

Liquidity Management for Illiquid Gulf Allocations

Funding the capital calls, managing the J-curve and sizing the reserve so an illiquid Gulf programme is never caught short

Chennakeshav Adya

Independent Researcher

June 2026

ABSTRACT

An illiquid private-markets programme draws capital before it returns it, on a schedule that is uneven and only partly predictable, and the most common way such a programme fails in practice is not poor investment selection but poor liquidity management. This paper sets out how a family-office or institutional investor should manage the liquidity of an illiquid Gulf allocation: how to anticipate the cash-flow profile, how to size the liquidity reserve, what sources of liquidity to rely on, and how to avoid the twin errors of holding too much idle cash and being caught short. Drawing on the literature on private-markets cash-flow modelling, the management of illiquidity in institutional portfolios, and the denominator effect, it advances five propositions concerning the liquidity management of an illiquid programme. Using a stylised, clearly-labelled framework, it models the cash-flow profile of a paced programme, shows how the reserve should be sized to the peak rather than the average net call, sets out the hierarchy of liquidity sources, and quantifies the cost of getting the reserve wrong in either direction. The analysis finds that liquidity management is the largest controllable operational risk in an illiquid programme; that the reserve must be sized to a stressed peak rather than an average; that a steadily paced programme becomes self-funding but only after a multi-year build; and that the denominator effect, the rise in an illiquid weight when liquid assets fall, must be planned for at the outset. The paper provides a liquidity-management framework, a reserve-sizing approach and a checklist, and discusses the limitations and avenues for further research.

JEL Classification: G11, G23, G32, C61, O53

Keywords: liquidity management, private markets, capital calls, J-curve, cash-flow modelling, denominator effect, Gulf allocation, illiquidity

1. INTRODUCTION

An illiquid private-markets programme has an awkward feature that distinguishes it from a portfolio of liquid assets: it draws capital before it returns it, on a schedule the investor does not fully control. Capital is called as managers find investments, often unevenly and at short notice; distributions arrive later and on their own timetable as those investments mature and are realised. Between the two lies a period in which the programme is a net consumer of cash, and managing the investor's liquidity through that period, and through the unpredictable cash flows of the whole programme, is a distinct discipline. It is also, on the evidence and in practice, the discipline whose failure most often damages an otherwise sound programme, because an investor caught short of liquidity is forced to sell other assets at the wrong time, or to default on a capital call, either of which can do more harm than a poor investment ever would.

This paper sets out how a family-office or institutional investor should manage the liquidity of an illiquid Gulf allocation. It is written for the chief investment officer, the head of alternatives and the operational team responsible for funding the programme and managing the investor's cash, whether for a first Gulf allocation or an established one. Its purpose is not to discourage illiquidity, which is a source of return for a long-horizon investor, but to set out how to bear it safely: how to anticipate the cash flows, size the reserve, marshal the sources of liquidity and avoid the errors that turn manageable illiquidity into a crisis.

The Gulf context sharpens the question in two ways. First, a Gulf allocation is expressed substantially through illiquid private assets, real estate, private credit, digital infrastructure, private equity, so the liquidity-management discipline applies with full force. Second, the region's secondary market for private positions is still developing, so the investor cannot assume it can sell its way out of a liquidity squeeze, which makes the planned reserve and the disciplined cash-flow management all the more important. The companion analyses establish that the Gulf merits an allocation and how to construct it; this paper addresses the operational discipline without which the allocation cannot be safely held.

The paper makes three contributions. First, it models the cash-flow profile of a paced Gulf programme and shows how to anticipate the net-call period and the eventual self-funding. Second, it sets out how the liquidity reserve should be sized, to a stressed peak rather than an average, and quantifies the cost of getting it wrong in either direction. Third, it establishes the hierarchy of liquidity sources and addresses the denominator effect, the rise in an illiquid allocation's weight when liquid assets fall. Throughout, the figures are modelled and clearly labelled; they illustrate the liquidity mechanics rather than forecasting outcomes, and they are not forecasts.

It is worth framing at the outset why this discipline is treated here as a subject in its own right rather than as a footnote to investment selection. In a portfolio of liquid assets, liquidity is largely a given: the investor can raise cash by selling at a market price more or less on demand, so the question of funding is rarely existential. In an illiquid programme it is the reverse. The investor has contractually committed to provide capital when called, cannot easily reverse that commitment, and cannot reliably raise cash from the positions themselves. Liquidity therefore moves from being an assumption to being a constraint that must be actively managed, and the investor that fails to manage it does not merely earn a lower return, it can be forced into actions, distress sales, defaults, that destroy value created elsewhere. The purpose of this paper is to give that constraint the structured treatment it deserves.

In situating these contributions, the paper aims to fill a gap between two bodies of writing. The academic literature on private-markets cash flows and on illiquidity is rigorous but pitched at a level of generality that a practitioner must translate before applying; the practitioner literature on funding private programmes is often firm-specific and not always grounded in the underlying mechanics. This paper tries to bridge the two, taking the structure the academic work establishes and rendering it as a set of disciplines an allocator can apply, with explicit attention to the Gulf context in which the companion papers situate the wider allocation. The intended reader is therefore the practitioner who wants the reasoning, not just the rules, so that the disciplines can be adapted to a particular programme rather than applied mechanically.

The remainder of the paper is structured as follows. Section 2 reviews the literature on private-markets cash-flow modelling, the management of illiquidity, and the denominator effect, and

develops five propositions. Section 3 describes the data and the stylised methodology. Section 4 presents and discusses the results, covering the cash-flow profile, reserve sizing, the sources of liquidity, the self-funding path, the denominator effect and the cost of mis-provisioning. Section 5 reports robustness. Section 6 sets out implementation considerations. Section 7 offers concluding comments, including limitations and directions for further research. A statement of disclosures, the references and appendices follow.

2. LITERATURE REVIEW AND PROPOSITIONS

The analysis draws on three strands of the literature: the modelling of private-markets cash flows; the management of illiquidity within institutional portfolios; and the denominator effect and its consequences in stress. This section reviews each and develops the propositions the analysis examines.

2.1 Modelling Private-Markets Cash Flows

A established literature models the cash flows of private-markets programmes, the timing and size of capital calls and distributions, and the resulting net cash-flow profile (Takahashi and Alexander, 2002). The central finding is that these cash flows, while uncertain in their precise timing, have a characteristic and modellable shape: calls front-loaded as capital is deployed, distributions back-loaded as investments mature, and a net-negative period, the J-curve, before the programme turns cash-positive. The literature establishes that an investor can and should model this profile forward, under conservative assumptions about call timing, to anticipate its liquidity needs. For a Gulf programme this means the net-call period and the eventual self-funding are predictable in shape if not in exact timing, which is the basis of the first proposition: that liquidity needs can be anticipated and planned for rather than merely reacted to.

The cash-flow modelling literature also clarifies why the timing of calls is so hard to predict at the level of an individual programme even when its aggregate shape is reliable. Calls are driven by the manager's deal flow, which depends on market conditions, competition for assets and the manager's own discipline, none of which the investor controls or can forecast precisely. Distributions are driven by exits, which depend on the maturity of investments and on the availability of buyers and financing, again outside the investor's control. The practical consequence, drawn out by the modelling literature, is that the investor should model the aggregate shape with reasonable confidence but treat the timing of any particular call or distribution as genuinely uncertain, and fund accordingly, with a reserve sized for the uncertainty rather than for the central path. This is the bridge from the modelling literature to the reserve-sizing question that the illiquidity-management literature addresses.

