleeve, rather than the position, as the unit of decision.

2.1 Portfolio Construction and Position Sizing

The foundational insight of modern portfolio theory, following Markowitz (1952), is that the contribution of an asset to a portfolio depends not on its own expected return and risk alone but on how it interacts with everything else the investor holds. A position that looks attractive in isolation can be unattractive in a portfolio if it adds risk that is already abundant, and a position that looks modest can be valuable if it diversifies. For pre-IPO secondaries this matters because such positions are highly idiosyncratic, driven by the fortunes of single companies, and because their correlation with the rest of a family office's holdings is neither zero nor one but uncertain and state-dependent. The sizing of the sleeve is therefore a portfolio question, not a position question.

The literature on position sizing under uncertainty adds a second discipline. The Kelly criterion and its practical descendants warn that even a favourable bet, sized too large, can ruin a portfolio through the mathematics of compounding and drawdown. Practitioners routinely size well below the theoretical optimum because the inputs are uncertain and the cost of overbetting is severe and asymmetric. The lesson for a pre-IPO sleeve is that conviction in a single name is not a reason to size it large; the right size is bounded by the damage a total loss would do to the whole portfolio, not by the upside a success would bring.

A further strand of the portfolio literature, on diversification and the law of active management, establishes that the reliability of an investment programme rises with the number of genuinely independent decisions it makes. An investor who concentrates in a handful of positions, however well chosen, holds a portfolio whose outcome is dominated by a small number of draws and is therefore unreliable even if its expected return is high. For an asset class whose individual outcomes are as dispersed as late-stage private companies, this argues strongly for breadth within the sleeve, and against the instinct to express conviction through concentration.

2.2 Illiquidity and Its Pricing

A second literature concerns illiquidity, the inability to sell an asset quickly at a fair price, and how it should be compensated and managed. The asset-pricing tradition, following Amihud and Mendelson (1986) and the later work on liquidity as a priced risk factor, holds that investors should demand a premium for bearing illiquidity and that the premium is real but variable. The corollary, developed in the literature on liability-driven and endowment investing, is that an investor can only harvest the illiquidity premium to the extent of its genuine tolerance for being locked in. An investor that may be forced to sell at an inopportune moment cannot safely hold illiquid assets, however attractive the premium, because the forced sale destroys the very return the premium was meant to provide.

This literature bears directly on the sleeve. Pre-IPO secondaries are among the least liquid positions a family office can hold: there is no continuous market, no reliable price, and no certainty about when, or whether, a liquidity event will occur. The size of the sleeve must therefore be governed by the family office's true liquidity tolerance, measured against its spending needs, its commitments elsewhere and its appetite for holding an asset it cannot value precisely or sell on demand. The illiquidity literature also cautions that marks on illiquid assets are smoothed and stale, which understates true volatility and can lull an investor into oversizing.

2.3 Private-Market Secondaries and Late-Stage Venture

A third literature examines the secondary market in private assets and the economics of late-stage venture. The secondary literature distinguishes the buyer's motivations, access to otherwise unavailable assets, a shorter path to liquidity than primary investing, and the chance to buy at a discount to the most recent valuation, from the buyer's risks, chiefly thin information and the difficulty of pricing against a last round that may be stale, structured or unrepresentative. The empirical work on venture returns, following the tradition of Cochrane (2005) and the later literature on return dispersion, establishes the central stylised fact that pre-IPO sleeve design must respect: returns to individual late-stage private companies are extraordinarily dispersed and skewed, with a minority of large winners, a body of modest outcomes and a meaningful tail of impairments.

Two implications follow. First, because outcomes are so dispersed, a single position is a poor proxy for the asset class; only a diversified set of positions, or a fund that holds many, gives the sleeve a chance of capturing the skew rather than landing in the tail. Second, because the entry price is set against a last round that may not reflect current value, the discount or premium to that round is an unreliable guide on its own and must be read alongside the quality of the information and the structure of the shares being bought. The secondary literature is clear that headline discounts often compensate for information asymmetry or for structural features, such as a later liquidation preference, that erode the value of the common shares a secondary buyer typically acquires.

2.4 Family-Office Investing and Behavioural Discipline

A fourth literature concerns how family offices actually invest and the behavioural risks they face. The family-office literature documents a strong appetite for direct and co-investment exposure, a preference for control and proximity, and a tendency to concentrate in areas of perceived expertise or personal interest. These traits are strengths: they bring patient capital and genuine engagement. But the behavioural-finance literature, following Kahneman and Tversky and the later work on overconfidence and the disposition effect, warns that the same traits produce predictable errors, oversizing positions in which the principal has conviction, anchoring on the entry price, holding losers too long and selling winners too early, and confusing access to a famous name with an edge in pricing it.

The remedy the literature proposes is structural rather than exhortatory. Rules decided in advance, position limits, a sizing method, a marking policy and a review schedule, protect the investor from its own later judgement under pressure. This is the deepest argument for the sleeve. The sleeve is a commitment device: by deciding the size, the limits and the rules before any specific opportunity arrives, the family office binds itself to a discipline that the individual deal, with its urgency and its narrative, would otherwise erode.

2.5 Synthesis and Propositions

Taken together, these literatures imply that a pre-IPO secondary allocation should be designed as a bounded, diversified and governed sleeve whose size reflects genuine liquidity tolerance, whose positions respect the dispersion of outcomes, and whose rules are set in advance to discipline behaviour. They yield five propositions that the rest of the paper develops and illustrates.

Proposition 1. The right unit of decision is the sleeve, not the individual position. Size, access, concentration and governance should be decided at the level of the allocation and in advance, so that any single opportunity is judged by whether it fits the sleeve.

Proposition 2. The size of the sleeve is bounded by the family office's true liquidity tolerance and by the damage a total loss would do to the whole portfolio, not by conviction in any single name.

Proposition 3. The route of access trades cost, control and diligence burden against directness, and should be matched to the size and the ambition of the sleeve rather than chosen by default or by what is offered first.

Proposition 4. Because outcomes are highly dispersed and skewed, the sleeve must be diversified across several positions or accessed through vehicles that diversify, and single-name concentration must be capped.

Proposition 5. The sleeve must be governed over a long and uncertain horizon through honest marking, scheduled review and explicit exit planning, because the principal risk is not only loss but the slow erosion of discipline and the surprise of stale marks.

3. DATA AND METHODOLOGY

This study is analytical and descriptive rather than empirical. It builds a framework for designing and governing a pre-IPO secondary sleeve and illustrates the framework with stylised figures whose inputs are stated and held constant. It does not estimate parameters from a dataset, test hypotheses statistically or report the performance of any actual position, fund or family office. Its purpose is to make the structure of the decision legible and to give a family office a disciplined way of reasoning about it.

