

The J-Curve Problem

Managing Liquidity and Pacing in a Family Office Private Markets Programme

Authored by CK Adya

Independent Researcher, London Business School

May 2026

Abstract

A private markets programme draws capital before it returns it. In the early years a fund calls capital and charges fees while its investments are still maturing, so the cumulative cash flow is negative before it turns positive, tracing the shape of a letter J. For a family office building a private markets programme, the J-curve is the central liquidity challenge: it must fund capital calls that are uneven and partly unpredictable, while distributions lag, and it must do so without holding so much idle cash that it drags the return or so little that it is caught short. This paper sets out a framework for managing the J-curve, the liquidity, and the commitment pacing of a family office private markets programme. It explains the J-curve and its causes, analyses the cash-flow profile of calls and distributions, develops a commitment-pacing model that smooths the programme across vintages, and addresses the denominator effect and allocation drift that complicate the build. It treats the liquidity buffers and over-commitment strategies that family offices use, the tools, including net-asset-value facilities and secondaries, that bridge and manage liquidity, the modelling that forecasts the cash flows, and the considerations specific to the GCC. The analysis finds that the J-curve is managed not by avoiding it but by pacing commitments across vintages to smooth it, sizing liquidity buffers to meet calls through a stress, and using financing tools to bridge, and that the family offices that manage it well build their programmes smoothly to target while those that manage it poorly face liquidity strain or allocation drift. Three family-office case studies, a sensitivity analysis, an international comparison and an implementation roadmap support the framework.

Keywords: Commitment pacing, denominator effect, family office, J-curve, liquidity, private markets, vintage diversification

1. Introduction

A private markets programme has an awkward shape in its early years. When a family office commits to a private markets fund, the fund does not take all the capital at once; it calls the capital over several years as it makes investments, and it charges fees from the outset, so the family office pays out capital and fees before the investments mature and return cash. The cumulative cash flow is therefore negative in the early years, reaching a trough, before the investments mature and distribute, turning the cash flow positive and tracing the shape of a letter J. This J-curve is the central liquidity challenge of building a private markets programme.

The J-curve matters because a family office must fund the capital calls when they come, and the calls are uneven and partly unpredictable, while the distributions that would offset them lag in the early years. A family office that has not planned for this can be caught short, unable to meet a call without selling other assets or scrambling for liquidity, or, at the other extreme, can hold so much idle cash against the calls that it drags the return of the whole portfolio. Managing the J-curve, the

liquidity and the pacing of the programme, is therefore essential to building a private markets allocation without strain or drag.


This paper sets out a framework for managing the J-curve. The central argument is that the J-curve is managed not by avoiding it, which is impossible, but by pacing commitments across vintages to smooth it, by sizing liquidity buffers to meet calls through a stress, and by using financing tools to bridge the timing gaps. A family office that paces its commitments steadily, holds appropriate buffers, and uses the available tools builds its programme smoothly to target, while one that commits in bursts, holds inappropriate buffers, or lacks the tools faces liquidity strain or allocation drift. The J-curve is a manageable challenge, and managing it well is part of building a private markets programme.

The paper complements the companion papers on the allocation shift, on building a private credit allocation, and on direct versus fund investing, in which the J-curve and liquidity management featured. Here the focus is specifically on the liquidity and pacing of the programme, the practical management of the cash flows that building a private markets allocation requires. The figures are modelled to illustrate the framework and are indicative. The paper proceeds from the J-curve explained (Section 2), through the conceptual foundations (Section 3), the cash-flow profile (Section 4), the commitment-pacing model (Section 5), sizing commitments to target (Section 6), the denominator effect and drift (Section 7), vintage diversification (Section 8), liquidity buffers (Section 9), over-commitment strategies (Section 10), the tools (Section 11), and cash-flow modelling (Section 12), before turning to risk, regional considerations, the fund perspective, cases, sensitivity, and implementation.

2. The J-Curve Explained

The J-curve is the shape of the cumulative net cash flow of a private markets fund over its life, illustrated in Figure 1. In the early years, the fund calls capital from its investors to make investments, and it charges management fees, so the investor cumulative cash flow is negative and falling, reaching a trough. As the investments mature and the fund distributes the returns, the cumulative cash flow rises, crosses zero, and turns positive, tracing the J shape. The depth of the trough and the time to recover define the liquidity challenge the J-curve poses.

Figure 1. The J-Curve: Cumulative Net Cash Flow of a Private Markets Fund



The J-curve arises from the structure of private markets funds, which call capital to invest, charge fees from the outset, and distribute only as the investments mature, typically over a fund life of a decade or more. The early negative cash flow reflects the capital called and the fees charged before the investments have matured to distribute, and the later positive cash flow reflects the distributions as the investments are realised. The J-curve is therefore inherent to the structure of private markets funds, not a defect, and it is a feature every private markets investor must manage.

The J-curve poses a liquidity challenge because the investor must fund the early negative cash flow, the capital calls net of the lagging distributions, from its other resources, and it must do so when the calls come, which is partly unpredictable. An investor that has not planned for the J-curve can be caught short when a call comes, while one that holds excessive liquidity against the calls drags its return. The J-curve challenge is therefore to fund the early negative cash flow without strain or excessive drag, which the commitment pacing and liquidity management examined in this paper address.

The J-curve is deepest and most challenging for a single fund or a single vintage, where the whole commitment follows one J-curve, and it is smoothed by diversifying across vintages, as Figure 2 in the next section shows and as the framework develops. A family office building a programme across many funds and vintages can smooth the aggregate J-curve, reducing the depth of the trough and the liquidity challenge, while one concentrating its commitments in a single vintage faces the full depth of the single J-curve. Managing the J-curve therefore begins with diversifying across vintages, which the commitment pacing achieves.

