

Mezzanine and Preferred Equity for GCC Development: Reaching Target Leverage in the Capital Stack

How the mezzanine and preferred layers bridge the gap between senior debt and common equity, lift returns to target leverage, and are structured and priced for Gulf real-estate development

Chennakeshav (CK) Adya

Independent Researcher

June 2026

ABSTRACT

Abstract. A real-estate development is rarely funded by senior debt and common equity alone. Between the two sits a family of intermediate capital, mezzanine debt and preferred equity, that fills the gap between the level a senior lender will fund and the level of leverage a sponsor wishes to reach. This paper sets out what that intermediate layer is, why it exists, how it is priced and structured, and how it lifts the return to common equity at the cost of greater risk. The setting is Gulf Cooperation Council development as it stands in 2026, where deep regional pools of capital, an active development pipeline across residential, hospitality and mixed-use assets, and a maturing private-credit market have made the mezzanine and preferred layer a practical instrument rather than a theoretical one. The analysis is built on a transparent, stylised framework whose assumptions are stated and tested rather than hidden. It makes four contributions. First, it locates the mezzanine and preferred layer precisely in the capital stack and explains the economic gap it fills. Second, it shows how the layer reaches a target loan-to-cost without diluting the sponsor and what that does to the equity return and its risk. Third, it sets out the intercreditor architecture, the order of payment and the principal negotiated terms, that governs how the layers coexist. Fourth, it draws the implications for sponsors, lenders and capital providers in the Gulf context. The framework is illustrative and forward-looking: the figures are chosen to expose the mechanics, not to forecast any particular transaction. The central finding is that the intermediate layer is best understood not as expensive debt or cheap equity but as a distinct position with its own risk, price and protections, and that the discipline of structuring it well is what separates leverage that creates value from leverage that destroys it.

JEL Classification: G32, G31, G24, G23, R33

Keywords: mezzanine finance, preferred equity, capital stack, target leverage, loan-to-cost, intercreditor, real-estate development, GCC, private credit, subordination

1. INTRODUCTION

A real-estate development of any scale, a residential tower, a hospitality asset, a mixed-use district, is rarely funded by a single kind of capital. It is funded by a stack of capital, each layer with its own claim on the asset and its cash flows, its own price, and its own appetite for risk. At the base sits senior debt, the largest and cheapest layer, secured first and repaid first. At the top sits common equity, the smallest

and dearest layer, which controls the project, bears the first loss and keeps whatever is left once everyone else has been paid. Between the two sits a less familiar but increasingly important family of capital: mezzanine debt and preferred equity, the intermediate layer that this paper is about.

The intermediate layer exists because of a gap. A senior lender will fund a development only up to a level it judges safe, commonly a loan-to-cost in the region of fifty-five to sixty-five per cent for a development asset in the Gulf in 2026, lower for riskier phases and asset types. A sponsor, by contrast, often wishes to commit far less common equity than the remaining thirty-five to forty-five per cent, both because common equity is the most expensive capital it can raise and because committing it concentrates the sponsor's resources in a single project. The mezzanine and preferred layer fills the space between the senior ceiling and the sponsor's preferred equity cheque. It is more expensive than senior debt, because it ranks behind it and is less well secured, but cheaper than common equity, because it ranks ahead of it and carries less of the project's residual risk.

This paper sets out the intermediate layer in the round: what it is, why it exists, how it is priced, how it is structured, what it does to the return on common equity, and how the layers are made to coexist through the intercreditor arrangements that govern the order of payment. The treatment is deliberately practical. The mezzanine and preferred layer is not an abstraction; it is a negotiated instrument with a coupon, a maturity, a set of rights and a position in a waterfall, and the paper treats it as such.

1.1 The Gulf Context in 2026

The Gulf setting bears on the intermediate layer in several specific ways. The region's development pipeline remains deep across the major markets, with residential, hospitality and mixed-use assets being delivered against national diversification programmes and continued population growth. At the same time, the cost of senior debt has been shaped by the global rate environment of the preceding years, and senior lenders have remained disciplined on loan-to-cost. The combination, an active pipeline and a constrained senior ceiling, is precisely the condition under which the intermediate layer becomes useful.