3.1 The Sleeve-Design Framework

The framework has four components, corresponding to the four levers of sleeve design. The first is a sizing method that links the target weight of the sleeve to the family office's risk tolerance, liquidity needs and information advantage, and sets a ceiling that no opportunity may breach. The second is an access spectrum that ranks the routes into pre-IPO secondaries, from listed pre-IPO funds through multi-manager secondary funds, single-name special-purpose vehicles and structured contracts to direct purchase, by cost, control and diligence burden. The third is a concentration and diligence discipline that caps single-name exposure and sets the depth of work each position requires. The fourth is a governance routine that specifies how the holdings are marked, reviewed and exited over the life of the sleeve.

3.1a What the Framework Does and Does Not Do

It is important to be clear about scope. The framework is a tool for reasoning about how a family office should size, structure and govern an allocation to pre-IPO secondaries. It is not a model that produces a recommended allocation, a target return or a valuation for any company, and it must not be read as doing so. It does not name companies, funds, platforms or managers, and it does not assess any actual opportunity. It is organised around the trade-offs that any family office faces, the trade-off between expected return and liquidity, between concentration and diversification, between control and cost, and between the upside of a single name and the damage its loss would do, rather than around a single optimum, because the right answer depends on facts the framework deliberately leaves open: the family office's wealth, its liquidity needs, its tolerance for loss, its expertise and its governance.

The choice of an analytical, illustrative method rather than an empirical one is deliberate and has two reasons. First, the data that would be needed to calibrate a pre-IPO secondary sleeve empirically, position-level returns, holding periods, discounts and impairment rates across a representative universe, are sparse, proprietary and unrepresentative, and any apparent precision drawn from them would be false. Second, the aim is to equip a family office to reason about its own situation, which a transparent framework does better than a fitted model whose assumptions are hidden. The cost of the choice is that the framework yields structure and relationships, not numbers to act on.

3.2 Assumptions and Their Justification

The stylised inputs used in the figures, the sizing bands by risk profile, the shape of the return dispersion, the relationship between entry price and expected return, the positioning of the access routes, and the swing factors in the sensitivity analysis, are indicative and illustrative. They are chosen to be plausible and to make the relationships visible, and they are held constant across the figures so that the framework is internally consistent. They are not forecasts, not sourced data and not the parameters of any actual sleeve. Table 1 records the principal base-case assumptions; Appendix A gives the full set.

Table 1. Base-Case Assumptions for the Sleeve-Design Framework

Parameter	Illustrative base case	Role in the framework
Sleeve target weight (balanced)	around 2 to 3% of total portfolio	Anchors the sizing ladder
Sleeve ceiling (aggressive)	high single-digit % of portfolio	Hard cap on the allocation
Single-name concentration cap	a small fraction of the sleeve	Limits idiosyncratic loss
Number of positions in a diversified sleeve	several to many, or one diversified fund	Captures the skew
Typical holding horizon	multi-year, uncertain	Drives the liquidity test
Entry reference	discount or premium to last round	Read with information quality
Access routes considered	five, listed to direct	Spans the cost-control trade-off

Indicative, illustrative calibrations used to make the structure legible; not forecasts, sourced data or the parameters of any actual allocation.

4. THE FRAMEWORK: SIZING, ACCESS, CONCENTRATION AND GOVERNANCE

This section presents the framework in the order of the propositions. It first establishes why the sleeve is the right unit of decision, then works through the four levers, sizing, access, concentration and liquidity, and governance, with the support of the illustrative figures. The aim throughout is to connect each design choice to the specific risks that distinguish pre-IPO secondaries from the public equities and the primary venture investments they sit between.

4.0a Why the Sleeve Is the Right Unit of Decision

Before sizing anything, it is worth establishing the proposition that motivates the whole framework. A pre-IPO secondary opportunity almost always presents itself as a single decision: a specific company, a specific price, a specific window. The instinct is to evaluate it on its own merits, and the better the company, the stronger the instinct. But evaluating the position in isolation answers the wrong question. The return and the risk of a single late-stage private company are both extreme and uncertain; the position could return several times the money or could be impaired entirely, and which of these happens is largely outside the buyer's control once the purchase is made. What the family office can control is how much of its wealth is exposed to that outcome, how many such outcomes it spreads its capital across, how it accesses them and how it governs them. These are sleeve-level decisions, and they dominate the position-level decision in their effect on the portfolio.

The same proposition carries a warning that runs through the analysis. The greatest danger in pre-IPO secondaries is not buying a bad company; it is sizing any company, good or bad, so large that its failure damages the whole portfolio, or accumulating so many positions through opportunistic enthusiasm that the sleeve quietly grows beyond the allocation the family office would have chosen deliberately. Both failures are failures of sleeve discipline, and both are avoided by deciding the size, the limits and the rules in advance.

4.1 Sizing the Sleeve

The first lever is size. Proposition 2 holds that the size of the sleeve is bounded by the family office's true liquidity tolerance and by the damage a total loss would do, not by conviction in any single name. Figure 2 sets out an illustrative sizing ladder that links the sleeve weight to the family office's risk profile.

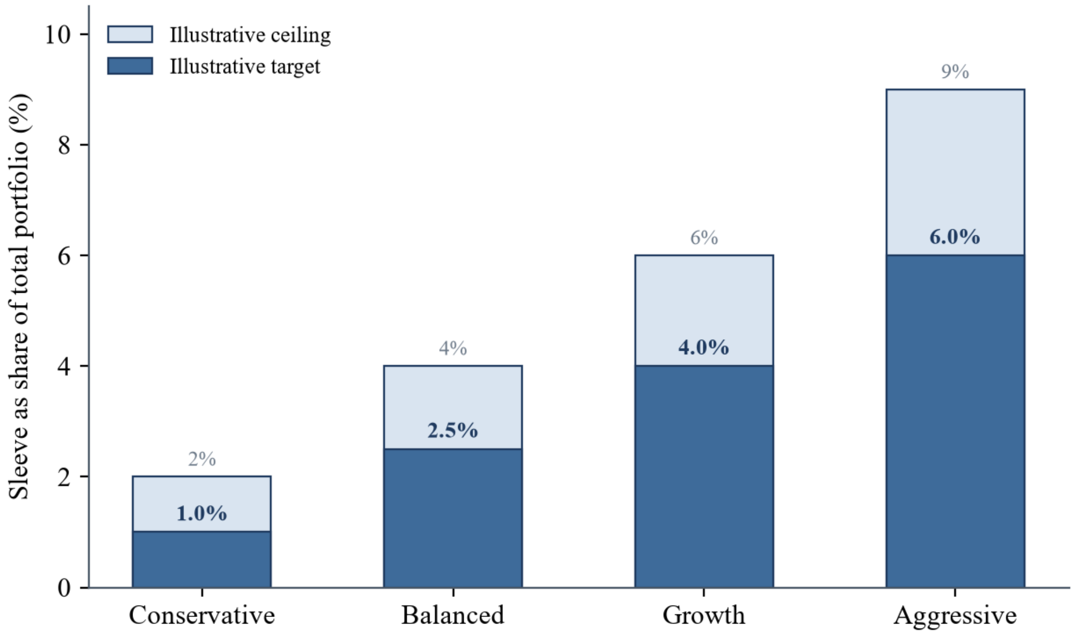

Figure 2. Illustrative Sleeve-Sizing Ladder by Risk Profile

Illustrative target weights and ceilings as a share of the total portfolio, rising from conservative to aggressive profiles. The bands are stylised; the right figure depends on the family office's own liquidity and loss tolerance.