3. Conceptual Foundations

The J-curve and its management rest on several conceptual foundations. The first is the distinction between committed and invested capital. When a family office commits to a fund, it commits a sum that the fund will call over time, so the committed capital exceeds the invested capital at any moment, and the family office must manage the gap, funding the calls as they come against the commitment. The distinction between committed and invested capital is central to the

liquidity management, because the family office must plan to fund its commitments as they are called, not just its current invested capital.

The second foundation is the vintage, the year in which a fund begins investing, and its role in diversification. Each vintage follows its own J-curve, and a programme concentrated in a single vintage follows one deep J-curve, while a programme diversified across vintages has overlapping J-curves that partly offset, as later vintages call capital while earlier vintages distribute. Diversifying across vintages therefore smooths the aggregate cash flow, reducing the depth of the J-curve, which is the central mechanism for managing it, and it is achieved by pacing commitments across years.


The third foundation is the commitment pacing, the rate at which the family office commits to new funds, which determines how the programme builds and how the J-curves overlap. Steady pacing, committing a portion of the target each year, builds the programme across vintages, smoothing the cash flow, while bursty pacing, committing in concentrated bursts, builds the programme in concentrated vintages, deepening the J-curve. The pacing is the lever the family office controls to manage the J-curve, and pacing steadily across vintages is the central technique for smoothing it.

The fourth foundation is the denominator effect and allocation drift, the tendency of the private markets allocation to drift from its target as the programme builds and as the portfolio value changes. Because committed capital is called over time, the invested allocation lags the target while the programme builds, and because private markets valuations move differently from public markets, the allocation can drift as the portfolio value changes, examined in Section 7. The denominator effect and drift complicate the build, requiring the family office to manage its commitments to reach and maintain its target allocation, which the pacing model addresses.

4. The Cash-Flow Profile

Managing the J-curve requires understanding the cash-flow profile of a private markets programme, the timing of the capital calls and the distributions, illustrated in Figure 4. The capital calls are concentrated in the early years, as the funds invest, and decline as the investment periods end, while the distributions are minimal early and rise as the investments mature, crossing the calls and turning the net cash flow positive. The profile of calls and distributions defines the liquidity the family office must provide and the timing of the J-curve.

Figure 4. Capital Calls and Distributions Over a Programme



The capital calls are partly unpredictable, because the funds call capital as they find investments, on their own timetable, not the investor, so the family office cannot know precisely when calls will come, only the broad profile. This unpredictability is part of the liquidity challenge, because the family office must be able to fund calls when they come, not just on an expected schedule, which requires a liquidity buffer and the flexibility to meet calls on the funds timetable. The unpredictability of the calls is therefore a feature the liquidity management must accommodate, through buffers and flexibility.

The distributions are also partly unpredictable, depending on when the funds realise their investments, which depends on the investments performance and the market conditions for exits. In a downturn, distributions can slow as exits are delayed, extending the J-curve and the period of negative cash flow, which is when the liquidity challenge is greatest. The family office must therefore plan for the possibility that distributions slow, extending the period it must fund the calls from its other resources, which is part of the stress the liquidity buffer must cover, examined in Section 9.


The cash-flow profile, with its uneven, partly-unpredictable calls and distributions, is what the family office must manage, providing the liquidity to fund the calls through the period before distributions catch up, including through a stress when distributions slow. Understanding the profile, modelling it, and planning the liquidity to fund it is the practical core of managing the J-curve, which the pacing model, the liquidity buffers, and the cash-flow modelling examined in this paper address. The cash-flow profile is the challenge, and managing it is the framework purpose.

5. The Commitment-Pacing Model

The commitment-pacing model is the central tool for managing the J-curve, smoothing the programme by committing steadily across vintages, illustrated in Figure 3. By committing a portion of the target allocation each year, the family office builds the programme across vintages, so that as the programme matures, earlier vintages distribute while later vintages call, the J-curves

overlapping and partly offsetting to smooth the aggregate cash flow. The pacing model determines how much to commit each year to build the programme to target while smoothing the cash flow.

Figure 3. Commitment Pacing Across Vintages



The pacing model accounts for the fact that committed capital is called and returned over time, so the family office must commit more than its target allocation over the build period to reach and maintain the target invested allocation, because at any moment some commitments are not yet called and some capital has been returned. The model calculates the commitment pace that builds the invested allocation to target and maintains it, accounting for the calls, the distributions, and the re-investment of distributions into new commitments, which sustains the programme at target once built.


Once the programme reaches its target, the pacing model sustains it by committing each year an amount that replaces the capital being returned by maturing vintages, maintaining the target allocation as a rolling programme. The mature programme is self-sustaining in its pacing: the distributions from maturing vintages fund new commitments, and the steady pacing maintains the target allocation and the smooth cash flow. The pacing model therefore both builds the programme to target and sustains it, and it is the central tool for managing the J-curve across the programme life.

The pacing model should be adjusted for the family office liquidity position and risk tolerance, pacing faster to reach target sooner where liquidity is ample, more cautiously where liquidity is tight, as the decision framework in Figure 8 illustrates. A family office with ample liquidity can pace faster, even over-committing as examined in Section 10, to reach target sooner, while one with tight liquidity should pace cautiously to protect its liquidity. The pacing model is therefore tailored to the family office liquidity position, balancing the speed of reaching target against the liquidity risk, which the framework structures.

6. Sizing Commitments to Reach Target

A particular challenge in the build is that committing exactly the target allocation does not reach the target invested allocation, because of the gap between committed and invested capital. Figure 5 illustrates the under-commitment problem: a family office that commits only its target allocation finds its invested allocation falling below target, because some commitments are not yet called and distributions are returning capital, so the invested allocation lags the commitments. To reach the target invested allocation, the family office must commit more than the target, accounting for the uncalled commitments and the distributions.