Three further features of the 2026 Gulf market matter. First, the region's pools of capital, sovereign, family-office and institutional, are deep and increasingly willing to take structured positions rather than only senior debt or outright equity. Second, the private-credit market in the region has matured, so that a sponsor seeking a mezzanine tranche or a preferred-equity partner now has a genuine market to approach rather than a handful of bilateral relationships. Third, much regional development is structured to be compatible with Islamic finance, which shapes how the intermediate layer is documented even when its economics are familiar. The paper notes these features where they bear on the analysis without treating any single jurisdiction's law as definitive.

1.2 Contribution and Structure

The paper makes four contributions. First, it locates the mezzanine and preferred layer precisely in the capital stack and explains, in terms of risk and price, the gap it fills. Second, it shows how the layer is used to reach a target leverage without diluting the sponsor, and quantifies, in a stylised way, what that does to the return on common equity and to its risk. Third, it sets out the intercreditor architecture that allows the layers to coexist: the order of payment, the standstill and cure rights, and the principal negotiated terms. Fourth, it draws out the implications for the three parties to a Gulf development, the sponsor, the senior lender and the intermediate-capital provider.

It is worth saying why the intermediate layer deserves treatment as a strategic question rather than a documentation detail. The choice of how much mezzanine or preferred equity to use, and on what terms, determines the sponsor's return, the project's resilience to stress, and the alignment between the parties. A stack that reaches for too much leverage through an aggressively priced mezzanine tranche can convert a sound project into a fragile one; a stack that uses too little leaves return on the table and ties up scarce common equity. The layer is, in short, a lever, and like any lever it must be sized to the load.

The remainder of the paper is structured as follows. Section 2 reviews the relevant literature, on capital structure, on subordinated and intermediate capital, and on the financing of real-estate development, and states the propositions that the framework tests. Section 3 sets out the data and methodology, including the stylised framework and its assumptions. Section 4 presents the results: the layers and their prices, the way the intermediate layer reaches target leverage, the effect on the equity return, the shift in the stack across the project phase, and the intercreditor waterfall. Section 5 tests the robustness of the conclusions to changes in the assumptions and to stress. Section 6 sets out the implementation considerations for the three parties. Section 7 concludes.

2. LITERATURE REVIEW AND PROPOSITIONS

The analysis draws on three strands of thought: the literature on capital structure, which explains why a project divides its funding into layers of differing seniority; the literature on subordinated and intermediate capital, which explains how instruments that sit between debt and equity are priced and protected; and the literature on the financing of real-estate development, which sets out how the stack is assembled and how it shifts as a project moves from land to stabilised asset. The three are complementary, and the propositions that close this section draw on all of them.

2.1 Capital Structure and Subordination

The foundational literature on capital structure examines how the division of a project's funding between debt and equity affects its value and the return to its owners. The classical starting point is that, under idealised conditions, the division does not affect the total value of the project: value is determined by the assets and their cash flows, not by how the claims on them are sliced. The interest of the real world lies in the departures from those idealised conditions. Taxes favour debt, because its cost is commonly deductible where equity's is not. Financial distress and the costs that attend it favour equity, because a project that has promised too much to its creditors can be tipped into default by a shortfall that a less levered project would absorb. The optimal structure trades the two off, and the trade-off is the organising idea of the whole field.

Subordination is the device that makes a layered stack possible. By agreeing that one claim ranks behind another, the parties create a junior position that bears more of the project's risk and is compensated with a higher return, and a senior position that bears less and accepts a lower return. The literature on subordinated capital treats the junior claim as a hybrid: it has the fixed-claim character of debt to the extent of its coupon and its priority over equity, and the residual character of equity to the extent that it is exposed to the project's downside before the senior claim is touched. The mezzanine and preferred layer is the practical expression of this hybrid, and the literature's central insight, that each layer must earn a return commensurate with the risk of its position, governs the whole of the present analysis.