The figure supports Proposition 2 by showing the sleeve as a small and bounded share of the total portfolio even for the most aggressive profile. The logic of the ladder is that the sleeve weight should rise with the family office's tolerance for illiquidity and loss and fall with its need for liquidity and its sensitivity to drawdown. A conservative family office with near-term spending needs or external commitments should hold a small sleeve or none; a wealthy, long-horizon family office with no forced-sale risk and genuine expertise can hold more. The key discipline is the ceiling: a hard cap, decided in advance, that the sleeve may not exceed however attractive the next opportunity appears. The ceiling protects the portfolio from the slow drift that opportunistic accumulation produces.

Three inputs govern where on the ladder a family office should sit. The first is liquidity tolerance: the share of wealth the family can lock away for years without affecting its spending, its other commitments or its peace of mind. The illiquidity literature is unambiguous that this is the binding constraint, because an investor forced to sell an illiquid asset destroys its own return. The second is loss tolerance: the size of total loss on the sleeve the family could absorb without material harm to its wealth or its plans. Because individual positions can be impaired entirely, the sleeve should be sized so that even a severe outcome is survivable. The third is information advantage: the extent to which the family office genuinely sees, prices and diligences these positions better than the market. A real advantage justifies a larger sleeve; the mere ability to access a famous name does not, and confusing access with advantage is among the most common and costly errors.

A practical way to set the size is to begin from the loss the family office could absorb and work backwards. If a total impairment of the sleeve would be painful but survivable, the sleeve is sized about right; if it would be catastrophic, the sleeve is too large, however confident the family is in its positions. This backward test is more robust than forward projections of return, because the loss is more certain in its effect than the gain is in its arrival.

4.1a A Worked Illustrative Sizing Example

It is worth working the sizing logic through a single illustrative case to show how the levers combine in practice. The numbers below are stylised and chosen only to make the reasoning concrete; they are not advice and not a recommendation. Consider a hypothetical family office with a long horizon, no near-term spending call on the relevant capital and a small in-house investment function. It judges, honestly, that it has no special edge in pricing individual late-stage companies but values the access that the asset class offers. Applying the ladder of Figure 2, it places itself between the balanced and growth profiles and adopts an illustrative target of around three per cent of the total portfolio, with a ceiling a little above that.

It then applies the backward loss test. If the entire sleeve were impaired, the family would lose around three per cent of its wealth, a painful but survivable outcome that would not disturb its spending or its other commitments. Had the same test produced a loss the family could not absorb, the target would have been too high and would have been cut until the worst case became survivable. Because the family lacks a pricing edge and a deep investment team, it chooses a diversified fund route rather than direct single-name purchases, accepting a fee layer in exchange for diversification and professional diligence. Within that route it still sets a look-through concentration preference, so that no single underlying company dominates the sleeve, and it commits to building the position over several years rather than in one window. The example shows the levers working together: the size follows from tolerance, the route follows from capacity, and the concentration and pacing follow from the dispersion of outcomes. None of these choices depends on a view about any particular company.

4.2 The Return Profile the Sleeve Must Respect

Sizing cannot be separated from the shape of the outcomes the sleeve is exposed to. Proposition 4 rests on the central stylised fact of late-stage private investing: returns to individual positions are highly dispersed and skewed. Figure 3 illustrates the shape.

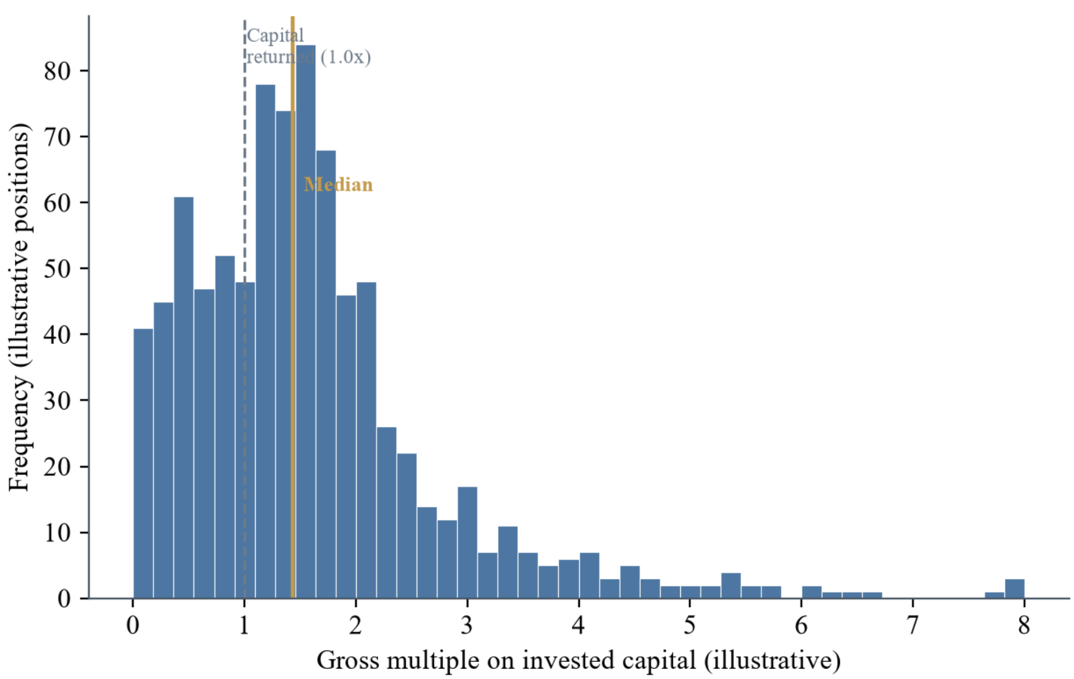

Figure 3. Illustrative Dispersion of Pre-IPO Secondary Outcomes

Illustrative distribution of gross multiples on invested capital across many hypothetical positions. A tail of impairments sits below capital returned, a body of modest outcomes around it, and a long right tail of winners above; the median sits below the mean.