Figure 5. Allocation Drift: Under-Commitment versus Target



The over-commitment needed to reach target depends on the call and distribution profile, and the pacing model calculates it, determining the commitment pace that builds the invested allocation to target despite the gap between committed and invested capital. A family office that does not account for this gap, committing only its target and finding its invested allocation lagging, under-builds its programme and fails to reach its target allocation, a common error the framework warns against. Sizing the commitments to reach target, accounting for the gap, is therefore essential to building the programme to its intended allocation.

The sizing must also account for the re-investment of distributions, because once the programme matures and distributes, the distributions can be re-invested into new commitments, sustaining the programme, so the commitment pace to maintain target falls once the programme is built. The mature programme requires less new commitment to maintain target, because the distributions fund part of the new commitments, which the pacing model accounts for in sizing the commitments across the programme life. The sizing therefore changes over the programme life, higher to build, lower to maintain, which the model calculates.

Sizing the commitments correctly, to build the invested allocation to target and maintain it, is therefore a calculation the family office must perform, accounting for the calls, distributions, and re-investment, which the pacing model provides. A family office that sizes its commitments correctly builds its programme to its intended allocation and maintains it, while one that does not under-builds, failing to reach target, or over-builds, exceeding target and risking liquidity strain.

The sizing, performed through the pacing model, is therefore central to building the programme to its intended allocation without under- or over-shooting.

7. The Denominator Effect and Allocation Drift

The denominator effect is a particular complication in managing the private markets allocation: the allocation, measured as a share of the total portfolio, can drift as the total portfolio value changes, independently of the private markets commitments. If the public markets portion of the portfolio falls in value, the private markets allocation rises as a share, because the denominator, the total portfolio, has fallen, even though the private markets value has not changed. This denominator effect can push the allocation above target in a public-market downturn, when the family office may least want to be over-allocated to illiquid private markets.

The denominator effect is most acute in a downturn, when public markets fall, pushing the private markets allocation above target, while private markets valuations, which lag and are smoothed, may not have fallen as much, exaggerating the effect. A family office can find itself over-allocated to illiquid private markets in a downturn, when it may need liquidity and when committing further to private markets is least attractive, a difficult position. The denominator effect therefore complicates the allocation management, and a family office should anticipate it, recognising that its private markets allocation will rise in a public-market downturn through no action of its own.


Managing the denominator effect requires the family office to take a measured view of its allocation, not over-reacting to the denominator-driven drift, and to plan its pacing and liquidity to accommodate it. A family office that over-reacts to the denominator effect, slashing its private markets pacing in a downturn when the allocation appears high, may under-commit in vintages that prove attractive, while one that takes a measured view, recognising the drift as partly a denominator artifact, maintains a more consistent pacing. The denominator effect should therefore be understood and accommodated, not over-reacted to, in managing the allocation through the cycle.

The allocation drift, from both the build dynamics and the denominator effect, means the family office must manage its allocation actively, adjusting its pacing to build and maintain the target while accommodating the drift. A family office that manages the drift, pacing to build the allocation while accommodating the denominator effect, maintains a sensible allocation through the cycle, while one that ignores the drift may under-build or over-react. The active management of the allocation, accommodating the drift, is therefore part of managing the programme, alongside the pacing and the liquidity management, and it requires the family office to monitor and adjust its allocation through the cycle.

8. Vintage Diversification

Vintage diversification is the central technique for smoothing the J-curve, illustrated in Figure 2, which compares the deep J-curve of a single vintage with the smoothed cash flow of a programme diversified across vintages. By committing across vintages, the family office builds overlapping J-curves that partly offset, as later vintages call while earlier vintages distribute, smoothing the aggregate cash flow and reducing the depth of the trough. The vintage diversification both smooths the liquidity challenge and diversifies the programme across the conditions of different years.

Figure 2. Single Vintage versus Diversified Vintages



Vintage diversification also diversifies the programme across the credit and market cycle, because the funds of different vintages invest in different conditions, so the programme is not concentrated in the investments of a single year, which may prove a poor vintage. A programme diversified across vintages spreads its investments across the cycle, reducing the risk that the whole programme is concentrated in a poor vintage, which is a genuine return benefit beyond the liquidity smoothing. The vintage diversification therefore both smooths the liquidity and diversifies the return across the cycle, a double benefit.

Achieving vintage diversification requires the steady pacing examined earlier, committing across years to build the programme across vintages, rather than committing in concentrated bursts. A family office that paces steadily builds the vintage diversification that smooths the J-curve and diversifies the return, while one that commits in bursts concentrates its vintages, deepening the J-curve and concentrating the return-cycle risk. The vintage diversification is therefore achieved through the pacing, and it is a central reason the steady pacing across vintages is the framework central technique for managing the programme.


The vintage diversification should span enough vintages to smooth the J-curve and diversify the cycle meaningfully, which the pacing across several years achieves. A programme paced across, say, four or five vintages achieves a meaningful smoothing and diversification, while one concentrated in one or two vintages does not. The family office should therefore pace its build across enough vintages to achieve the diversification, which the pacing model plans, and should sustain the vintage diversification in the mature programme by continuing to commit across vintages. The vintage diversification, achieved through pacing, is the foundation of a smooth, diversified programme.

9. Liquidity Buffers

Managing the J-curve requires the family office to hold liquidity buffers to meet the capital calls, including through a stress, illustrated in Figure 6. The family office should hold a buffer of cash and liquid assets sufficient to meet the expected calls, plus a reserve for the unpredictability of the calls, plus a stress buffer for the possibility that distributions slow and calls continue in a

downturn. The buffers ensure the family office can meet its calls without being forced to sell other assets or scramble for liquidity, which is the core of managing the J-curve liquidity.

Figure 6. Liquidity Buffers for a Private Markets Programme



The buffers must be sized to balance two risks: too small a buffer risks being caught short when calls come, particularly in a stress when distributions slow, while too large a buffer holds excessive idle cash that drags the return of the whole portfolio. The family office should size the buffers to meet the expected and stressed calls without excessive idle cash, which requires modelling the calls and distributions, including under a stress, to determine the buffer needed. The buffer sizing is therefore a balance, informed by the cash-flow modelling examined in Section 12, between the liquidity risk and the return drag.