The modelling literature further establishes that the most useful output of a cash-flow projection is not a single expected path but a distribution of possible paths, from which the investor can read the plausible worst case against which to provision. A projection that yields only a central estimate invites the investor to fund to that estimate and to be caught when the actual path runs worse, whereas a projection that yields a range, even a roughly estimated one, directs the investor to fund toward the adverse end. For an illiquid Gulf programme, where the consequences of being caught short are severe and the escape routes limited, this distributional reading is the prudent one. The investor need not compute a precise probability distribution; it is enough to ask, repeatedly and conservatively, how much worse than the central case the calls and distributions could plausibly run, and to provision for that. The modelling literature thus supplies not just a

tool but a disposition: to fund the programme against an adverse plausible path rather than a comfortable central one.

2.2 Managing Illiquidity in Institutional Portfolios

A second literature addresses how institutions should bear illiquidity, balancing the illiquidity premium against the need to meet obligations and to rebalance (Ang, 2014; Swensen, 2009). The consistent guidance is that a long-horizon investor can and should bear illiquidity to harvest its premium, but must size it against its genuine liquidity needs and maintain a reserve to meet calls and obligations without forced selling. The literature also warns of the asymmetry of the error: under-provisioning, holding too little liquidity, risks forced sales and call defaults at the worst time, while over-provisioning, holding too much idle cash, drags returns; both are costly, but the forced-sale error is the more dangerous (Ang, 2014). This underpins the second and fifth propositions, on sizing the reserve to a stressed peak and on the asymmetric cost of getting it wrong.

The illiquidity-management literature also distinguishes between two kinds of liquidity an investor needs: liquidity to meet the programme's own calls, the focus of this paper, and liquidity to meet the investor's other obligations and to rebalance the wider portfolio. The two interact, because the same reserve and the same liquid assets may be called upon for both, and a stress that raises one demand often raises the other. A complete liquidity policy considers the programme's calls alongside the investor's wider needs, sizing the total liquidity provision against the combined demand rather than against either in isolation. This paper concentrates on the programme's own liquidity because that is the demand the illiquid allocation creates and the one most often under-appreciated, but the reader should hold in mind that the reserve sized here sits within a wider liquidity policy and should be reconciled with it.

2.3 The Denominator Effect

A third strand concerns the denominator effect, the phenomenon by which a fall in the value of an investor's liquid assets raises the weight of its illiquid allocation above target, precisely when the investor is least able to rebalance, because the illiquid assets cannot be quickly sold and may even require further capital calls. The experience of institutional investors in past crises documents how the denominator effect can push illiquid allocations well above target, straining liquidity and forcing difficult choices. The lesson is that an investor must anticipate the denominator effect at the construction and liquidity-planning stage, sizing the illiquid allocation with headroom and holding a reserve, rather than discovering the effect in a crisis. This is the fourth proposition.

2.3a Why Liquidity Failure, Not Selection, Ends Programmes

A point that runs through all three literatures, and through practitioner experience, deserves to be stated plainly: the most common way an illiquid programme damages an investor is not through poor investment selection but through a liquidity failure. A programme of perfectly chosen investments can still inflict serious harm if the investor cannot fund its calls, because the consequences of a liquidity failure, forced sales at depressed prices, default penalties, forfeiture of investments, and damaged manager relationships, are immediate and often irreversible, whereas a poor investment merely underperforms. This asymmetry is why the liquidity-management discipline deserves at least as much attention as the selection discipline, and arguably more, because it is both more controllable and more consequential when it fails. The

literatures on cash-flow modelling, illiquidity management and the denominator effect are, in this light, a single body of guidance on how to avoid the failure that does the most damage.

The Gulf context heightens this further. Because a Gulf allocation is concentrated in illiquid assets and the regional secondary market is still developing, the investor has fewer escape routes from a liquidity squeeze than it would in a deeper, more liquid market. The disciplines this paper sets out are therefore not optional refinements but the conditions under which a Gulf allocation can be safely held at all, which is why the analysis treats liquidity management as the operational foundation of the whole programme rather than as an administrative afterthought.

2.4 Synthesis and Propositions

Taken together, these literatures imply that the liquidity of an illiquid programme can be anticipated and planned, that the reserve must be sized to a stressed peak and the denominator effect anticipated, and that the asymmetric cost of mis-provisioning argues for prudent but not excessive liquidity. From this the paper develops five propositions:

Proposition 1. The cash-flow profile of a paced Gulf programme, its net-call period and eventual self-funding, is predictable in shape and can be modelled and planned for in advance.

Proposition 2. The liquidity reserve must be sized to a stressed peak net call rather than the average, because it is the peak, not the average, that causes a liquidity failure.

Proposition 3. An investor should rely on a hierarchy of liquidity sources, the reserve, distributions, a credit line, a liquid sleeve and, last, the secondary market, rather than on any single source, especially given the Gulf's developing secondary market.

Proposition 4. The denominator effect must be anticipated at the outset, by sizing the illiquid allocation with headroom and holding a reserve, not discovered in a crisis.

Proposition 5. The cost of mis-provisioning is asymmetric, under-provisioning risks forced sales and is more dangerous than the return drag of over-provisioning, so the reserve should err toward prudence.

3. DATA AND METHODOLOGY

This study is analytical, using a transparent, stylised cash-flow framework whose assumptions are stated explicitly rather than estimated from a proprietary dataset, which suits a question about the prospective liquidity of a programme. The figures are modelled and clearly labelled; they illustrate the liquidity mechanics an investor can apply to its own programme, and they are not forecasts.

3.1 The Liquidity Framework

The framework has three components: a cash-flow model projecting the calls, distributions and net cash flow of a paced Gulf programme over its life; a reserve-sizing analysis showing how the reserve should be set against a stressed peak net call; and an analysis of the hierarchy of liquidity sources and the denominator effect. Together these address the anticipation, provisioning and management of the programme's liquidity.

3.1a What the Framework Does and Does Not Do

It is important to be clear about the framework’s scope. It is a tool for reasoning about the shape and magnitude of an illiquid programme’s liquidity demands and about how to provision for them; it is not a forecasting model, and it does not attempt to predict the precise calls and distributions of any particular programme. The figures should be read as illustrations of mechanics, the way calls lead distributions, the way the peak exceeds the average, the way a market fall raises the illiquid weight, rather than as numerical predictions. The value of such a framework is that it makes the structure of the problem visible, so that an investor can apply the same reasoning to its own programme with its own, better-informed inputs. A reader who substitutes the framework’s stylised numbers for its own circumstances would be misusing it; a reader who adopts its disciplines and populates them with its own conservative estimates would be using it as intended.

The framework is deliberately transparent for a second reason. Liquidity management is an area where false precision is dangerous, because it can lull an investor into provisioning to a central estimate that a stress then overwhelms. By keeping the assumptions explicit and the mechanics simple, the framework keeps the investor’s attention on the questions that matter, how bad could the stressed peak be, and is the reserve sized for it, rather than on the spurious accuracy of a complex model. In liquidity management, a simple model used conservatively beats a complex one used trustingly.