The figure supports Proposition 4 in two ways. First, it shows why a single position is a poor proxy for the asset class. The distribution is wide enough that any one position is as likely to land in the body or the tail as among the winners, and the winners, though they drive the average, are a minority. A family office holding one or two positions is not taking the asset class's return; it is taking a lottery whose expected value is irrelevant to its actual, single draw. Second, the skew shows why diversification within the sleeve is not optional. Only by spreading capital across several positions, or by accessing the asset through a vehicle that holds many, does the sleeve give itself a reasonable chance of capturing the right tail rather than landing in the left.

The practical implication for concentration is direct. The number of positions in the sleeve, or the breadth of the fund through which it is accessed, should be large enough that the sleeve's outcome reflects the distribution rather than a single draw from it. This is the deepest reason a single-name concentration cap matters: it forces the diversification that the shape of the outcomes demands.

4.3 Entry Price and Expected Return

A recurring temptation in secondaries is to treat a discount to the last round as a margin of safety. Figure 4 shows why the relationship between entry price and expected return, though real, must be read with care.

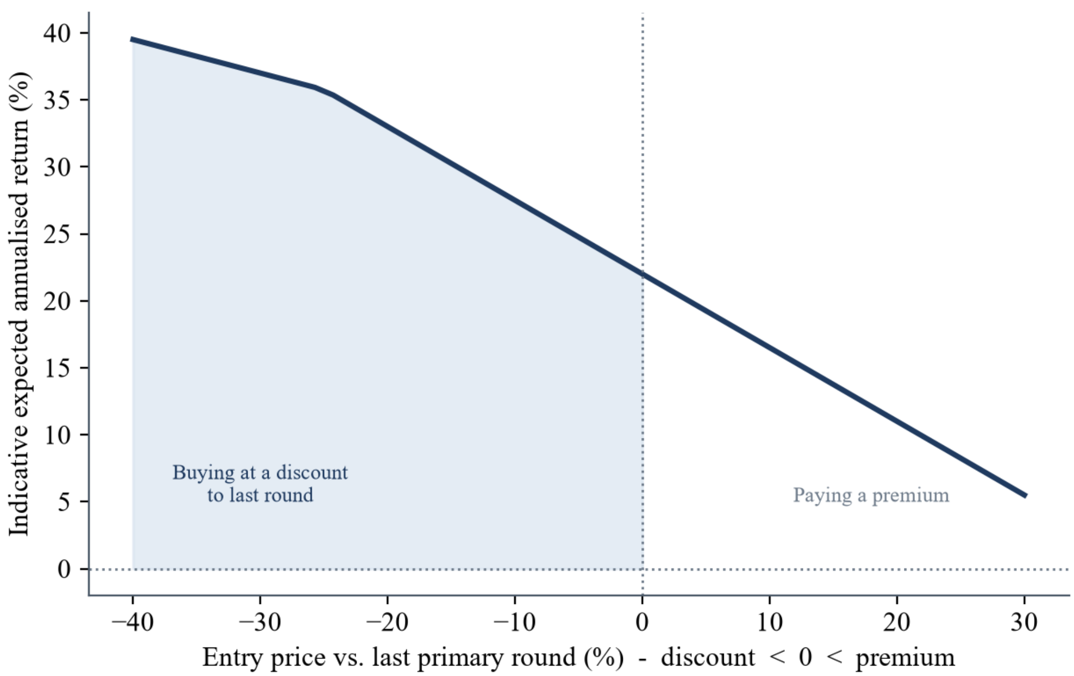

Figure 4. Illustrative Relationship Between Entry Price and Expected Return

Illustrative schematic. Buying at a discount to the last primary round raises the indicative expected return and paying a premium lowers it, but the relationship holds only when the last round is a fair reference and the shares carry the same rights.

The figure illustrates the intuitive point that, all else equal, a lower entry price raises the expected return and a higher one lowers it. But the secondary literature is emphatic that all else is rarely equal. The last primary round, against which the discount or premium is measured, may be stale, struck in different market conditions, or structured with preferences and protections that flatter the headline valuation. The shares a secondary buyer typically acquires, often common stock, may sit behind layers of preferred stock with liquidation preferences that take the first proceeds in any sale. A large apparent discount can therefore compensate for information the seller has and the buyer lacks, or for a structural disadvantage in the shares themselves, rather than representing a bargain.

The discipline the figure implies is to read the discount alongside two other things: the quality of the information behind the last round and the rights attaching to the specific shares on offer. A modest discount on well-understood common shares in a company with recent, representative financials may be more attractive than a large discount on shares whose structure or freshness the buyer cannot verify. The entry price is a lever, but it is the least reliable of the four when used alone.

4.4 Accessing the Sleeve

The third lever is access. Proposition 3 holds that the route into pre-IPO secondaries trades cost, control and diligence burden against directness, and should be matched to the size and ambition of the sleeve. Figure 5 sets out the access spectrum.

Figure 5. The Access Spectrum: Routes Into Pre-IPO Secondaries

Illustrative positioning. Routes run from listed pre-IPO funds and multi-manager secondary funds, through single-name special-purpose vehicles and structured contracts, to direct purchase of shares, trading rising cost, diligence and control against directness and access.

The figure arranges the principal routes by two dimensions: the cost, diligence burden and control they entail, and the directness of the exposure they provide. At one end, a listed pre-IPO fund or investment trust offers the lightest route, instant diversification, daily liquidity and no diligence burden on individual names, at the cost of fees, a diluted exposure and often a premium or discount to net asset value that has nothing to do with the underlying companies. A multi-manager secondary fund offers diversification and professional sourcing and diligence, at the cost of a fee and carry layer and a loss of control over which names enter the portfolio. A single-name special-purpose vehicle or feeder gives concentrated exposure to one company with shared diligence, but concentrates risk and often layers a sponsor's fee. A structured or forward contract can shape the exposure but adds counterparty and structural complexity. At the far end, direct purchase of shares offers the most control and the cleanest economics but demands the most: sourcing, legal work, diligence and the willingness to bear single-name risk alone.

The framework does not declare a winner. The right route depends on the size of the sleeve, the family office's expertise and capacity, and its tolerance for the fee layering that the lighter routes impose. A small sleeve, or one run by a family office without dedicated investment staff, is usually better served by a diversified fund route, which buys diversification and diligence that the family could not assemble itself. A large sleeve run by a capable team can justify the direct route, which removes the fee layer and restores control, but only if the team can genuinely source, price and diligence the positions. The most common error is to choose the direct route for its glamour and its apparent cheapness while lacking the capacity it demands, and so to take concentrated, under-diligenced single-name risk that a fund would have diversified away.