The stress buffer is particularly important because the J-curve liquidity challenge is greatest in a downturn, when distributions slow, extending the period of negative cash flow, while the denominator effect may also push the allocation above target. A family office that holds a stress buffer can meet its calls through a downturn without forced selling, while one that holds only an expected-case buffer may be caught short when distributions slow. The stress buffer is therefore the protection against the downturn scenario that most stresses the J-curve liquidity, and sizing it for a realistic stress is essential to managing the programme through the cycle.

The buffers can be supplemented by the financing tools examined in Section 11, which can bridge calls and reduce the buffer needed, allowing the family office to hold less idle cash while retaining the ability to meet calls. A family office that uses financing tools to bridge calls can hold a smaller buffer, reducing the drag, while retaining the ability to meet calls through the tools, which is an efficient way to manage the J-curve liquidity. The buffers and the financing tools therefore work together, the buffers providing the first line of liquidity and the tools the backstop, to meet the calls efficiently without excessive idle cash.

10. Over-Commitment Strategies

A family office with ample liquidity can use an over-commitment strategy to reach and maintain its target allocation more efficiently, committing more than its target allocation to account for the gap between committed and invested capital and to keep the programme fully invested. Because committed capital is called over time and distributions return capital, committing exactly the target leaves the programme under-invested, as Section 6 showed, so committing somewhat more, over-committing, keeps the programme closer to its target invested allocation. The over-commitment is a deliberate strategy to manage the gap between committed and invested capital.

The over-commitment must be sized carefully, because committing too much risks a liquidity strain if many funds call capital at once, particularly in a downturn when distributions slow, leaving the family office unable to meet the calls. The over-commitment therefore balances the benefit of keeping the programme fully invested against the risk of a liquidity strain if calls coincide, and it should be sized to the family office liquidity, with a larger over-commitment for a family office with ample liquidity and little or none for one with tight liquidity. The over-commitment is a tool for the liquidity-rich, not the liquidity-constrained.


The over-commitment strategy is most appropriate for a mature programme with a steady distribution flow that funds part of the new commitments, reducing the net liquidity demand, and for a family office with ample liquidity to meet the calls if distributions slow. A mature, liquidity-rich family office can over-commit to keep its programme fully invested, relying on the distribution flow and its liquidity to meet the calls, while a young programme or a liquidity-constrained family office should not over-commit, because it lacks the distribution flow and the liquidity to meet the calls if they coincide. The over-commitment is therefore a strategy for the mature, liquidity-rich programme.

The over-commitment, used appropriately, keeps the programme fully invested and efficient, capturing the full return of the target allocation, while used inappropriately it risks a liquidity strain, the danger the framework warns against. A family office that over-commits within its liquidity capacity captures the efficiency, while one that over-commits beyond its liquidity capacity risks the strain, particularly in a downturn when calls continue and distributions slow. The over-commitment should therefore be sized to the family office liquidity capacity, including a stress, and a family office should be cautious about over-committing beyond what it can fund through a stress, which is the prudent use of the strategy.

11. Tools: Facilities, Secondaries and Credit Lines

Several tools help a family office manage the J-curve liquidity, bridging calls and managing the cash flow. The first is the net-asset-value facility, examined in a companion paper, which lends against the private markets portfolio to provide liquidity, bridging capital calls and funding new commitments without selling assets or holding excessive cash. Figure 7 illustrates a NAV facility bridging a capital call, providing the liquidity to meet the call alongside the family office own resources. The NAV facility is a powerful tool for managing the J-curve liquidity, reducing the buffer needed and bridging the calls.

Figure 7. Bridging a Capital Call with a NAV Facility



The second tool is the secondary market, which allows the family office to buy or sell fund interests, managing its programme and its liquidity. The family office can buy secondary interests to gain immediate exposure or vintage diversification, reducing the J-curve by buying interests that are already partly through their J-curve, or sell interests to raise liquidity or rebalance, managing its allocation and liquidity. The secondary market is therefore a tool for both building and managing the programme, and for raising liquidity if needed, complementing the primary commitments and the financing tools.

The third tool is the credit line or subscription facility, which provides short-term liquidity to bridge capital calls and smooth the cash flow, allowing the family office to meet calls from the facility and repay it from distributions or other resources. The credit line is a simple bridging tool, reducing the buffer needed and smoothing the cash flow, and it complements the NAV facility and the secondary market in the family office liquidity-management toolkit. The three tools together, the NAV facility, the secondary market, and the credit line, give the family office a range of options for managing the J-curve liquidity.

The tools, used together, allow the family office to manage the J-curve liquidity efficiently, bridging calls, managing the allocation, and raising liquidity when needed, without holding excessive idle cash or being caught short. A family office that builds the capability to use these tools manages its J-curve liquidity more efficiently than one relying only on cash buffers, holding less idle cash while retaining the ability to meet calls and manage the programme. The tools are therefore part of the sophisticated management of the J-curve liquidity, and building the capability to use them is part of building the liquidity-management capability the programme requires.

12. Cash-Flow Modelling and Forecasting

Managing the J-curve requires the family office to model and forecast its cash flows, the calls and distributions across its commitments and vintages, so that it can plan its liquidity and pacing. The

family office should model the expected calls and distributions of each commitment, aggregate them across the programme, and forecast the net cash flow over time, including under a stress where distributions slow. The cash-flow model is the foundation of the liquidity planning, telling the family office how much liquidity it needs and when, so that it can size its buffers and plan its pacing.