3.2 Assumptions and Their Justification

The stylised inputs, the pace of commitments, the timing of calls and distributions, and the stressed-scenario assumptions, are calibrations informed by the private-markets cash-flow literature (Takahashi and Alexander, 2002) and by the typical behaviour of private programmes; they are indicative rather than precise estimates and are recorded in Appendix A. The call-timing and stress assumptions are treated as the central uncertain inputs and are tested in Section 5. The base-case assumptions are set out in Table 1.

Table 1. Base-Case Assumptions for the Liquidity Framework

Parameter	Assumption
Programme type	Paced, illiquid Gulf private-markets programme
Build / commitment period	5 years
Net-negative (J-curve) period	Years 1 to 4
Self-funding onset	Around years 5 to 6
Average annual net call	~25% of annual commitment
Stressed peak net call (reserve basis)	~60% of annual commitment
Secondary market	Developing; not a primary liquidity source

Indicative, forward-looking assumptions used to illustrate liquidity mechanics; not forecasts. See Appendix A.

4. RESULTS AND DISCUSSION


This section presents the framework in the order of the propositions: the cash-flow profile (Proposition 1); reserve sizing (Proposition 2); the hierarchy of liquidity sources (Proposition 3);

the self-funding path and the denominator effect (Proposition 4); and the cost of mis-provisioning (Proposition 5).

4.1 The Cash-Flow Profile

The starting point is to anticipate the programme’s cash-flow profile, shown in Figure 1.

Figure 1. The Liquidity Challenge: Calls, Distributions and Net Flow



Indicative; calls are front-loaded, distributions back-loaded, and the net flow is negative before it turns positive.

The profile supports Proposition 1. The capital calls are front-loaded, peaking in the early years as the programme deploys; the distributions are back-loaded, building as investments mature; and the net cash flow is negative in the early years before turning positive, the J-curve. The crucial point is that this shape is predictable even though the exact timing is not, which means an investor can model it forward and plan its liquidity around the anticipated net-call period rather than reacting to calls as they arrive. An investor that models its programme’s cash flows, under conservative call-timing assumptions, knows in advance roughly how much liquidity it will need and when, which transforms liquidity management from a reactive scramble into a planned provision. The first discipline of liquidity management is therefore to build and maintain this forward cash-flow model, updating it as commitments are made and as actual calls and distributions arrive.

The distinction between calls and distributions also shapes how the investor should think about the two flows when planning. Calls are an obligation: once committed, the investor must meet them or face the consequences of default, so they must be funded with near-certainty. Distributions are an expectation: they are likely but not guaranteed in timing or amount, so they should be relied upon cautiously in any funding plan. This difference in character, certain obligation against uncertain expectation, argues for an important asymmetry in how the investor treats the two: fund the calls as though they will arrive on the early, heavy end of their plausible range, and count the distributions as though they may arrive on the late, light end. Planning to this conservative pairing protects the investor against the combination that causes trouble, and it costs little when the flows turn out closer to the central case, because the surplus reserve simply earns

its modest return until it is needed. This asymmetric treatment of obligation and expectation is one of the simplest and most valuable habits in liquidity management.

4.1a Reading the Profile in Practice

The practical use of the cash-flow profile is to convert it into a funding plan. An investor that has modelled the profile knows, year by year, roughly how much net cash the programme will demand, and can arrange to have that cash available before it is called rather than scrambling when the call arrives. The model also tells the investor when the demand peaks and when it eases, which allows the reserve to be sized to the demanding years and, cautiously, run somewhat leaner once self-funding is established. The discipline is to let the model drive the funding plan, updating both as the programme runs, so that the investor is always looking forward to the next demand rather than reacting to the last one. A programme funded from a forward model is calm; a programme funded from the latest call notice is perpetually anxious and prone to error.

It is worth emphasising that the profile's precise timing is genuinely uncertain, and the model should be built and used in that spirit. The investor should not treat the central projection as a forecast to be relied upon, but as the centre of a range, and should fund to the conservative end of that range. This is not pessimism; it is the recognition that the cost of being ready for a call that arrives early is small, while the cost of being unready for it is large, an application of the asymmetry that runs through the whole analysis.

4.1b The Gulf Dimension of the Liquidity Profile


The cash-flow profile of a Gulf programme carries some features the investor should weigh. The region's private markets span real estate, private credit, digital infrastructure and private equity, and these have differing call and distribution rhythms: development-linked real estate and data-centre projects can call capital in large, construction-driven tranches, while private credit may call more steadily and distribute current income that helps fund other calls. A Gulf programme that blends these will have a composite profile whose shape depends on the mix, and the investor should model that composite rather than assuming a generic pattern. The income-distributing element, in particular, can be a useful internal source of liquidity that brings forward the self-funding point, which is one reason a balanced Gulf programme can be easier to fund than one concentrated purely in capital-appreciation strategies that distribute only on exit. The discipline is to understand the cash-flow character of each sleeve and to build the composite profile from them, so that the reserve is sized to the programme the investor actually holds.

The second Gulf feature is the developing state of the regional secondary market, which the analysis has flagged as the reason the secondary sale sits last in the source hierarchy. As the market deepens over time, the last resort will become somewhat more reliable and the discount in a stress somewhat narrower, which will modestly ease the liquidity constraint. But an investor sizing a reserve today should not anticipate that improvement; it should provision as though the secondary market may offer little in a stress, treating any future deepening as an upside rather than a planning assumption. This conservatism is the prudent posture given the asymmetry of the consequences, and it is consistent with the analysis's general bias toward provisioning against the adverse case.

4.2 Sizing the Liquidity Reserve

Having anticipated the cash-flow profile, the investor must size the reserve it holds against it. Figure 2 illustrates the principle.

Figure 2. Sizing the Reserve: Average Net Call versus Stressed Peak



Indicative; sizing to the average leaves the programme exposed in a stressed year.


The figure supports Proposition 2. The average annual net call is modest, but the peak net call in a stressed year, when calls bunch and distributions slow, is materially larger. An investor that sizes its reserve to the average is adequately funded in a normal year but exposed in precisely the stressed year when liquidity matters most, because a liquidity failure is caused by the peak, not the average. The reserve must therefore be sized to a conservative, stressed estimate of the peak net call over a meaningful horizon, not to the average. This is the single most important quantitative decision in liquidity management, and the most common error is to anchor on the average because it is the more comfortable number. The reserve is not idle waste; it is the price of being able to meet every call without forced selling, and it should be sized accordingly.

A practical approach is to model the net call under a stress scenario, calls accelerated and distributions slowed simultaneously, and to size the reserve to cover the resulting peak with a margin. The reserve need not be held entirely in cash; it can be held in high-quality liquid assets that can be sold quickly without material loss, which recovers some of the return drag while preserving the liquidity. The discipline is to ensure the reserve is genuinely available, not notionally available assets that would themselves be impaired in the stress that triggers the call.