A point worth emphasising about fees is that they compound across the holding period and across the layers of the structure. A fund of secondaries that itself charges fees and carry, holding positions that may already carry a sponsor's economics, can impose a drag that consumes much of the illiquidity premium the sleeve was meant to harvest. The access decision is therefore also a fee decision, and the family office should trace the full stack of fees, its own, the fund's, and any underlying sponsor's, before committing.

4.5 Concentration and Diligence

The fourth design choice within the sleeve is concentration: how much of the sleeve any single name may occupy, and how deeply each position must be diligenced. Proposition 4 requires a cap on single-name exposure, and the dispersion of outcomes shown in Figure 3 is the reason. A sleeve in which one position dominates is not a diversified allocation to an asset class with a skewed return; it is a concentrated bet whose outcome is governed by a single draw from a wide distribution. The concentration cap, decided in advance, forces the breadth that the shape of the outcomes demands and protects the sleeve from the behavioural pull to oversize the name in which the principal has most conviction.

Diligence depth should scale with the directness of the route and the size of the position. A position accessed through a diversified fund inherits the manager's diligence and the diversification of the vehicle, and the family office's work is concentrated on the manager rather than the names. A direct, single-name position bears the full diligence burden alone: the quality and freshness of the company's information, the rights attaching to the specific shares, the identity and motivation of the seller, the structure of the cap table and the realism of the path to liquidity. The asymmetry of information that the secondary literature identifies, the seller usually knows more than the buyer, can only be offset by diligence, and the depth of that diligence is itself a reason to prefer the fund route for a family office that cannot sustain it.

4.6 Liquidity and the Holding Period

The fifth and most underappreciated feature of the sleeve is its liquidity profile over time. Proposition 5 holds that the sleeve must be governed over a long and uncertain horizon, and the shape of that horizon is unlike anything in the public portfolio. Figure 6 illustrates the profile.

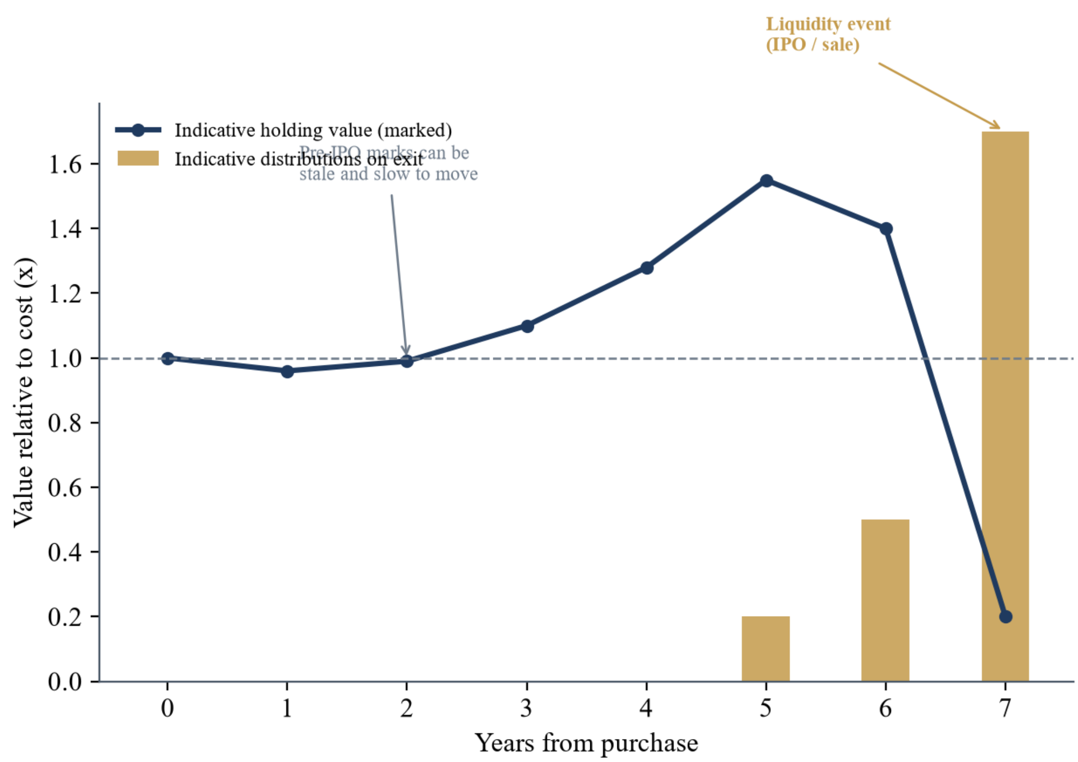

Figure 6. Illustrative Liquidity and Value Profile of a Sleeve Position

Illustrative path. The marked value of a pre-IPO position can move slowly and stay stale for years before a liquidity event, an initial public offering or a sale, finally realises value or loss; distributions arrive late and uncertainly.

The figure makes three points that govern how the sleeve must be sized and supervised. First, the holding period is long and its end is uncertain. A position may be held for many years before any liquidity event, and the event may never come on the terms or the timetable the buyer hoped. The sleeve must therefore be sized as money the family office can genuinely lock away, reinforcing the liquidity test of Section 4.1. Second, the marked value between purchase and exit is unreliable. Pre-IPO positions are marked infrequently, often to the last round or to a manager's estimate, and the marks are smoothed and stale. The apparent stability of the marks understates the true volatility of the position and can lull a family office into a false sense of safety. Third, the realisation, when it comes, is lumpy and binary: a large distribution on a successful listing or sale, or a write-down on a disappointment, with little in between.

The governance implication is that the family office must mark the sleeve honestly, treating the smoothed marks with scepticism and stress-testing the portfolio for the possibility that the true value is well below or above the carried mark. It must also plan for the holding period rather than be surprised by it, neither expecting an early exit nor being forced into a late one, and it must keep the sleeve small enough that its illiquidity never compromises the liquidity of the whole portfolio.

4.7 Putting the Levers Together

The four levers interact, and Figure 7 places their sensitivity in a single view. It shows, in illustrative terms, how the net outcome of the sleeve swings with each of the principal choices.