The model should be stochastic or scenario-based, recognising the unpredictability of the calls and distributions, modelling a range of outcomes rather than a single forecast, so that the family office can plan for the stress as well as the expected case. A single-point forecast understates the liquidity risk, because the calls and distributions are uncertain, and a family office that plans only for the expected case may be caught short in a stress. The model should therefore capture the range of outcomes, particularly the stress where distributions slow and calls continue, so that the liquidity planning covers the stress, not just the expected case.

The model should be updated as the programme evolves, incorporating the actual calls and distributions and adjusting the forecast, so that the liquidity planning reflects the current state of the programme. A family office that updates its model maintains an accurate forecast of its liquidity needs, adjusting its buffers and pacing as the programme evolves, while one that relies on a static model may find its liquidity planning diverging from reality. The model is therefore a living tool, updated as the programme evolves, providing the current forecast that informs the ongoing liquidity and pacing management.

The cash-flow modelling, well done, gives the family office the foresight to manage its J-curve liquidity proactively, planning its buffers, pacing, and tool usage to meet its calls smoothly through the cycle. A family office that models its cash flows can plan its liquidity proactively, avoiding both strain and excessive drag, while one that does not manages its liquidity reactively, risking being caught short or holding excessive cash. The cash-flow modelling is therefore the analytical foundation of managing the J-curve, and building the capability to model and forecast the cash flows is part of building the liquidity-management capability.

13. Risk Management: Liquidity Stress

The principal risk in managing the J-curve is a liquidity stress, in which capital calls continue while distributions slow, typically in a downturn, leaving the family office struggling to meet its calls. This stress is the scenario the liquidity buffers and tools must cover, and managing it is the core of the J-curve risk management. A family office that plans only for the expected case, with calls and distributions as forecast, may be caught short in the stress, when distributions slow and the net liquidity demand rises, which is why the stress scenario must be planned for explicitly.

The stress is exacerbated by the denominator effect and over-commitment, which can compound in a downturn: public markets fall, pushing the private markets allocation above target through the denominator effect, while distributions slow and calls continue, and any over-commitment adds to the calls, straining the liquidity. A family office that has over-committed and holds a thin buffer can face a severe liquidity strain in this compound stress, forced to sell assets at depressed prices or default on a call, which is the danger the framework warns against. The risk management must therefore plan for this compound stress, sizing the buffer and limiting the over-commitment to survive it.

Managing the stress requires the stress buffer examined in Section 9, sized to meet the calls through a realistic downturn when distributions slow, plus the financing tools examined in Section 11, which provide a backstop liquidity, plus a prudent over-commitment that does not exceed the family office capacity to fund the calls through the stress. A family office that holds a stress buffer, has access to financing tools, and limits its over-commitment can meet its calls through the stress without forced selling, while one that holds a thin buffer, lacks the tools, and has over-committed may be caught short. The combination of buffer, tools, and prudent over-commitment is the stress management.

The stress management also includes contingency planning: knowing in advance how the family office would meet its calls in a severe stress, through the buffer, the tools, asset sales, or other resources, so that it is not improvising in a crisis. A family office that has planned its stress contingencies can act decisively in a stress, drawing on its buffer and tools in a planned sequence, while one that has not may be forced into hasty, costly actions. The contingency planning, alongside the buffer and tools, is part of the stress management, ensuring the family office can meet its calls through even a severe stress without forced selling or default.

14. Considerations Specific to the GCC

GCC family offices have characteristics that shape their J-curve management. The first is liquidity: many regional family offices command substantial liquidity, from operating businesses or accumulated wealth, which gives them the capacity to fund the J-curve and to use over-commitment strategies, as the over-commitment section discussed. A liquidity-rich regional family office can pace its build faster and over-commit to keep its programme fully invested, managing the J-curve from a position of liquidity strength, which is an advantage in building a private markets programme.

The second consideration is the relative youth of many regional family offices private markets programmes, examined in the companion papers, which means many are in the early, J-curve-trough years of their build, facing the deepest liquidity challenge. A regional family office building its programme is in the early years when the J-curve is deepest and the liquidity challenge greatest, and it should pace its build to manage the trough, diversifying across vintages from the start and holding appropriate buffers. The youth of the regional programmes therefore makes the J-curve management particularly relevant, as many are in the challenging early years.

The third consideration is the availability of the financing tools in the region, the NAV facilities, credit lines, and secondary market access, which are developing in the region as its private markets ecosystem matures. A regional family office can increasingly access these tools to manage its J-curve liquidity, though the regional market for them is younger than the international one, and it should build relationships with the providers of these tools as part of its liquidity-management capability. The developing availability of the tools in the region supports the J-curve management, and a regional family office should access them as they develop.

The fourth consideration is the compliant requirement, which applies to the financing tools as to the underlying investments, and which a portion of regional capital requires. A regional family office requiring compliance must ensure its financing tools, the NAV facilities and credit lines, are compliant, as the companion paper on NAV facilities discussed, and compliant versions of these tools are available. The compliant requirement therefore shapes the family office choice of financing tools, favouring compliant structures, which are available in the region, for family

offices that require them. The compliant dimension applies to the J-curve management tools as to the broader programme.

15. The Fund Perspective

Understanding the fund perspective helps a family office anticipate and manage the calls and distributions that drive the J-curve. The fund calls capital as it finds investments, on its own timetable, seeking to deploy the committed capital over its investment period, and it distributes as it realises investments, on a timetable driven by the investments performance and the exit conditions. The fund timetable, not the investor, drives the calls and distributions, which is why they are partly unpredictable from the investor perspective and why the family office must plan for a range of timing.

The fund has an interest in calling capital efficiently, to deploy it and earn returns, and in distributing as it realises investments, to return capital and crystallise its carry, so its incentives broadly align with deploying and returning capital on a sensible timetable. But the fund timetable is driven by the investment opportunities and the exit conditions, which the investor cannot control, so the calls and distributions remain partly unpredictable, and the family office must accommodate the fund timetable rather than expecting the fund to accommodate its liquidity. Understanding the fund incentives and constraints helps the family office anticipate the broad timetable while planning for its unpredictability.