A further practical question in sizing the reserve is the horizon over which the stressed peak should be measured. A reserve sized to the worst single quarter is different from one sized to the worst rolling year, and the investor must choose a horizon that matches how quickly its back-stop sources can be mobilised. If the credit line and liquid sleeve can be accessed within days, the reserve need only cover the shortest period before those sources arrive; if they take weeks or are uncertain, the reserve must cover a longer stretch unaided. The discipline is to align the reserve horizon with the realistic speed of the back-stops, and to be conservative about that speed, because the back-stops are slowest and least certain in precisely the stress that triggers the call. An investor that assumes instant access to a credit line that in fact takes a fortnight to draw has mis-sized its reserve regardless of how carefully it estimated the peak.

4.3 The Hierarchy of Liquidity Sources

An investor should not rely on a single source of liquidity. Figure 3 sets out the hierarchy.



Indicative ordering by reliability and cost; the secondary market is the last resort, not the first.

The hierarchy supports Proposition 3. The first and most reliable source is the dedicated reserve, available on demand at no cost beyond its return drag. The second is the programme’s own distributions, which fund later calls once the programme matures, the self-funding mechanism discussed below. The third is a committed credit line or subscription facility, which can bridge a timing mismatch at a known cost. The fourth is a liquid sleeve elsewhere in the portfolio that can be sold to meet a call, at the cost of disturbing the wider allocation. The last resort is the secondary market, selling the private positions themselves, which in the Gulf is still a developing market and may offer poor prices in a stress. The discipline is to rely on the hierarchy in order, reserving the lower, costlier and less reliable sources for genuine stress, and never to plan on the secondary market as a primary source, because it is least available exactly when it is most needed.

Table 2 summarises the sources, their reliability and their cost.

Table 2. Liquidity Sources Ranked by Reliability and Cost

Source	Reliability	Cost	Role
Dedicated reserve	Highest	Return drag on idle assets	First line; meets routine calls
Programme distributions	Medium (matures over time)	None	Self-funds later calls
Committed credit line	High (if committed)	Interest / commitment fee	Bridges timing mismatches
Liquid sleeve sale	Medium	Disturbs wider allocation	Reserve back-stop
Secondary sale	Lowest (developing market)	Potentially deep discount	Last resort only

Indicative. Rely on sources in order; do not plan on the secondary market as primary.

The liquid sleeve elsewhere in the portfolio deserves a parallel caution. It is a genuine back-stop, but using it to meet a call means selling other assets, which disturbs the wider allocation and may crystallise losses if the sale comes in a down market. The investor should therefore treat the liquid sleeve as a back-stop behind the reserve and the credit line, to be used when those are exhausted, and should be clear in advance about which assets would be sold and in what order, so that the decision is made on a considered policy rather than under pressure. A pre-agreed sell-down order, favouring the most liquid and least strategically important holdings first, turns a potentially disorderly scramble into a controlled, reversible adjustment, and it is far better settled in calm conditions than invented in a stress.


4.3a The Credit Line as a Bridge, Not a Crutch

The committed credit line deserves particular comment because it is both valuable and easily misused. Used as a bridge, to smooth a short timing mismatch between a call and an expected distribution or reserve top-up, a credit line is an efficient tool that reduces the idle cash the investor must hold and so recovers some return. Used as a crutch, to fund calls that the investor has no near-term means of repaying, it converts a liquidity problem into a leverage problem and can leave the investor worse off when the facility must be repaid or is withdrawn. The discipline is to use the credit line for genuine timing bridges within a clear repayment path, and never to let it substitute for an adequate reserve. A credit line is also a contract that can be repriced or withdrawn, particularly in the stress when it is most wanted, so it should be a layer behind the reserve, not the reserve itself.

4.4 The Self-Funding Path

A steadily paced programme eventually funds its own calls from its distributions. Figure 4 shows the cumulative net cash position over the programme's life.

Figure 4. The Path to Self-Funding: Cumulative Net Cash Position



Indicative; the cumulative position troughs before recovering as distributions overtake calls.

The figure illustrates the self-funding dynamic central to Proposition 1 and relevant to Proposition 4. The cumulative net cash position declines through the build years as calls exceed


distributions, troughs around the point of peak deployment, and then recovers as the maturing programme's distributions overtake its calls and eventually fund them entirely. The crucial practical points are two. First, the trough is the maximum cumulative liquidity the programme demands, and the reserve plus other sources must be able to cover it. Second, self-funding is not immediate; it arrives only after a multi-year build during which the investor must fund the programme from its reserve and other sources. An investor that expects self-funding too early, or that does not provision for the trough, will be caught short. The discipline is to model the cumulative position, identify the trough, provision for it, and recognise that the relief of self-funding is a reward for patience, not an early feature.

It is worth dwelling on why the multi-year wait for self-funding is so often underestimated in practice. An investor that sees other, mature programmes funding their own calls may assume its new programme will do likewise, forgetting that those programmes reached self-funding only after their own build years. A new programme has no distributions to draw on at the start and cannot self-fund until its early investments mature, which takes years. The error of expecting self-funding too early is therefore an error of forgetting the programme's own youth, and it is corrected by modelling the cumulative position explicitly and reading off when, on conservative assumptions, distributions will actually overtake calls. Until that point the programme must be funded from the reserve and back-stops, and the investor that has modelled the trough knows how much that requires and for how long.

4.5 The Denominator Effect

A specific risk an illiquid programme must plan for is the denominator effect. Figure 5 illustrates it.

Figure 5. The Denominator Effect: Illiquid Weight Rises When Liquid Assets Fall



Indicative; a fall in liquid assets raises the illiquid weight above target without any change in the illiquid assets themselves.


The figure supports Proposition 4. When the liquid part of a portfolio falls in value, in a market decline, the weight of the illiquid allocation rises mechanically, even though the illiquid assets

themselves have not changed, because they are a larger share of a smaller total. The investor may find its illiquid allocation well above target precisely when it cannot rebalance, because the illiquid assets cannot be quickly sold and may even be calling further capital. The denominator effect thus compounds a market stress with a liquidity and allocation stress. The lesson is to anticipate it at the outset: to size the illiquid allocation with headroom below any hard ceiling, so that a denominator-effect rise does not breach the limit; to hold a reserve that does not itself collapse in the stress; and to plan rebalancing around the illiquidity rather than assuming it can be done on demand. An investor that plans for the denominator effect treats it as a known feature; one that does not discovers it as a crisis.

The denominator effect also has a less obvious second-round consequence worth noting. When a market fall raises the illiquid weight above target, an investor governed by a hard ceiling may feel obliged to reduce the illiquid allocation, but the only way to do so quickly is to sell into the developing secondary market at a discount, crystallising a loss and surrendering the illiquidity premium it was being paid to bear. The better course, available only to an investor that anticipated the effect, is to have set the target with enough headroom that a denominator-effect rise does not breach the ceiling, so that no forced reduction is required and the investor can simply wait for the liquid assets to recover. This is why the denominator effect must be planned for at construction: the investor that built in headroom rides it out, while the investor that did not is forced to sell exactly when selling is most costly. Planning converts the denominator effect from a forced-sale trigger into a tolerable fluctuation.

4.6 The Cost of Mis-Provisioning


Finally, the analysis quantifies the asymmetric cost of getting the reserve wrong. Figure 6 contrasts over- and under-provisioning.