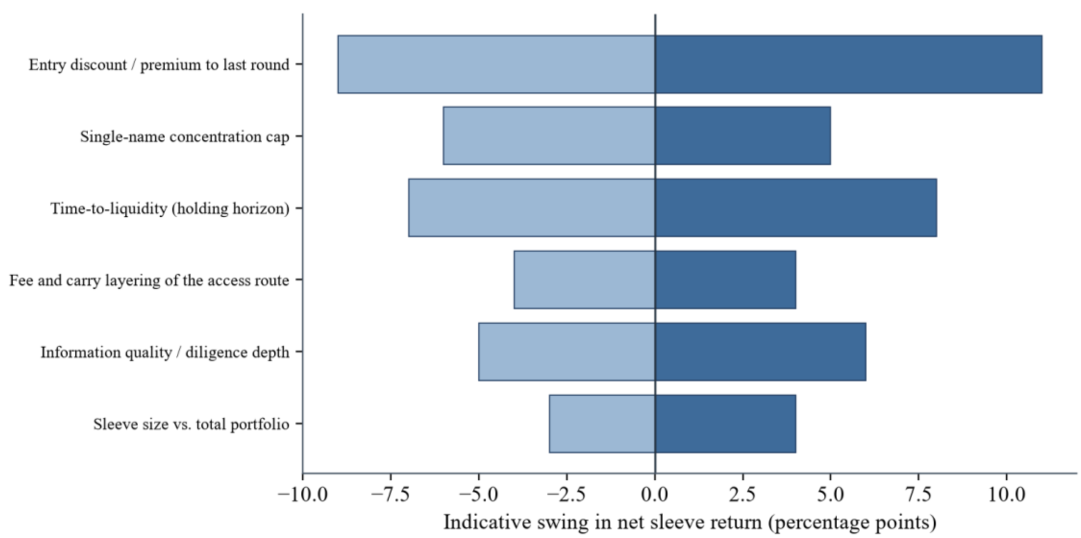

Figure 7. Illustrative Sensitivity of the Sleeve Outcome to Key Choices

Illustrative swing in the net outcome of the sleeve as each driver moves across a plausible range. Entry price, concentration, holding horizon, fee layering, information quality and sleeve size each shift the result; the magnitudes are stylised.

The figure draws the framework together. Several drivers, the entry discount or premium, the holding horizon and the concentration of the sleeve, swing the outcome substantially, which is why they deserve the most disciplined treatment. Others, such as the fee layering of the access route and the depth of information, swing it less in any single period but compound across the long holding period and so matter more than their per-period magnitude suggests. The sizing of the sleeve relative to the portfolio swings the sleeve's own return least, because it is a scaling choice, but it governs the damage the sleeve can do to the whole portfolio, which is the variable that matters most. The lesson is that the family office should spend its discipline where the swings and the compounding are largest, on price, concentration, horizon and fees, while sizing the sleeve conservatively to protect the portfolio from the irreducible uncertainty of the individual outcomes.

4.8 The Sleeve in a Gulf Family-Office Context

The framework is general, but its application in a Gulf family office in 2026 carries some distinctive features that are worth drawing out. The first is the depth and patience of the capital. Many Gulf family offices have long horizons, limited near-term liquidity needs and a genuine tolerance for being locked in, which is precisely the condition the illiquidity literature requires for safely holding pre-IPO secondaries. This depth can justify a sleeve at the upper end of the ladder, but it also creates a temptation: the very absence of a liquidity constraint can encourage oversizing, because the discipline that a spending need would impose is missing. The ceiling and the backward loss test matter more, not less, when liquidity is abundant, because they supply the discipline that circumstances do not.

The second feature is the strong appetite among Gulf principals for direct and co-investment exposure and for proximity to the underlying asset. This appetite is a real strength, bringing engagement and patient ownership, but it interacts with the access decision of Section 4.4. A family office drawn to the direct route for reasons of preference and control should test that preference against its actual capacity to source, price and diligence single-name positions, and should be honest about whether a diversified fund route would serve the sleeve better despite offering less control. The framework does not discourage the direct route; it asks that it be chosen for capacity rather than for instinct.

The third feature is structural and governance preference. Where a family office requires Shariah-aware exposure, the access route and the underlying positions must be screened accordingly, which narrows the available universe and may favour certain structures over others; this is a matter for qualified advice rather than for a general framework. More broadly, Gulf family offices often combine a close, principal-led decision culture with a growing professional investment function, and the sleeve discipline proposed here, rules set in advance, honest marking and scheduled review, is well suited to that combination, because it lets the principal stay close to the decisions while binding the allocation to a structure that survives the pressure of any single opportunity. None of these features changes the mechanics of a pre-IPO secondary; each shapes how the sleeve should be sized, accessed and governed in the Gulf context.

5. IMPLEMENTATION: BUILDING AND RUNNING THE SLEEVE

The framework leads to a process. This section sets out a practical sequence for a family office to design, build and govern a pre-IPO secondary sleeve. The sequence is deliberately ordered so that the bounding decisions, size and rules, are made before any capital is committed, and so that the behavioural discipline the literature recommends is built into the structure rather than relied upon in the moment.

5.1 Step One: Decide Whether to Have a Sleeve at All

The first step is to decide whether a pre-IPO secondary sleeve belongs in the portfolio at all. This is a function of the family office's liquidity tolerance, its loss tolerance and its genuine information advantage, as set out in Section 4.1. A family office with near-term spending needs, external commitments, a low tolerance for illiquidity or no real edge in pricing these positions should conclude that the sleeve is not for it, and there is no shame in that conclusion. The honest default is no sleeve; the sleeve must earn its place.

5.2 Step Two: Set the Size and the Ceiling

If a sleeve is warranted, the next step is to set its target weight and its hard ceiling, using the sizing ladder of Section 4.1 and the backward loss test. The target is where the sleeve sits in normal conditions; the ceiling is the line it may never cross. Both are decided now, in writing, as part of the portfolio policy, so that they constrain every later opportunity rather than being renegotiated under the pressure of a specific deal.

5.3 Step Three: Choose the Access Route

With the size set, the family office chooses how it will access the sleeve, using the access spectrum of Section 4.4. The choice follows from the size and the family office's capacity: a diversified fund route for a small sleeve or a family office without dedicated staff, a direct or single-name route only for a larger sleeve run by a team that can genuinely source, price and diligence positions. The full stack of fees is traced and accepted before any commitment, and the route is chosen for its fit rather than for the glamour of the direct deal.

5.4 Step Four: Set the Concentration and Diligence Rules

Before any position is bought, the family office sets the single-name concentration cap and the diligence standard each position must meet, following Section 4.5. The cap forces the diversification the dispersion of outcomes demands; the diligence standard ensures that the asymmetry of information is offset by work. These rules are part of the sleeve's constitution, decided once and applied to every position, so that no single name, however compelling, escapes them.

5.5 Step Five: Build the Sleeve Gradually

With the rules in place, the family office builds the sleeve towards its target weight over time rather than all at once. Gradual deployment spreads the entry across market conditions and vintages, avoids the concentration in time that deploying in a single window would create, and gives the family office the chance to learn from its early positions before committing the full allocation. The discipline of building towards a target, rather than reacting to whatever is offered, keeps the sleeve aligned with the plan.