2.2 Mezzanine and Preferred Equity

A second strand concerns the instruments themselves. Mezzanine debt is, in form, a loan: it has a principal, a coupon and a maturity, and it ranks as a creditor's claim, albeit behind the senior loan. Its distinguishing features are its subordination, its higher coupon, and the frequent presence of a deferred or payment-in-kind component that allows interest to accrue rather than be paid in cash during the development phase, when the asset produces no income. Preferred equity is, in form, an equity interest: it ranks behind all debt, but ahead of common equity, and it commonly carries a preferred return that must be satisfied, whether by current payment or by accretion, before common equity receives anything. The literature treats the two as points on a spectrum rather than as wholly distinct instruments; the choice between them turns on tax, on the senior lender's requirements, on the law of the jurisdiction and, in the Gulf, on compatibility with Islamic-finance structures.

The pricing of the intermediate layer follows from its position. Because it ranks behind the senior loan, it bears loss before the senior loan does, and it must be compensated for that exposure with a return above the senior cost. Because it ranks ahead of common equity, it bears loss after common equity does, and it can accept a return below the common-equity cost. The literature describes the layer's return as the sum of a return for time and a return for the risk of its position, with the latter rising as the layer is pushed further up the stack and as the project's leverage increases. The same literature emphasises that the layer's protections, its rights on default, its controls over the senior debt, its security where it has any, are as much a part of its value as its coupon.

The choice between mezzanine debt and preferred equity, although the two sit close together in the stack, is not a matter of indifference. The distinction matters for at least four reasons. The first is the senior lender's tolerance: a senior lender may permit a layer of additional debt beneath it only on strict conditions, or may prefer that the gap be filled by an instrument that does not rank as debt at all, in which case preferred equity is the natural answer. The second is tax and deductibility, where the treatment of a coupon as interest rather than as a distribution can change the after-tax cost of the layer materially, though the significance of this varies across the Gulf jurisdictions. The third is control: preferred equity, being an equity interest, can carry governance rights, consent thresholds and step-in mechanics that a loan typically does not, which makes it attractive to a provider that wishes to influence the project rather than merely to lend to it. The fourth, specific to the region, is compatibility with Islamic finance, which can make an equity-based or profit-sharing structure simpler to document than a conventional interest-bearing loan. In practice many Gulf structures blend the two, using a subordinated facility for the portion nearest the senior debt and a preferred interest for the portion nearest the common equity, so that the labels matter less than the economics and the rights they carry.

2.3 Leverage and the Cost of Capital

A further strand of the literature examines how leverage affects the return to equity and the overall cost of capital, and it is directly relevant to the use of the intermediate layer. The central result is that leverage raises the expected return to equity, because a fixed-claim layer takes only its contracted return and leaves the residual to a smaller equity base, but that it also raises the risk of that return and the probability of distress. The relationship is not linear: as leverage rises, each additional unit adds more risk than the last, because the cushion of capital beneath the senior and intermediate layers thins and the project's ability to absorb a shortfall diminishes. The literature on the cost of financial distress explains why there is a limit to useful leverage, and why the limit is reached well before the point at which the base-case return is maximised. The intermediate layer is the practical instrument through which a

sponsor moves along this relationship, and the literature's caution, that the relationship is convex in risk, is the reason the present paper insists that leverage be sized to the stressed case rather than the base case.

2.4 Financing Real-Estate Development

A third strand concerns the financing of development specifically, where the stack is assembled against an asset that does not yet exist and a cash flow that has not yet begun. The development literature stresses that the right stack changes as the project de-risks. In the land and predevelopment phase, when the project's value is most uncertain, the stack is equity-heavy, because lenders will fund little against an unentitled or unbuilt asset. As construction completes and the asset is leased and stabilised, the stack becomes debt-heavy, because a stabilised asset with a contracted income can support far more leverage at a far lower cost. The intermediate layer is most useful in the middle of this journey, when the project is too advanced to be funded by common equity alone but too immature to support full senior leverage.

The development literature also stresses the role of the take-out, the refinancing or sale that repays the development capital once the asset is stabilised. The intermediate layer is, in most cases, a bridge to that take-out: it funds the gap during the riskiest phase and is repaid when the stabilised asset can be refinanced on senior terms or sold. This bridging character explains many of the layer's features, its shorter maturity, its deferred interest, its exit fees, and it is central to the way the present paper frames the layer's purpose.