The fund also uses its own liquidity tools, such as subscription lines, that can affect the timing of calls, sometimes delaying calls by funding investments from the subscription line and calling capital later, which can compress the calls when they come. A family office should understand that the fund use of subscription lines can affect the timing and pattern of its calls, sometimes delaying and then compressing them, which affects the family office liquidity planning. The fund liquidity practices therefore affect the investor cash flows, and understanding them helps the family office anticipate the call pattern.

Understanding the fund perspective therefore helps the family office anticipate and plan for the calls and distributions, accommodating the fund timetable while planning for its unpredictability, which is part of managing the J-curve. A family office that understands the fund incentives, constraints, and liquidity practices can anticipate the broad call and distribution pattern, plan its liquidity accordingly, and accommodate the fund timetable, while one that does not may be surprised by the timing of calls and distributions. The understanding of the fund perspective is therefore part of the J-curve management, informing the family office liquidity planning.

16. Case Studies

Three family-office cases illustrate the framework applied to different liquidity positions. The figures are modelled for analytical clarity and are not drawn from any specific family office.

16.1 Case A: the cautious family office

Case A is a family office with a measured liquidity position that paces its private markets build cautiously, committing steadily across vintages without over-committing, and holding ample liquidity buffers. It reaches its target allocation more slowly, over several years, but it manages the J-curve with little liquidity stress, meeting its calls comfortably from its buffers and the smoothed cash flow of its diversified vintages. The case illustrates the cautious approach,

prioritising liquidity safety over speed of reaching target, suiting a family office that values liquidity security.

16.2 Case B: the balanced family office

Case B is a family office with a balanced liquidity position that paces its build steadily across vintages, with a modest over-commitment to reach target efficiently, and holds appropriate buffers supplemented by financing tools. It reaches its target allocation at a reasonable pace, manages the J-curve through its diversified vintages, buffers and tools, and maintains a balanced liquidity position. The case illustrates the balanced approach, reaching target at a reasonable pace while managing the J-curve liquidity with buffers and tools, suiting a family office balancing speed and liquidity safety.

16.3 Case C: the over-committing family office

Case C is a liquidity-rich family office that paces its build faster, over-committing to keep its programme fully invested and reach target sooner, relying on its ample liquidity, its distribution flow, and financing tools to meet the calls. It reaches its target allocation quickly and keeps its programme fully invested, but it bears a higher liquidity stress risk if calls coincide in a downturn, which its ample liquidity and tools are sized to cover. The case illustrates the over-committing approach, reaching target quickly from a position of liquidity strength, suiting a liquidity-rich family office that can bear the higher stress risk.

Figure 8. Time to Target and Liquidity Stress Risk by Approach





Figure 8 compares the three cases on the time to reach target and the liquidity stress risk. The cautious family office reaches target slowly with low stress risk, the over-committing family office reaches target quickly with higher stress risk, and the balanced family office sits between. The comparison illustrates the trade-off the framework structures, between the speed of reaching target and the liquidity stress risk, and shows that the right approach depends on the family office liquidity position and risk tolerance, with the liquidity-rich able to pace faster and the liquidity-constrained pacing cautiously. The cases show the framework producing different approaches for different liquidity positions, each managing the J-curve appropriately.

17. Sensitivity and Scenario Analysis

A tornado analysis identifies the variables that most influence the liquidity resilience of a private markets programme. Figure 9 presents the result.

Figure 9. Sensitivity of Liquidity Resilience to Key Variables



The analysis shows that the over-commitment level and the commitment pace most influence the liquidity resilience, with the distribution speed, the liquidity buffer, and the vintage spread also significant. The over-commitment matters because excessive over-commitment most threatens the liquidity resilience, increasing the calls beyond the family office capacity, while the commitment pace matters because bursty pacing concentrates the calls. The distribution speed matters because slow distributions extend the liquidity demand, and the buffer and vintage spread provide the resilience. The sensitivity confirms that the liquidity resilience depends most on the family office own choices, the over-commitment and pacing, which it controls.

Table 3. Scenario Matrix for Liquidity Resilience

Scenario	Over-commitment	Distributions	Resilience
Resilient	None / modest	Normal	High
Balanced	Modest	Normal	Good
Stressed	High	Slow (downturn)	Strained
Crisis	Excessive	Frozen	At risk of default

Modelled scenarios. The crisis case shows excessive over-commitment with frozen distributions risking default.

The scenario matrix shows the danger the framework warns against: the crisis scenario, excessive over-commitment with frozen distributions in a severe downturn, can put the family office at risk of defaulting on a capital call, a serious consequence that can damage its standing and forfeit its investment. This confirms that the over-commitment must be limited to what the family office can fund through a severe stress, and that the liquidity buffer and tools must cover the stress, because the consequence of being caught short, defaulting on a call, is severe. The matrix reinforces the framework prescription to pace prudently, limit over-commitment to the stress capacity, and hold a stress buffer, because the downside of mismanaging the J-curve liquidity is severe.

18. International Comparison

The J-curve and its management are well understood by sophisticated institutional and family office investors internationally, who have developed the pacing models, the over-commitment strategies, the liquidity buffers, and the tools that the framework describes. The international experience confirms the framework key lessons: that the J-curve is managed by pacing across vintages, that over-commitment keeps the programme invested but must be sized to the liquidity capacity, and that the financing tools and secondary market help manage the liquidity. The international practice validates the framework approach to managing the J-curve.

The international experience also documents the consequences of mismanaging the J-curve, including the cases in past downturns where over-committed investors faced liquidity strains, were forced to sell fund interests at deep discounts in the secondary market, or defaulted on calls, which confirms the danger the framework warns against. The cautionary international cases, of investors caught short by excessive over-commitment and frozen distributions in a downturn, reinforce the framework prescription to pace prudently and size the over-commitment and buffers for a stress. The GCC family offices can learn from both the established practice and the cautionary cases, managing the J-curve with the discipline the framework describes.