Indicative; under-provisioning risks forced sales and default, a larger and more dangerous cost than the drag of over-provisioning.

The figure supports Proposition 5. Over-provisioning, holding more reserve than needed, costs the return drag on the excess idle assets, a real but bounded and recoverable cost. Under-provisioning, holding too little, risks forced sales of other assets at depressed prices, or, worse, default on a capital call, which can trigger penalties, forfeiture of the investment, and reputational damage with managers. The cost of under-provisioning is therefore both larger and more dangerous than the cost of over-provisioning, and it strikes in stress when it is least affordable. This asymmetry argues for sizing the reserve toward prudence: an investor should prefer to carry a little too much liquidity than a little too little, because the two errors are not symmetric. The art is to provision prudently without hoarding, but when in doubt, the prudent error is the safer one.

Figure 7 places the sensitivity of the programme’s peak liquidity need to the key assumptions in a single view.




Indicative; call-timing and distribution-pace assumptions dominate the peak liquidity need.

The sensitivity analysis confirms that the call-timing and distribution-pace assumptions dominate the programme’s peak liquidity need, which is why the reserve must be sized to a stressed combination of both, and why these are the assumptions the investor should monitor most closely as the programme runs.

Figure 8 contrasts three approaches to liquidity management on a common set of measures.

Figure 8. Three Approaches to Liquidity Management Compared



Indicative; the planned-reserve approach balances safety against drag better than either under- or over-provisioning.

The comparison draws the analysis together. The under-provisioned approach carries the least drag but the highest risk of forced sales and default; the over-provisioned approach is safe but carries a heavy return drag; the planned-reserve approach, sized to a stressed peak with a hierarchy of sources behind it, achieves safety at a far lower cost than blanket over-provisioning. The planned-reserve approach is the one the analysis recommends, because it manages the liquidity risk deliberately rather than by either hoping for the best or hoarding against it.

A point implicit in the asymmetric cost deserves to be made explicit, because it changes how an investor should weigh the reserve decision. The return drag of over-provisioning is paid in good times and bad alike, a small, steady cost that is easy to see and therefore easy to resent; the cost of under-provisioning is paid only in the stress, but when paid it is large, sudden and often irreversible. Human nature tends to over-weight the visible, recurring cost and under-weight the rare, severe one, which biases investors toward holding too little reserve. Recognising this bias is itself a tool: an investor that knows it will be tempted to under-provision can set a policy that sizes the reserve to the stress and resists the temptation to trim it when the drag is salient and the stress feels remote. The framework, in this sense, is partly a device for committing the investor in calm times to the prudence it will be glad of in difficult ones, much as any sound risk policy binds the future self against a predictable lapse.

4.6a A Worked Illustration of Reserve Sizing

To make the reserve-sizing logic concrete, consider an investor committing to a paced Gulf programme and reasoning through the provision in the spirit of the framework. The investor first models the programme's aggregate cash-flow profile and reads off two numbers: the average annual net call, which sets the routine funding need, and the stressed peak net call, which sets the reserve. Suppose the model shows an average net call of roughly a quarter of the annual commitment but a stressed peak, calls accelerated and distributions slowed together, of around three-fifths. An investor that funds to the average holds enough for a normal year and is short by

more than a third of a commitment in the stressed year. An investor that funds to the stressed peak, with a margin, holds enough for the bad year too, at the cost of carrying somewhat more reserve in the normal years. The framework makes the choice explicit: the additional reserve is the premium the investor pays to be certain of meeting every call, and the asymmetry of the consequences, a bounded drag against an unbounded forced-sale-and-default risk, makes that premium worth paying.

The same illustration shows how the back-stops change the sum. If the investor has a committed credit line it can draw quickly and reliably, part of the gap between the average and the stressed peak can be covered by the line rather than by idle cash, reducing the drag while preserving the safety, provided the line is genuinely committed and the investor has a clear path to repay it from distributions or reserve top-ups. The reserve then covers the portion of the stressed peak that must be met before the line and other sources can be mobilised, and the line covers the remainder. This layering, reserve first, committed line behind it, liquid sleeve behind that, is how a disciplined investor achieves safety at a lower cost than holding the entire stressed peak in cash, and it is the practical expression of the source hierarchy.

4.7 Bringing the Six Mechanics Together

The six elements of the analysis, the cash-flow profile, reserve sizing, the source hierarchy, the self-funding path, the denominator effect and the cost of mis-provisioning, are not separate techniques but a single connected discipline. The cash-flow profile tells the investor what the programme will demand; reserve sizing converts that demand into a provision against the stressed peak; the source hierarchy arranges the back-stops behind the reserve; the self-funding path shows when the demand eases; the denominator effect warns that a market stress will compound the liquidity stress; and the asymmetry of mis-provisioning sets the bias toward prudence that ties the whole together. An investor that applies them as a connected system, rather than picking one or two, manages the programme's liquidity deliberately. The recurring failure mode, in the Gulf and elsewhere, is to model the profile but under-size the reserve, or to size the reserve but neglect the denominator effect, leaving a gap that a stress finds. The value of treating the six as one discipline is that it closes those gaps.

It is also worth restating where judgement enters. The framework quantifies the mechanics, but the investor must still choose the stress severity to provision against, the balance between reserve and credit line, and the headroom to leave for the denominator effect. These are judgements about risk appetite and circumstances, not outputs of a formula, and they should be made explicitly and revisited as the programme and the investor's wider position evolve. The framework disciplines the judgement; it does not replace it.

5. ROBUSTNESS AND SCENARIO ANALYSIS

Because the framework is stylised, its value lies in whether its conclusions hold under different assumptions. This section varies the central inputs.

5.1 Faster and Slower Call Pacing

The pace at which capital is called is the principal driver of the early liquidity need. If calls arrive faster than the base case, the net-call period is shorter but deeper, raising the peak and the required reserve; if slower, the period is longer but shallower. Under either, the qualitative conclusions hold: the reserve must be sized to the peak, and the peak is sensitive to call timing,

which is precisely why a stressed call-timing assumption should drive the reserve. Faster pacing strengthens the case for a larger reserve; slower pacing does not remove the need for one.

5.2 Slower Distributions

If distributions arrive more slowly than the base case, the self-funding point is delayed and the cumulative trough is deeper, raising the maximum liquidity the programme demands. This is the more dangerous stress, because it extends the period over which the investor must fund the programme from its reserve and other sources. The conclusion is reinforced: the reserve and credit line must be sized for a scenario in which distributions disappoint, not only for the base case, and the self-funding relief must not be assumed to arrive on schedule.

Pacing interacts with the reserve in a way worth drawing out. Because a paced programme smooths and partly offsets its own flows, it lowers both the average and the stressed peak net call, which reduces the reserve the investor must hold for a given level of commitment. An unpaced programme, by contrast, concentrates its calls and so demands a larger reserve and leans more heavily on the back-stops. The investor therefore has two levers on the same risk: it can reduce the liquidity demand at source through disciplined pacing, and it can provision against the demand that remains through the reserve and hierarchy. The most efficient approach uses both, pacing to shrink the problem and provisioning to cover what is left, rather than relying on provisioning alone to absorb a lumpy, self-inflicted profile. This linkage is why the construction and liquidity disciplines should be designed together rather than in sequence, and why an investor that paces well has already done much of its liquidity management before the reserve is sized.