2.5 Synthesis and Propositions

Taken together, these literatures imply that the intermediate layer is a distinct position with its own risk, price and protections, that it exists to fill the gap between the senior ceiling and the sponsor's equity, and that its use lifts the return to common equity while raising the project's fragility. The framework that follows tests five propositions drawn from this synthesis.

Proposition 1. Each layer of the stack, senior debt, mezzanine, preferred equity and common equity, earns a return commensurate with the risk of its position, so that required return rises monotonically with subordination.

Proposition 2. The mezzanine and preferred layer fills the gap between the senior loan-to-cost ceiling and the sponsor's target leverage, and does so at a blended cost between the senior and common-equity costs.

Proposition 3. Adding the intermediate layer to reach a higher target leverage lifts the expected return to common equity, but raises its risk and the probability of a poor or negative outcome under stress.

Proposition 4. The right stack shifts across the project phase, equity-heavy and preferred-heavy in development, senior-heavy once stabilised, and the intermediate layer is most valuable in the middle phases.

Proposition 5. The coexistence of the layers depends on the intercreditor arrangements, the order of payment and the negotiated rights, as much as on their pricing, and a well-structured intercreditor regime is what allows leverage to be added safely.

3. DATA AND METHODOLOGY

This study is analytical rather than empirical. It uses a transparent, stylised framework whose assumptions are stated openly so that a reader can see how each conclusion follows from its inputs and can substitute their own. The framework is not a model of any particular transaction, and its figures are illustrative. Its purpose is to expose the mechanics of the intermediate layer clearly enough that the propositions can be examined, and to do so in a way calibrated to Gulf development conditions as they stand in 2026.

3.1 The Capital-Stack Framework

The framework has four components. The first is a layer analysis that sets out, for each layer of the stack, its share of total cost, its indicative cost or required return, and its position in the order of payment. The second is a leverage analysis that shows how the intermediate layer is used to move from the senior ceiling to a target loan-to-cost, and what that does to the return on common equity under base and stressed cases. The third is a phase analysis that shows how the composition of the stack shifts as the project moves from land through construction and lease-up to stabilisation. The fourth is an intercreditor analysis that sets out the order of payment and the principal negotiated rights that allow the layers to coexist.

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 structure of the stack, not a pricing engine and not a forecast. It does not predict the cost of any layer in any market at any date; the costs it uses are indicative and chosen to expose relationships, not to quote a market. It does not model the cash flows of a specific asset; it abstracts from the particulars of any one development to show the mechanics that all developments share. It does not constitute advice on any transaction or any instrument, and nothing in it should be read as a recommendation to use or to avoid any particular structure.

The framework is organised around trade-offs rather than around a single optimum. There is no universally correct amount of mezzanine or preferred equity; the right amount depends on the sponsor's risk appetite, the senior lender's terms, the asset's phase and the market's pricing. The framework is built to make those trade-offs visible, so that a sponsor or a capital provider can reason about them, rather than to assert a conclusion that holds in every case.

3.2 Assumptions and Their Justification

The stylised inputs, the layer shares and costs, the leverage-return relationship, and the phase profile, are set out in Table 1 and used consistently throughout. They are chosen to be plausible for a Gulf development asset in 2026 and, more importantly, to be internally consistent, so that the relationships the framework illustrates hold regardless of whether any single figure matches a particular deal. The senior cost reflects a development-debt margin over the prevailing reference rate; the mezzanine cost reflects its subordination and its deferred component; the preferred return reflects its position behind all debt; and the common-equity return reflects the residual risk it bears. The leverage-return relationship is calibrated so that the equity return rises with leverage in the base case and falls, steeply, under stress.

Table 1. Base-Case Assumptions for the Capital-Stack Framework

Indicative, forward-looking assumptions used to illustrate capital-structure mechanics for a stylised GCC development asset in 2026; not quotes, forecasts or advice. All figures are illustrative.