5.6 Step Six: Govern, Mark and Review

Once positions are held, the sleeve enters its longest phase: governance over the holding period. The family office marks the holdings honestly, treating smoothed marks with scepticism, and reviews the sleeve on a schedule rather than only when an exit or a crisis forces attention. The review checks that the sleeve remains within its size ceiling and its concentration cap, that the thesis behind each position still holds, and that the liquidity of the whole portfolio is not compromised by the illiquidity of the sleeve. The marking and review routine is the governance discipline that Proposition 5 demands.

5.7 Step Seven: Plan the Exit

Exit planning begins before purchase, not after. For each position and for the sleeve as a whole, the family office forms a view of how value will be realised, through a listing, a sale or a later secondary, and over what horizon, and it holds that view loosely, ready to revise it as events unfold. Because realisations are lumpy and their timing is outside the buyer's control, the family office plans to be patient and avoids any structure or commitment that would force a sale at the wrong moment. The aim is to be a willing rather than a forced seller whenever liquidity finally arrives.

5.8 Common Pitfalls and How the Process Avoids Them

The staged process is designed to avoid the pitfalls that recur when family offices invest in pre-IPO secondaries. The first is sizing by conviction rather than by tolerance, oversizing the name the principal believes in; the size and ceiling, set in advance, prevent it. The second is opportunistic drift, accumulating positions until the sleeve quietly exceeds the intended allocation; the ceiling and the scheduled review catch it. The third is confusing access with edge, paying up for a famous name without any advantage in pricing it; the information-advantage test and the diligence standard expose it. The fourth is anchoring on a stale mark and being surprised by the eventual realisation; honest marking and stress-testing soften it. The fifth is being forced to sell an illiquid position at the wrong time; conservative sizing and exit planning guard against it. Each pitfall is behavioural, and each is met by a rule decided before the pressure arrives.

6. RISKS, CAVEATS AND LIMITATIONS

A framework for a risky activity is only useful if it is honest about its own limits and about the dangers of the territory. This section sets out the principal risks of a pre-IPO secondary sleeve and the limits of the framework offered here.

6.1 The Limits of an Illustrative Framework

The figures and tables in this paper are stylised and illustrative. They are designed to show structure, relationships and trade-offs, not to forecast returns, value any company or recommend any allocation. The sizing bands, the shape of the return dispersion, the positioning of the access routes and the sensitivity swings are indicative calibrations chosen to make the framework legible, and they should not be read as data, forecasts or the parameters of any actual sleeve. A family office applying the framework must replace the illustrative inputs with its own circumstances and with current, specific and properly advised analysis.

6.2 Market and Asset Risks

Pre-IPO secondaries carry the full set of risks of late-stage private companies and add several of their own. The companies may fail, may be worth far less than their last round implied, or may never reach a liquidity event. The shares acquired may sit behind preferences that take the first proceeds. The last round against which price is measured may be stale or structured. The secondary market itself may be thin, so that a position cannot be sold even if the buyer wishes to exit before a listing. And the broad environment for listings and private valuations can shift, lengthening holding periods and depressing realisations across the whole sleeve at once. These risks are real, material and only partly within the family office's control.

6.3 Liquidity and Concentration Risks

The two risks the framework works hardest to manage are liquidity and concentration, and they remain even when the framework is followed. Liquidity risk is the risk that the family office cannot realise a position when it needs to, and that the marks it relies on in the meantime are unreliable. Concentration risk is the risk that a single position, or a cluster of correlated ones, comes to dominate the sleeve or the portfolio and exposes it to a single failure. The sizing ceiling, the concentration cap and the liquidity test reduce these risks but cannot remove them, because the underlying assets are genuinely illiquid and genuinely dispersed.

6.4 Information, Structural and Behavioural Risks

The asymmetry of information between a motivated seller and a later buyer is a permanent feature of the secondary market, and diligence reduces but never eliminates it. Structural risks, the rights attaching to specific shares, the layers of preference above them, and the terms of the access vehicle, can erode value in ways that a headline price conceals. And the behavioural risks the framework is built to discipline, oversizing by conviction, anchoring on price, confusing access with edge, never disappear; they are managed by the rules but pressed constantly by the nature of the opportunities. A family office should treat its own discipline as the asset most in need of protection.

A particular structural risk deserves emphasis because it is easy to overlook. The shares that change hands in a secondary are often subject to transfer restrictions, rights of first refusal held by the company or its early investors, and information rights that the buyer may not inherit. A buyer can find that the position it thought it had acquired is encumbered, that the company can block or pre-empt the transfer, or that it will receive little ongoing information about the business it now owns a stake in. These features do not appear in the headline price and are discoverable only through diligence on the specific shares and the specific company, which is another reason the diligence standard of Section 4.5 is part of the sleeve's constitution rather than an optional extra.

6.5 Scope and Suitability

Finally, this paper is an educational and analytical document for a general audience of principals, family-office executives and advisers. It does not name companies, funds, platforms or managers, does not assess any actual opportunity, and does not constitute investment, legal, tax or accounting advice. Whether a pre-IPO secondary sleeve is suitable for any particular family office, and how it should be sized, structured and governed, depends on facts specific to that family and requires professional advice that no general framework can replace. Nothing in this paper is an offer, a solicitation or a recommendation of any security, strategy or transaction.

7. CONCLUSION

Late-stage private companies now create much of their value before they list, and a deep secondary market lets a family office buy into them before any public offering. The opportunity is real, but it arrives in a form that encourages the wrong question. The right question is not which famous private company to buy; it is how large a standing allocation to such positions the portfolio should carry, how that allocation should be accessed and structured, how it should be diversified and how it should be governed over the long and uncertain years before value is realised.

The argument of this paper has been that the sleeve, not the individual position, is the right unit of decision, and that a sleeve is designed through four levers, size, access, concentration and governance, each set in advance and each tied to the specific risks that distinguish pre-IPO secondaries from public equity and from primary venture. The size is bounded by genuine liquidity and loss tolerance, not by conviction. The access route trades cost and control against directness and should match the sleeve's scale and the family office's capacity. The concentration cap forces the diversification that the dispersion of outcomes demands. And the governance routine marks the holdings honestly and reviews them on a schedule over a holding period whose end the buyer does not control.

The deeper point is that almost everything that goes wrong with private secondaries is a failure of sleeve discipline rather than a failure of stock selection. Positions are oversized by conviction, allowed to drift beyond the intended allocation, accessed through the wrong route, under-diligenced or surprised by stale marks. Each failure is met not by a better forecast but by a rule decided before the pressure of a specific deal arrives. The sleeve is, in the end, a commitment device: a way for a family office to bind its future self to the discipline its present self can see is wise.