5.2a The Role of Pacing in Reducing Liquidity Risk

A robustness consideration that sits across the scenarios is the role of pacing, the deliberate spreading of commitments across years and vintages. A programme that commits steadily over several years, rather than all at once, smooths its cash-flow profile: at any time it holds a mix of young commitments calling capital and older ones distributing it, which partly offsets the two flows and reduces the net demand and its volatility. Pacing therefore reduces the liquidity risk the reserve must cover, and a well-paced programme can be safely run with a smaller reserve relative to commitments than a lumpy one. The robustness implication is that the investor's pacing policy and its liquidity policy are linked: disciplined pacing is itself a liquidity-management tool, and an investor that paces well has already reduced the problem this paper addresses before the reserve is even sized. The scenarios above assume a paced programme; an unpaced, lumpy programme would show deeper troughs and a larger required reserve, which is a further argument for pacing.

It is also instructive to consider the scenario the framework is not designed for: a systemic event in which the investor faces simultaneous demands on every source, calls accelerating, distributions stopping, the credit line withdrawn, the liquid sleeve impaired and the secondary market closed. No reserve sized to a normal stress will fully cover such a tail event, and it would be dishonest to claim otherwise. The framework's response is twofold: size the reserve to a severe but plausible stress rather than to a normal one, which covers the great majority of cases; and recognise that for the true tail, the protection lies in the conservatism of the whole posture, the headroom for the denominator effect, the refusal to lean on the credit line, the patience not to assume early self-funding, rather than in any single number. An investor that has been prudent across all the disciplines will enter even a systemic event better placed than one that provisioned to a comfortable case, even if no one can be fully insulated from the most extreme outcomes.

Acknowledging this limit honestly is better than a false promise of complete safety, and it points the investor toward the conservatism that offers the most protection available.

5.3 A Combined Stress

The most demanding scenario combines accelerated calls with slowed distributions, the pattern that tends to occur in a market stress, when managers find opportunities and call capital while exits slow and distributions dry up. This combined stress is exactly what the reserve should be sized against, and it is also when the denominator effect bites. Under this scenario the planned-reserve approach with a credit-line back-stop holds, while the under-provisioned approach fails. The robustness analysis thus confirms the central recommendation: size the reserve to a combined stress, hold a hierarchy of sources behind it, and never plan on the secondary market or on self-funding arriving early.

A final scenario worth considering is a benign one, in which calls arrive slowly and distributions arrive early and strong. It might seem that such a case poses no liquidity problem, and in the narrow sense it does not. But it carries a subtler risk: an investor that enjoys an easy early experience may conclude that its reserve is excessive and run it down, only to be caught when a later, more demanding programme or a market turn restores the stress. The discipline in good times is therefore not to mistake a favourable draw for a permanent reduction in the liquidity need, and to maintain the reserve and back-stops sized to the stress that has merely not yet arrived. This is the same conservatism that governs the adverse scenarios, applied in the easy ones, and it guards against the complacency that an unbroken run of comfortable years can breed.

5.4 When the Framework Is Most and Least Reliable

A candid robustness assessment should say where the framework is most and least reliable. It is most reliable as a guide to the structure of the liquidity problem, the lead of calls over distributions, the gap between peak and average, the mechanics of the denominator effect, and the asymmetry of mis-provisioning, because these follow from the nature of an illiquid programme rather than from any particular numbers. It is least reliable as a source of specific magnitudes, because the actual calls, distributions and stress severity of any programme depend on managers, markets and timing that no framework can predict. The correct use, therefore, is to take the structure as robust and to populate the magnitudes with the investor's own conservative estimates, updated as the programme runs. An investor that does this is using the framework where it is strong; an investor that treats its stylised numbers as predictions is leaning on it where it is weak. This honest division of reliability is itself part of good practice, because it keeps the investor's confidence aligned with what the analysis can actually support.

6. IMPLEMENTATION CONSIDERATIONS

Translating the framework into practice involves building the cash-flow model, sizing and holding the reserve, arranging the back-stop sources, and governing the whole over time.

6.1 Building and Maintaining the Cash-Flow Model

The foundation is a forward cash-flow model of the programme, projecting calls and distributions under conservative assumptions and updated as commitments are made and as actual flows arrive. The model need not be elaborate; a disciplined, regularly updated projection is far better than none. It should drive the reserve sizing and be stress-tested for the combined scenario. The discipline is to maintain it as a living tool, not to build it once and file it.

6.2 Sizing and Holding the Reserve

The reserve should be sized to the stressed peak net call with a margin, held in genuinely liquid assets that will not themselves be impaired in the stress that triggers a call, and ring-fenced from other uses so it is available when needed. The investor should resist the temptation to run the reserve down in good times to chase return, because the call it must meet may arrive in bad times.

6.3 Arranging the Back-Stop Sources

Behind the reserve, the investor should arrange the back-stop sources before they are needed: a committed credit line or subscription facility negotiated in calm conditions, clarity on which liquid assets would be sold if necessary, and a realistic, conservative view of what the secondary market might offer in a stress. Arranging these in advance is far cheaper and more reliable than scrambling for them in a crisis.

Beyond the first year, the discipline becomes one of maintenance and review. As the programme matures and begins to distribute, the investor should reassess the reserve in light of the now-better-understood cash-flow character of its particular managers and strategies, neither clinging to an over-large reserve that the maturing programme no longer needs nor running it down prematurely on the strength of a few good distributions. The cumulative position should be tracked against the modelled trough, and the approach to self-funding confirmed before any easing of the reserve. Through the programme's later years, when it is self-funding and distributing, the liquidity focus shifts toward managing the reinvestment of distributions and the funding of any new commitments, which begins the cycle again. Liquidity management is thus a continuous discipline across the programme's life, not a one-time set-up, and the investor that treats it so will find each successive commitment easier to fund than the last, because it is layered onto a maturing base that is already generating its own liquidity.

6.3a Sequencing the First Year of a Programme

For an investor beginning a Gulf programme, the first year sets the liquidity discipline for its life, and a clear sequence helps. Before committing, the investor should build the forward cash-flow model and size the reserve and back-stop sources to the resulting stressed peak, so that the funding is arranged before the first call rather than after it. On committing, it should ring-fence the reserve and put the credit line and liquid-sleeve arrangements in place while conditions are calm and terms are favourable. Through the first year it should track actual calls and distributions against the model, refining the projection and confirming that the provisions remain adequate as the picture sharpens. Establishing this rhythm early means the programme is funded from a plan from the outset, which is far easier than retrofitting discipline onto a programme that has been funded reactively and has drifted into an under-provisioned position.

A related discipline is to document the liquidity plan and its assumptions at the start, so that the rationale for the reserve size and the back-stop arrangements is recorded and can be reviewed as circumstances change. This matters most when the people managing the programme change, or when a stress arrives and decisions must be made quickly: a documented plan lets the investor act on a considered policy rather than improvising under pressure. The documentation need not be lengthy, but it should capture the cash-flow model, the reserve basis, the source hierarchy and the denominator-effect headroom, the four elements on which the programme's liquidity safety rests.