19. Building the Liquidity-Management Capability

Managing the J-curve well requires the family office to build a liquidity-management capability, spanning the cash-flow modelling, the pacing planning, the buffer sizing, and the use of the financing tools. This capability, like the manager-selection and portfolio-construction capabilities examined in the companion papers, is part of the broader private markets capability the family office builds as it matures, and it is essential to building a private markets programme without liquidity strain. A family office that builds the liquidity-management capability can build its programme smoothly to target and manage it through the cycle, while one that lacks it may face strain or drift.

The capability begins with the cash-flow modelling examined in Section 12, the analytical foundation that forecasts the liquidity needs, then the pacing planning that builds the programme to target across vintages, then the buffer sizing that ensures the calls can be met through a stress, and the use of the financing tools that bridge and manage the liquidity. Building these components, the modelling, the pacing, the buffers, and the tools, gives the family office the capability to manage the J-curve liquidity proactively, planning and executing its liquidity management rather than reacting to calls as they come.

The capability should be built progressively as the programme grows, starting with the basic cash-flow modelling and pacing for a young programme and developing the over-commitment strategies and tool usage as the programme matures and the family office liquidity-management sophistication grows. A family office building its programme should build its liquidity-management capability alongside it, developing the modelling, pacing, buffers, and tools as the programme grows, so that the capability keeps pace with the programme. The progressive build of the capability, alongside the programme, ensures the family office can manage the J-curve at each stage of the programme growth.

The liquidity-management capability, once built, is a durable asset that allows the family office to manage its private markets programme through cycles, building to target, maintaining the

allocation, and meeting its calls through stresses, which is essential to a successful private markets programme. A family office that builds the capability manages its programme smoothly through the cycle, while one that lacks it may face the strains and drifts the framework warns against. The liquidity-management capability is therefore part of the broader private markets capability the family office builds, and it is essential to building and sustaining a private markets programme, which the framework helps it develop.

20. Common Errors and How to Avoid Them

A recognisable set of errors recurs in managing the J-curve.

- **Bursty commitment.** Committing in concentrated bursts rather than steadily across vintages deepens the J-curve and concentrates the calls and the cycle risk. The remedy is steady pacing across vintages.
- **Under-committing.** Committing only the target allocation, ignoring the gap between committed and invested capital, under-builds the programme. The remedy is to size commitments to reach target, accounting for the gap.
- **Excessive over-commitment.** Over-committing beyond the liquidity capacity risks a strain or default when calls coincide in a downturn. The remedy is to size over-commitment to the stress capacity.
- **Thin buffer.** Holding a thin buffer that does not cover a stress risks being caught short when distributions slow. The remedy is a stress buffer plus financing tools.
- **Over-reacting to drift.** Over-reacting to the denominator-driven allocation drift in a downturn can lead to under-committing in attractive vintages. The remedy is a measured view of the drift.

Each of these errors is avoidable through the disciplined approach the framework encourages: pace steadily across vintages, size commitments to reach target, limit over-commitment to the stress capacity, hold a stress buffer with financing tools, and take a measured view of the drift. The family office that does so builds its programme smoothly to target and manages the J-curve liquidity through the cycle, while the one that does not faces a deep J-curve, under-builds, over-commits, holds a thin buffer, or over-reacts to drift. The discipline is what manages the J-curve from a challenge into a manageable feature of the programme.

21. Implementation Roadmap

1. Model the expected and stressed cash flows, the calls and distributions, across the programme, as the analytical foundation for the liquidity planning.
2. Pace commitments steadily across vintages, building the vintage diversification that smooths the J-curve and diversifies the cycle.
3. Size commitments to reach and maintain the target allocation, accounting for the gap between committed and invested capital and the re-investment of distributions.
4. Hold liquidity buffers sized to meet the calls through a realistic stress, balancing the liquidity risk against the return drag.
5. Use over-commitment only within the liquidity capacity, sized to fund the calls through a stress, and only for a mature, liquidity-rich programme.
6. Build access to the financing tools, NAV facilities, credit lines and the secondary market, to bridge calls and manage the liquidity efficiently.

7. Build the liquidity-management capability, the modelling, pacing, buffers and tools, and manage the programme actively through the cycle.

22. Strategic Perspective: Liquidity as an Enabler

The strategic insight of this paper is that managing the J-curve liquidity well is what enables a family office to build and sustain a private markets programme at all, because a programme that cannot fund its calls cannot be built. The liquidity management is therefore not a peripheral operational matter but the enabler of the whole private markets allocation, and a family office that masters it can build the substantial private markets programme that the allocation shift requires, while one that does not is constrained in how much it can allocate to private markets. The liquidity management enables the allocation, and mastering it is foundational to the private markets programme.

The liquidity management also enables the family office to allocate more to private markets than it otherwise could, because a family office that can manage the J-curve liquidity well, with efficient buffers, prudent over-commitment, and the financing tools, can sustain a larger private markets allocation without strain than one that manages the liquidity poorly and must hold large buffers or limit its allocation. The liquidity-management capability therefore expands the family office capacity for private markets, allowing it to capture more of the private markets opportunity, which is a strategic benefit beyond merely avoiding strain.

The liquidity management connects to the family office broader private markets capability, examined in the companion papers, as the enabler that allows the manager selection, portfolio construction, and co-investment capabilities to be deployed at scale. A family office that has built the liquidity-management capability can deploy its other private markets capabilities at the scale its allocation intends, while one that has not is constrained by its liquidity management, unable to build the programme its other capabilities could support. The liquidity management is therefore the enabler of the broader private markets capability, and building it is foundational to the family office private markets ambitions.