Layer / parameter	Share of cost	Indicative cost / return	Position
Senior debt	55-60%	6-8% (single digit)	Repaid first; bears loss last
Mezzanine	10-14%	11-16%	Behind senior; ahead of equity
Preferred equity	11-15%	11-15%	Behind all debt; ahead of common
Common equity	15-20%	16-28%+	Paid last; bears first loss
Senior LTC ceiling	60%	n/a	Limit of senior funding
Target leverage	85% LTC	blended	Reached via intermediate layer

Two points about the assumptions deserve emphasis. First, the costs are gross and indicative; they are not quotes, and an actual transaction would price each layer against its own market and its own risk. Second, the relationships matter more than the levels. If a reader judges the senior cost too high or the mezzanine share too low, they may adjust the inputs; the framework's conclusions, that required return rises with subordination, that the intermediate layer reaches target leverage at a blended cost, that leverage lifts and risks the equity return, survive any reasonable adjustment.

4. RESULTS AND DISCUSSION

This section presents the framework in the order of the propositions: the layers and their prices, the way the intermediate layer reaches target leverage, the effect on the return to common equity, the shift in the stack across the phase, and the intercreditor waterfall that allows the layers to coexist.

4.0a Why the Same Asset Yields Different Returns

Before setting out the layers, it is worth establishing the proposition that motivates the whole analysis: that the same asset, producing the same cash flow, yields different returns to different providers of capital, because they hold different positions. A senior lender to a development earns a single-digit return, because it is repaid first and bears loss last. A common-equity holder in the same development may earn a high-teens or twenties return, because it is paid last and bears loss first. The asset has not changed; the position has. The intermediate layer occupies the ground between, and earns a return between, precisely because it occupies a position between. This is not a market inefficiency to be arbitrated; it is the price of risk, correctly charged.

The same proposition carries a warning that runs through the analysis. Because the junior positions are paid last, they are the first to be impaired when the project disappoints. Leverage that lifts the equity return in the base case magnifies the loss in the stressed case. The intermediate layer is a tool for reaching higher returns, but it is also a tool for reaching them with less of one's own capital at stake, and the two are the same thing seen from opposite sides.

4.1 The Layers of the Stack

The stack is built from layers of differing risk and cost. Figure 1 sets out their share of total cost and their indicative cost or required return for a stylised Gulf development.

The Development Capital Stack: Share and Cost




Figure 1. The Development Capital Stack: Share and Cost

Illustrative; senior debt is the largest and cheapest layer; common equity the smallest and dearest, with the intermediate layer priced between.

The figure supports Proposition 1. The stack runs from senior debt, the largest layer at around fifty-five per cent of cost and the cheapest at a single-digit return, through the intermediate layers, mezzanine at a low-double-digit cost and preferred equity at a similar level, to common equity, the smallest layer and the dearest, with a required return in the twenties. Required return rises with subordination at every step, because each step down the stack means being paid later and bearing loss sooner. The intermediate layer is, by share, modest, around a quarter of the cost between them, but it is precisely the layer that allows the sponsor to reach its target leverage without committing the full balance as common equity.

A practical implication of the risk-return ladder is that the layers must be priced and sized together, not separately. A senior lender's willingness to lend at a low cost depends on there being a substantial cushion of junior capital beneath it; a mezzanine provider's willingness to lend depends on there being common equity beneath it in turn. The stack is an interlocking structure, and a change to any one layer ripples through the others.

4.2 How the Intermediate Layer Reaches Target Leverage

The central purpose of the intermediate layer is to bridge the gap between the senior ceiling and the sponsor's target leverage. Figure 2 shows the bridge for a stylised development that targets eighty-five per cent loan-to-cost against a senior ceiling of sixty per cent.




Figure 2. Reaching Target Leverage with the Mezzanine and Preferred Layer

Illustrative; the mezzanine and preferred layers fill the gap between the senior loan-to-cost ceiling and the sponsor's target leverage.

The figure supports Proposition 2. The senior loan funds the first sixty per cent of cost, the limit the senior lender will reach against a development asset. The sponsor wishes to reach eighty-five per cent of cost in capital that ranks ahead of, or alongside, its common equity, leaving only fifteen per cent to be funded by the common equity that bears the first loss. The mezzanine and preferred layers fill the twenty-five-percentage-point gap, mezzanine taking the portion nearest the senior debt and preferred equity the portion nearest the common equity. The result is that the sponsor reaches its target leverage while committing common equity equal to only fifteen per cent of cost, far less than the forty per cent it would have committed had it funded the gap itself.