6.3b Common Pitfalls and How to Avoid Them

It helps to name the pitfalls that recur in practice, because naming them is the first step to avoiding them. The first is anchoring the reserve on the average net call, which feels reasonable but leaves the programme exposed in the stressed year; the remedy is to size to a stressed peak as a matter of policy. The second is treating the credit line or the secondary market as part of the reserve rather than as back-stops behind it; the remedy is to hold the source hierarchy explicitly and to fund routine calls from the reserve alone. The third is expecting self-funding too early and running the reserve down before distributions reliably arrive; the remedy is to model the cumulative trough and provision for it. The fourth is neglecting the denominator effect until a market fall pushes the illiquid weight through its ceiling; the remedy is to set the target with headroom from the outset. The fifth is letting the cash-flow model go stale, so that the funding plan drifts away from the programme's actual flows; the remedy is the regular update. Each pitfall is a failure of discipline rather than of knowledge, and each has a simple, known remedy, which is why a structured, governed approach is so effective at preventing them.

A final pitfall worth flagging is over-correction: an investor that has been frightened by a near-miss may swing to holding so large a reserve that the return drag materially impairs the programme, defeating the purpose of bearing illiquidity in the first place. The remedy is the same asymmetry that argues for prudence, read in the other direction: the reserve should be sized to the stressed peak with a sensible margin and a hierarchy of back-stops behind it, not to an extreme that no plausible scenario requires. Good liquidity management is prudent, not fearful, and the difference shows in a reserve that is adequate for the stress without being so large that it hollows out the programme's return. Holding the right reserve, neither too little nor far too much, is the practical art the framework is designed to support.

6.4 Governance and Monitoring

Liquidity management should be governed explicitly: a clear owner, a regular review of the cash-flow model and reserve adequacy, monitoring of the illiquid weight against its ceiling for the denominator effect, and a documented plan for a liquidity stress. The governance should be proportionate but real, because liquidity management is the operational discipline whose failure does the most damage, and it deserves commensurate attention.

6.5 A Note on Communicating Liquidity Risk Internally

A practical but often neglected part of implementation is communicating the liquidity position to the investor's own decision-makers, whether a family principal, an investment committee or a board. The reserve, by design, looks like idle cash earning little, and there will be recurring pressure to put it to work, especially in good times when the stress feels remote. The liquidity manager must therefore be able to explain, in terms the decision-makers accept, why the reserve exists, what it protects against, and what would happen without it. Framing the reserve as insurance against forced sales and call defaults, rather than as uninvested cash, helps; so does showing the cash-flow model and the stressed peak it is sized against, so the number has a visible rationale. An investor whose decision-makers understand and endorse the liquidity policy will hold to it through the temptations of a long bull market; one whose decision-makers see only the drag will be pressed to erode the reserve until a stress exposes the gap. Communicating the policy well is thus part of maintaining it, and it should be done before the pressure arrives, not in the middle of it.

The same communication discipline applies when a stress does arrive. A documented policy and a clear cash-flow picture let the liquidity manager show decision-makers what is happening,

which sources are being drawn in what order, and why the programme remains funded, which preserves calm and prevents the panicked, value-destroying decisions that stress can prompt. Liquidity management, in the end, is as much about institutional behaviour as about numbers: the best-sized reserve fails if it is run down under pressure, and the best plan fails if it is abandoned in a crisis. Building shared understanding of the policy in advance is what makes it durable, and durability is what turns a good plan into actual protection.

7. CONCLUDING COMMENTS

An illiquid Gulf programme draws capital before it returns it, and managing the liquidity through that period is the discipline whose failure most often damages an otherwise sound programme. This paper has set out how to manage it: anticipate the cash-flow profile, size the reserve to a stressed peak, rely on a hierarchy of sources, plan for the denominator effect, and respect the asymmetric cost of getting the reserve wrong.

The evidence and analysis support five conclusions. First, the cash-flow profile of a paced programme is predictable in shape and can be modelled and planned for. Second, the reserve must be sized to a stressed peak net call, not the average, because the peak is what causes a failure. Third, the investor should rely on a hierarchy of liquidity sources and never plan on the developing secondary market as primary. Fourth, the denominator effect must be anticipated at the outset, with headroom and a reserve, not discovered in a crisis. Fifth, the cost of mis-provisioning is asymmetric, so the reserve should err toward prudence.

It is fitting to close the substantive analysis by returning to the point with which it began. An illiquid programme draws capital before it returns it, and the gap between the two is where investors are most often hurt, not by choosing the wrong investments but by being unable to fund the right ones. Everything in this paper, the forward model, the stressed reserve, the source hierarchy, the headroom for the denominator effect, the bias toward prudence, serves the single end of ensuring that the investor is never forced to act against its own interest for want of liquidity. That end is achievable with discipline rather than sophistication, and it is the difference between an illiquid allocation that quietly compounds its premium over a long horizon and one that is periodically forced into value-destroying actions by a liquidity squeeze. For a Gulf allocator building exposure across real estate, private credit, digital infrastructure and private equity, the liquidity discipline set out here is what makes the long-horizon harvest of the illiquidity premium safe to pursue.

Before turning to the limitations, it is worth drawing out the practical disposition the analysis recommends. Good liquidity management is, in the end, a habit of looking forward and provisioning prudently: maintaining the cash-flow model, sizing the reserve to a stress rather than a hope, arranging the back-stops before they are needed, and leaving headroom for the denominator effect. None of this is technically difficult, and that is rather the point. The failures that damage illiquid programmes are seldom failures of sophistication; they are failures of discipline, of provisioning to a comfortable central case, of assuming self-funding will arrive early, of treating the credit line or the secondary market as a reserve, of forgetting the denominator effect. The contribution of a structured framework is to make those disciplines explicit and routine, so that the investor is protected by habit rather than relying on judgement in the moment of stress, when judgement is hardest. An investor that builds these habits early will

find that liquidity management recedes from a recurring worry into a managed, almost quiet, part of running the programme, which is exactly what good liquidity management should feel like.

These conclusions are subject to limitations that also mark out the avenues for further work. The framework is stylised and illustrative, and a programme's actual cash flows will differ; the value is in the disciplines, not the specific numbers. The analysis treats one investor's programme in isolation, whereas a large allocator runs many programmes whose cash flows partly offset, which would refine the aggregate reserve. The Gulf secondary market is developing, and as it deepens the last resort in the hierarchy will become more reliable, which further work should revisit. And the framework does not model the interaction between the programme's liquidity and the investor's wider obligations, which a complete treatment would integrate. These are extensions rather than corrections: the central message, that liquidity management is the largest controllable operational risk in an illiquid programme and must be planned deliberately, is robust. An investor that anticipates its cash flows, sizes its reserve to a stressed peak, marshals a hierarchy of sources and plans for the denominator effect can bear the illiquidity of a Gulf programme safely, and so harvest its premium without being caught short.