Two observations close the argument. The first concerns humility about outcomes. The dispersion of returns to individual late-stage companies is so wide that no family office, however expert, can rely on selecting winners; the most it can do is size and diversify so that the inevitable losses are survivable and the occasional winners matter. The second concerns the primacy of liquidity. The single most important discipline is to hold only what the family can genuinely lock away, because the forced sale of an illiquid position at the wrong moment destroys more value than any selection error. A family office that sizes conservatively, diversifies honestly and governs patiently will, on the evidence of how these positions behave, be better served than one that backs its conviction in a single name, however brilliant that name appears today. The framework offered here is a way of building that discipline into the portfolio before the next opportunity tests it.

REFERENCES

[1] Amihud, Y. and Mendelson, H. (1986). Asset Pricing and the Bid-Ask Spread. Journal of Financial Economics, 17(2), 223-249.

[2] Ang, A. (2014). Asset Management: A Systematic Approach to Factor Investing. New York: Oxford University Press.

[3] Cochrane, J. H. (2005). The Risk and Return of Venture Capital. Journal of Financial Economics, 75(1), 3-52.

[4] Ewens, M. and Farre-Mensa, J. (2020). The Deregulation of the Private Equity Markets and the Decline in IPOs. Review of Financial Studies, 33(12), 5463-5509.

[5] Harris, R. S., Jenkinson, T. and Kaplan, S. N. (2014). Private Equity Performance: What Do We Know? Journal of Finance, 69(5), 1851-1882.

[6] Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. Econometrica, 47(2), 263-291.

[7] Kaplan, S. N. and Schoar, A. (2005). Private Equity Performance: Returns, Persistence and Capital Flows. Journal of Finance, 60(4), 1791-1823.

[8] Lerner, J., Schoar, A. and Wongsunwai, W. (2007). Smart Institutions, Foolish Choices: The Limited Partner Performance Puzzle. Journal of Finance, 62(2), 731-764.

[9] Markowitz, H. (1952). Portfolio Selection. Journal of Finance, 7(1), 77-91.

[10] Metrick, A. and Yasuda, A. (2010). The Economics of Private Equity Funds. Review of Financial Studies, 23(6), 2303-2341.

[11] Nadauld, T. D., Sensoy, B. A., Vorkink, K. and Weisbach, M. S. (2019). The Liquidity Cost of Private Equity Investments: Evidence from Secondary Market Transactions. Journal of Financial Economics, 132(3), 158-181.

[12] Sahlman, W. A. (1990). The Structure and Governance of Venture-Capital Organizations. Journal of Financial Economics, 27(2), 473-521.

[13] Sharpe, W. F. (1964). Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. Journal of Finance, 19(3), 425-442.

[14] Thaler, R. H. (1980). Toward a Positive Theory of Consumer Choice. Journal of Economic Behavior and Organization, 1(1), 39-60.

ABOUT THE AUTHOR

Chennakeshav (CK) Adya is an independent researcher and corporate-finance practitioner with more than twenty-five years of experience across capital markets, private capital and cross-border advisory. His research focuses on how capital is structured, allocated and governed in the Gulf and the wider emerging markets, with a particular interest in private markets, alternatives and the practical questions facing family offices and institutional allocators. He writes to make complex financing and allocation decisions legible to the principals and advisers who must make them.

APPENDIX A. BASE-CASE ASSUMPTIONS

The table below records the stylised assumptions used in the figures and the framework. They are indicative, illustrative and held constant for internal consistency. They are not sourced data, forecasts or the parameters of any actual allocation.

Table A1. Full Base-Case Assumptions

Item	Illustrative value or range	Notes
Conservative sleeve target	around 1% of portfolio	Near-term liquidity needs
Balanced sleeve target	around 2 to 3% of portfolio	Long horizon, some illiquidity tolerance
Growth sleeve target	around 4% of portfolio	High illiquidity tolerance
Aggressive sleeve target	around 6% of portfolio	Wealthy, expert, no forced-sale risk
Sleeve ceiling (aggressive)	high single digits %	Hard cap, never breached
Single-name cap	a small fraction of the sleeve	Forces diversification
Return dispersion	skewed, long right tail	Minority of winners drive the mean
Entry reference	discount or premium to last round	Read with information and rights
Access routes	five, listed to direct	Cost and control trade-off
Holding horizon	multi-year, uncertain end	Drives the liquidity test
Realisation pattern	lumpy, binary, late	Listing, sale or write-down

Illustrative assumptions; see Section 3. Not sourced data, forecasts or the parameters of any actual allocation.

APPENDIX B. SLEEVE-CONSTRUCTION CHECKLIST

The following checklist distils the framework into the questions a family office should answer before and while running a pre-IPO secondary sleeve.

- Have we decided, honestly, that a sleeve belongs in our portfolio at all, given our liquidity and loss tolerance?
- Have we set a target weight and a hard ceiling in advance, and would a total loss of the sleeve be survivable?
- Have we tested our size by the backward loss test rather than by our conviction in any single name?
- Have we distinguished genuine information advantage from mere access to a famous name?
- Have we chosen an access route that fits the sleeve's size and our own capacity, rather than by default?
- Have we traced the full stack of fees, our own, the fund's and any underlying sponsor's, before committing?
- Have we set a single-name concentration cap and a diligence standard that every position must meet?
- Have we planned to build the sleeve gradually across vintages rather than in a single window?
- Do we mark the holdings honestly and review the sleeve on a schedule, treating stale marks with scepticism?

- Have we planned each exit in advance and arranged our affairs so that we are never a forced seller?
- Have we taken qualified legal, tax and investment advice specific to our circumstances?

APPENDIX C. GLOSSARY OF KEY TERMS

The following glossary defines the principal terms used in this paper for readers less familiar with private markets and secondaries.

Pre-IPO secondary. The purchase of shares in a late-stage private company from an existing holder, before any public listing.

Sleeve. A sized, structured and governed pocket of a portfolio with its own mandate, limits and rules, treated as the unit of decision.

Last round. The most recent primary financing of a company, against which a secondary entry price is often measured as a discount or premium.

Liquidation preference. A right attaching to preferred shares to receive proceeds ahead of common shares in a sale, which can reduce the value of common stock a secondary buyer holds.

Illiquidity premium. The additional return an investor expects for holding an asset that cannot be sold quickly at a fair price.

Special-purpose vehicle. A company or fund formed to hold a single investment, often used to give several investors shared access to one position.

Multi-manager secondary fund. A pooled vehicle that buys secondary positions across many companies, providing diversification and professional diligence for a fee.

Return dispersion. The width of the range of possible outcomes across individual positions, which for late-stage private companies is wide and skewed.

Concentration cap. A pre-set limit on how much of the sleeve any single position may occupy.

Marking. The process of assigning a current value to an illiquid holding, often infrequent, smoothed and stale for private positions.

Liquidity event. An initial public offering, sale or later secondary through which an investor finally realises value from a private holding.

Family office. A private organisation that manages the wealth and affairs of a single family or a small group of families.