The blended cost of the stack is the share-weighted average of the layer costs. Because the intermediate layer is priced between the senior and common-equity costs, adding it to reach target leverage raises the blended cost above the senior cost but holds it below the cost of an all-common-equity gap. The sponsor accepts a higher blended cost of capital in exchange for a smaller common-equity cheque and, as the next subsection shows, a higher return on that cheque. Whether the trade is worthwhile depends on the spread between the project's expected return and the cost of the intermediate layer, and on the sponsor's tolerance for the risk that leverage adds.

4.2a A Worked Illustration

A simple worked illustration makes the mechanics concrete. Consider a stylised Gulf development with a total cost of one thousand units. A senior lender funds six hundred units, sixty per cent of cost, at a cost of seven per cent. Without an intermediate layer, the sponsor would fund the remaining four hundred units as common equity. With the intermediate layer, the sponsor instead raises one hundred and twenty units of mezzanine at a cost of fourteen per cent and one hundred and thirty units of preferred equity at a cost of thirteen per cent, leaving only one hundred and fifty units, fifteen per cent

of cost, to be funded by common equity. The blended cost of the whole stack is the share-weighted average: roughly seven per cent on sixty per cent, fourteen on twelve, thirteen on thirteen, and the common-equity cost on the remaining fifteen, which produces a blended cost modestly above the senior cost but well below an all-equity gap.

The value of the structure to the sponsor lies in what happens to the common-equity return. Suppose the project, once stabilised and sold, returns its total cost plus a surplus. In the all-equity case, that surplus is spread over four hundred units of common equity. In the levered case, the senior and intermediate layers take only their contracted returns, and the remaining surplus accrues to just one hundred and fifty units of common equity. The same surplus, divided over a smaller base, produces a higher return per unit of common equity. This is the arithmetic of leverage, and it is the reason the intermediate layer is used. The illustration also shows the limit: if the project returns less than the contracted cost of the senior and intermediate layers, those layers must still be paid, and the shortfall falls entirely on the one hundred and fifty units of common equity, which can be wiped out by a loss that the four-hundred-unit equity base would have absorbed. The figures here are illustrative and are chosen only to expose the arithmetic.

4.3 Return by Tranche

Each tranche earns a return commensurate with its position. Figure 3 sets out the indicative gross return for each layer as a range rather than a point, because the return to any tranche depends on the asset, the phase and the market.




Figure 3. Required Return by Tranche

Illustrative; required return rises with subordination, and the range of outcomes widens for the more junior positions.

The figure reinforces Proposition 1 by setting out the ladder explicitly. Senior debt earns the lowest return, a single-digit range that compensates it for time and for a small, well-secured credit risk. Mezzanine earns a low-double-digit range, compensating it for ranking behind the senior loan and for the deferred character of much of its return. Preferred equity earns a similar range, with its return depending heavily on whether its preferred coupon is paid currently or accretes. Common equity earns the widest and highest range, because it captures the project's upside and bears its first loss; its return can be very high if the project succeeds and negative if it does not. The widening of the ranges down the

stack is itself informative: the junior the position, the more its outcome depends on how the project actually performs.

4.4 How Leverage Lifts and Risks the Equity Return

The central structuring trade-off is leverage. Figure 4 shows its effect on the return to common equity in a base case and a stressed case as total leverage rises from zero to the eighty-five per cent target.




Figure 4. How Leverage Lifts and Risks the Equity Return

Illustrative; higher leverage raises the base-case equity return but lowers, and ultimately turns negative, the stressed-case return.

The figure supports Proposition 3. In the base case, the equity internal rate of return rises as leverage increases, because a given project return is spread over a smaller common-equity base while the intermediate and senior layers take only their contracted return. At the eighty-five per cent target, the base-case equity return is materially higher than it would be with no leverage at all. But the stressed case tells the other half of the story. As leverage rises, the stressed-case equity return falls, and beyond a point it turns negative: the contracted returns to the senior and intermediate layers must be paid whether or not the project performs, and at high leverage a disappointing outcome leaves nothing, or less than nothing, for common equity. The gap between the two lines is the risk that leverage adds, and it widens as leverage rises.