DISCLOSURES

This paper is provided for general information and educational purposes only. It does not constitute investment, financial, legal, tax or other professional advice, nor an offer or solicitation to buy or sell any security or to pursue any investment strategy. The figures and scenarios are stylised, forward-looking illustrations used to convey liquidity-management mechanics; they are not forecasts, do not represent the past or expected performance of any actual programme, and should not be relied upon as such. Illiquid private-market investments carry the risk of capital calls and of being unable to realise positions when required. The author is an independent researcher; the views are his own and do not represent those of any institution. Readers should obtain professional advice tailored to their circumstances before making any decision.

REFERENCES

References

- [1] Ang, A. (2014). *Asset Management: A Systematic Approach to Factor Investing*. New York: Oxford University Press.
- [2] Ang, A., Papanikolaou, D. and Westerfield, M. M. (2014). Portfolio Choice with Illiquid Assets. *Management Science*, 60(11), 2737-2761.
- [3] Brown, G. W., Gredil, O. and Kaplan, S. N. (2019). Do Private Equity Funds Manipulate Reported Returns? *Journal of Financial Economics*, 132(2), 267-297.
- [4] Harris, R. S., Jenkinson, T. and Kaplan, S. N. (2014). Private Equity Performance: What Do We Know? *Journal of Finance*, 69(5), 1851-1882.
- [5] Kaplan, S. N. and Schoar, A. (2005). Private Equity Performance: Returns, Persistence, and Capital Flows. *Journal of Finance*, 60(4), 1791-1823.
- [6] Phalippou, L. (2014). Performance of Buyout Funds Revisited? *Review of Finance*, 18(1), 189-218.
- [7] Robinson, D. T. and Sensoy, B. A. (2016). Cyclicity, Performance Measurement, and Cash Flow Liquidity in Private Equity. *Journal of Financial Economics*, 122(3), 521-543.
- [8] Swensen, D. F. (2009). *Pioneering Portfolio Management: An Unconventional Approach to Institutional Investment (Fully Revised and Updated)*. New York: Free Press.
- [9] Takahashi, D. and Alexander, S. (2002). Illiquid Alternative Asset Fund Modeling. *Journal of Portfolio Management*, 28(2), 90-100.

ABOUT THE AUTHOR

About the Author

Chennakeshav Adya is an independent researcher and corporate-finance practitioner with more than 25 years of industry experience as both an operator and an investment banker. He holds a Bachelor of Engineering from Visvesvaraya Technological University (VTU), an MBA from London Business School, and is currently pursuing an LLM (Master of Laws) at University College London (UCL). His research focuses on capital formation, private markets and cross-border investment across the Gulf, the United Kingdom, Europe and the United States, with particular attention to real estate, private credit, digital infrastructure and alternative assets. He writes to bridge institutional investment practice and the practical questions facing allocators and capital seekers.

APPENDIX A. BASE-CASE ASSUMPTIONS

The table below records the stylised assumptions used in the framework. They are indicative, forward-looking calibrations chosen to illustrate liquidity-management mechanics, informed by the private-markets cash-flow literature; they are not forecasts and not estimates from a proprietary dataset.

Table A1. Full Base-Case Assumptions

Parameter	Base case	Rationale
Build / commitment period	5 years	Typical paced programme build
J-curve (net-negative) period	Years 1-4	Calls lead distributions
Self-funding onset	Years 5-6	Distributions overtake calls
Average annual net call	~25% of commitment	Illustrative paced average
Stressed peak net call	~60% of commitment	Calls bunch, distributions slow
Stress scenario	Calls +, distributions -	Combined market stress
Secondary market	Developing	Not a primary liquidity source
Reserve basis	Stressed peak + margin	Sized to peak not average

Indicative assumptions; see Section 3.

APPENDIX B. LIQUIDITY-MANAGEMENT CHECKLIST

The following checklist distils the framework into the questions an investor should answer when managing the liquidity of an illiquid Gulf programme.

1. Have we built a forward cash-flow model of the programme, projecting calls and distributions under conservative assumptions? Do we update the model as commitments are made and as actual calls and distributions arrive? Have we identified the net-call period and the cumulative trough, the maximum liquidity the programme will demand? Have we sized the reserve to a stressed peak net call with a margin, rather than to the average? Is the reserve held in genuinely liquid assets that will not be impaired in the stress that triggers a call? Is the reserve ring-fenced from other uses so it is available when needed? Have we arranged the back-stop sources, credit line, liquid sleeve, in advance, in calm conditions? Have we a realistic,

conservative view of what the developing secondary market would offer in a stress? Have we sized the illiquid allocation with headroom below any ceiling, to absorb the denominator effect? Have we a documented liquidity-stress plan and a clear owner for liquidity management? Do we monitor the illiquid weight against its ceiling and review reserve adequacy regularly? Do we resist running the reserve down to chase return in good times?

APPENDIX C. GLOSSARY OF KEY TERMS

The following glossary defines the principal technical terms used in this paper for readers who may be less familiar with private-markets liquidity management.

Capital call. A demand by a fund manager that an investor transfer a portion of its committed capital, issued as the manager finds investments to make.

Commitment. The total amount an investor contractually agrees to provide to a fund over its life, called in instalments rather than all at once.

Distribution. A return of capital and profit from a fund to its investors, paid as the fund realises its investments.

Net cash flow. In a given period, distributions received minus capital calls paid; negative while the programme is deploying, positive as it matures.

J-curve. The characteristic shape of a private programme's cumulative net cash flow, declining early as calls exceed distributions, then recovering.

Drawdown period. The early phase of a programme during which capital is being called and deployed; the net-negative period.

Self-funding. The state a paced programme reaches when its distributions are large enough to fund its ongoing calls without new external cash.

Liquidity reserve. Assets an investor holds specifically to meet capital calls and obligations without forced selling of other holdings.

Peak net call. The largest net cash outflow the programme is expected to demand in a single period under a stressed scenario; the basis for sizing the reserve.

Cumulative trough. The lowest point of the programme's cumulative net cash position; the maximum total liquidity the programme demands over its life.

Denominator effect. The mechanical rise in an illiquid allocation's portfolio weight when the value of liquid assets falls, pushing it above target without any change in the illiquid assets.

Over-provisioning. Holding more liquidity reserve than the programme requires, incurring a return drag on the idle assets.

Under-provisioning. Holding too little liquidity reserve, risking forced sales of other assets or default on a capital call.

Capital-call default. Failure to meet a capital call, which can trigger penalties, interest, and forfeiture of part or all of the investment.

Subscription / credit facility. A committed line of credit a fund or investor can draw on to bridge a timing mismatch between calls and available cash.

Liquid sleeve. The portion of a wider portfolio held in readily saleable assets, available as a back-stop source of liquidity.

Secondary market. The market in which existing private-fund interests or positions are bought and sold before the fund's natural end; developing in the Gulf.

Illiquidity premium. The additional return an investor expects to earn for bearing the inability to sell an asset quickly.

Pacing. The deliberate spreading of commitments across years and vintages to smooth the cash-flow profile and diversify timing risk.

Vintage year. The year in which a fund or commitment begins investing; diversifying across vintages spreads timing and market risk.

Stress scenario. A deliberately adverse set of assumptions, typically accelerated calls and slowed distributions, used to size the reserve.

Ring-fencing. Setting aside the reserve so it cannot be used for other purposes and is genuinely available to meet calls.