The strategic prize is a family office that can build and sustain a substantial private markets programme, capturing the private markets opportunity at scale, enabled by a liquidity-management capability that funds the J-curve smoothly through the cycle. A family office that masters the liquidity management can pursue its private markets ambitions fully, while one that does not is constrained by the J-curve liquidity challenge, limiting its allocation or facing strain. The liquidity management is therefore the enabler of the family office private markets ambitions, and mastering it, through the framework this paper provides, is foundational to capturing the private markets opportunity the companion papers describe.

23. Conclusion

The J-curve, the early negative cash flow of a private markets programme before distributions catch up, is the central liquidity challenge of building a private markets allocation, and managing it well is essential to building the allocation without strain or drag. This paper has shown that the J-curve is managed not by avoiding it but by pacing commitments across vintages to smooth it, sizing liquidity buffers to meet calls through a stress, using financing tools to bridge the timing gaps, and managing the over-commitment and the denominator-driven drift that complicate the build.

The central conclusions are that vintage diversification through steady pacing is the core technique for smoothing the J-curve; that commitments must be sized to reach target, accounting for the gap between committed and invested capital; that over-commitment keeps the programme invested but must be limited to the stress liquidity capacity; that the liquidity buffers and financing tools must cover a realistic stress; and that the liquidity management is the enabler of the whole private markets programme. The family office that builds the liquidity-management capability and manages the J-curve with the discipline the framework describes can build and sustain a substantial private markets programme through the cycle, and the frameworks in this paper are intended to help it do so.

24. Limitations and Directions for Further Research

This paper is framework-oriented and relies on modelled figures, and its conclusions are directional rather than precise. The cash-flow profiles, pacing models and buffer sizes are calibrated to observable conditions but are not empirical estimates, and they vary across funds, strategies and programmes. The behaviour of calls and distributions in a stress, central to the framework, is uncertain and varies across downturns.

Several extensions would strengthen the analysis. An empirical study of capital call and distribution patterns across funds and strategies would replace the modelled profiles with data. An analysis of how private markets programmes managed their liquidity through past downturns, when distributions slowed, would test the framework stress assumptions. And a study of the use and effectiveness of the financing tools, the NAV facilities, credit lines and secondaries, in managing the J-curve would illuminate their role. Each is a natural subject for a later paper in this series.

Appendix A. Base Case Assumptions

Table 3. Base Case Assumptions

Parameter	Value	Parameter	Value
Fund life	~10 years	J-curve trough	~years 2-3
Investment period	~5 years	Pacing	Steady, 4-5 vintages
Over-commitment (liquidity-rich)	Modest	Liquidity buffer	Expected + stress
Denominator effect	Acute in downturn	Tools	NAV facility, credit line, secondaries
Modelling	Scenario / stochastic	Compliant tools	Available
Binding factor	Liquidity capacity	Currency	USD (peg) core

Modelled parameters used to generate the figures and cases. Not transaction-specific.

Appendix B. Liquidity-Management Checklist

Table 4. Liquidity-Management Checklist

Area	What to do
Cash-flow model	Forecast calls and distributions, expected and stressed
Pacing plan	Commit steadily across vintages to build and maintain target
Commitment sizing	Size to reach target, accounting for the committed-invested gap
Liquidity buffer	Hold expected-call reserve plus a stress buffer
Over-commitment	Limit to the stress liquidity capacity
Financing tools	Arrange NAV facility, credit line, secondary access
Denominator effect	Take a measured view of allocation drift
Stress contingency	Plan how to meet calls in a severe stress
Monitoring	Update the model and adjust as the programme evolves
Capability	Build the modelling, pacing and tool-usage capability

A practical checklist for managing the J-curve liquidity.

Appendix C. Glossary of Terms

Table 5. Glossary of Key Terms

Term	Definition
J-curve	The early negative then positive cumulative cash-flow shape.
Capital call	A fund demand for committed capital.
Distribution	A return of capital and gains from a fund.
Committed capital	Capital pledged to a fund, called over time.
Invested capital	Capital actually called and invested.
Vintage	The year a fund begins investing.
Commitment pacing	Committing across vintages to smooth the programme.
Denominator effect	Allocation drift from changes in total portfolio value.
Over-commitment	Committing above target to keep the programme invested.
NAV facility	A loan against the portfolio to provide liquidity.

Definitions provided for reference.

References

- [1] Brealey R., Myers S., Allen F., Principles of Corporate Finance, 13th Edition, McGraw-Hill Education, 2020.
- [2] Demaria C., Introduction to Private Equity, Debt and Real Assets, 3rd Edition, Wiley, 2020.

- [3] Takahashi D., Alexander S., Illiquid Alternative Asset Fund Modeling, Journal of Portfolio Management, 2002.
- [4] Preqin, Private Capital Cash Flow and Liquidity, Market Overview, 2024.
- [5] Swensen D., Pioneering Portfolio Management, Free Press, 2009.
- [6] Institutional Limited Partners Association, Commitment Pacing and Liquidity Management, ILPA, 2023.
- [7] Cambridge Associates, Managing Private Investment Programs and the J-Curve, 2022.
- [8] Ang A., Asset Management: A Systematic Approach to Factor Investing, Oxford University Press, 2014.
- [9] Bank for International Settlements, Liquidity in Private Markets, BIS Papers, 2024.
- [10] Accounting and Auditing Organization for Islamic Financial Institutions, Shariah Standards, AAOIFI, 2017.

About the Author

Chennakeshav Adya is an independent researcher and corporate finance practitioner with more than twenty years of international experience across business strategy, transformation, and strategic finance. He has originated and led transactions across real estate, technology, and clean energy in the UAE, the United Kingdom, India, Europe, and the Americas, and has advised developers, corporates, family offices, and fund managers on capital raising and structuring. He is a guest speaker at London Business School and writes on startups, capital markets, and the application of technology in finance. The views expressed in this paper are his own and do not constitute investment advice.