The practical consequence is that the sponsor should think in terms of a leverage it can survive rather than a leverage that maximises the base-case return. The intermediate layer makes high leverage attainable, but attainable is not the same as advisable. The discipline of the structure lies in choosing a target leverage at which the stressed case is survivable, so that a disappointing outcome is a poor return rather than a default. The same figure that makes the case for the intermediate layer also marks its limit.

4.5 How the Stack Shifts Across the Phase

The right stack changes as the project de-risks. Figure 5 shows the shift from the land and predevelopment phase, through construction and lease-up, to stabilisation.




Figure 5. How the Capital Stack Shifts Across the Project Phase

Illustrative; the stack is equity- and preferred-heavy in development and senior-heavy once the asset is stabilised.

The figure supports Proposition 4. In the land and predevelopment phase, when the project's risks are highest and lenders will fund least, the stack is equity-heavy and preferred-heavy: common and preferred equity together fund the majority of cost, and senior debt is modest. As construction proceeds and the asset takes shape, senior debt grows and the equity layers shrink. By stabilisation, when the asset is built, leased and producing a contracted income, the stack is debt-heavy: a stabilised asset supports far more senior leverage at a far lower cost, and the intermediate and common-equity layers can be partly repaid through a refinancing or sale. The intermediate layer is most valuable in the middle phases, where it bridges a project that is too advanced for equity alone but too immature for full senior leverage.

The phase shift also changes which risks the structure must guard against. In development, the dominant risks are completion and cost; the structure must ensure the project can be finished. In lease-up, the dominant risk is absorption; the structure must survive a slower-than-expected path to stabilisation. Once stabilised, the dominant risk is refinancing; the structure must ensure the take-out can be executed. The intermediate layer's terms, its maturity, its deferred interest, its exit fee, are shaped by this sequence, because the layer's job is to bridge from the development risks to the take-out.

4.6 The Intercreditor Waterfall

The layers coexist only because their relative rights are defined in advance. Figure 6 sets out the order in which cash flow is applied, the waterfall that governs who is paid when.

Intercreditor Cash-Flow Waterfall




Figure 6. Intercreditor Cash-Flow Waterfall

Illustrative; cash flow is applied in order of priority, from senior debt service down to common-equity distributions.

The figure supports Proposition 5. Cash flow from the project is applied first to senior debt service, then to senior amortisation and the funding of any reserve, then to the mezzanine layer's current pay, then to its deferred or payment-in-kind accrual, then to the preferred return, and only then to distributions to common equity. The waterfall is the operational expression of subordination: it is the mechanism by which the junior layers' agreement to rank behind the senior layers is given effect month by month. The order is not merely a matter of who is paid in a default; it governs the ordinary flow of cash through the life of the project, and it is the single most important term the parties negotiate.

Around the waterfall sits a structure of rights that protects each layer's position. The senior lender commonly restricts the intermediate layer's ability to accelerate or enforce while the senior debt is outstanding, through a standstill. The intermediate layer commonly secures the right to cure a senior default, so that it is not wiped out by a default it could have remedied, and sometimes the right to purchase the senior debt at par. The preferred-equity holder commonly secures consent rights over major decisions and a mechanism to take control if its preferred return is not met. These rights, examined in Section 6, are as much a part of the layer's value as its coupon, and a stack whose intercreditor terms are poorly drawn can fail even when its pricing is sound.

5. ROBUSTNESS AND SCENARIO ANALYSIS

A stylised framework earns its keep only if its conclusions survive reasonable changes to its assumptions. This section tests the framework against changes in the cost of the layers, the senior ceiling, and the project's performance, and asks in each case whether the propositions hold.

5.1 Sensitivity to the Cost of the Intermediate Layer

The most uncertain input is the cost of the intermediate layer, which moves with the market's appetite for subordinated real-estate risk. If the mezzanine and preferred costs are higher than the base case, the blended cost of the stack rises and the case for reaching a high target leverage weakens, because the spread between the project's return and the cost of the marginal layer narrows. If they are lower, the case strengthens. The proposition that the layer reaches target leverage at a blended cost between the senior and common-equity costs holds regardless, because it follows from the ordering of the costs rather than their levels; what changes is whether the trade is worthwhile, and that is exactly the judgement the framework is meant to inform.

5.2 Sensitivity to the Senior Ceiling

The senior loan-to-cost ceiling determines the size of the gap the intermediate layer must fill. A higher senior ceiling, as might prevail for a lower-risk asset or in an easier credit market, leaves a smaller gap and a smaller role for the intermediate layer; a lower ceiling, as might prevail for a riskier phase or a tighter market, leaves a larger gap and a larger role. In the Gulf in 2026, with senior lenders disciplined on development loan-to-cost, the gap is wide enough that the intermediate layer is genuinely useful, which is why the layer has developed as a market rather than remaining a bilateral arrangement. The proposition that the layer fills the gap is unaffected by the ceiling's level; the ceiling sets the size of the layer, not its existence.

5.3 Stress Testing the Equity Return

The most important robustness test is the stressed case already shown in Figure 4. The framework's central caution, that leverage lifts the base-case return but raises the risk of a poor or negative outcome, is not a feature of any single set of inputs; it follows from the fixed-claim character of the senior and intermediate layers. Under any plausible calibration, there is a leverage beyond which the stressed-case equity return turns negative, and the location of that point, not the height of the base-case return, is what should discipline the choice of target leverage. A sponsor that sizes its leverage to the stressed case rather than the base case will choose a lower target than the one that maximises the base-case return, and will be glad of it when the project disappoints.

5.4 The Value of the Layer's Protections

A robustness analysis that looked only at pricing would miss half of what makes the intermediate layer work. The layer's protections, its rights on default, its controls over the senior debt, its security where it has any, are what allow it to accept a return below the common-equity cost despite bearing real risk. A mezzanine provider with a robust cure right, the ability to remedy a senior default and so prevent the senior lender from enforcing and wiping out the junior position, holds a materially less risky claim than one without, even at the same coupon. A preferred-equity holder with a step-in right, the ability to take control of the project if its preferred return is not met, holds a materially better-protected claim than one whose only recourse is to wait. The implication for the framework is that the cost of the layer cannot be read off its position alone; it depends on the protections, and two layers at the same point in the stack can be priced very differently if their protections differ. A robust structure is one in which price and protection are matched, so that the provider is neither overpaid for a well-protected position nor underpaid for an exposed one.

The same point bears on the senior lender and the sponsor. For the senior lender, the value of the junior cushion depends on the junior layers being unable to disrupt the senior position, which is why the

standstill is central. For the sponsor, the value of the structure depends on retaining enough control that the project can be managed through a difficult period, which is why the consent thresholds and step-in triggers in the preferred-equity terms are negotiated so carefully. The robustness of a structure, in short, is as much a matter of its rights as of its pricing, and a stress test that ignored the rights would give a false sense of either security or danger.

5.5 Limitations

The framework has limitations that bound its conclusions. It is stylised, and abstracts from the particulars of any asset, market or jurisdiction. It treats the cost of each layer as given, where in practice the costs are negotiated and interdependent. It does not model the timing of cash flows in detail, and so understates the importance of the deferred and payment-in-kind features that matter so much in development. It says nothing about the legal enforceability of the intercreditor terms in any particular jurisdiction, which in the Gulf can turn on the choice of law and the structure used to achieve compatibility with Islamic finance. These limitations do not undermine the propositions, which concern relationships rather than levels, but they mean the framework is a guide to reasoning rather than a substitute for transaction-specific analysis.

6. IMPLEMENTATION CONSIDERATIONS

The framework's conclusions translate into practical considerations for the three parties to a Gulf development: the sponsor that assembles the stack, the senior lender that anchors it, and the intermediate-capital provider that fills the gap. This section sets out the considerations for each, and then the principal negotiated terms that bind them.

6.1 For the Sponsor

For the sponsor, the intermediate layer is a tool for reaching a target leverage without committing the full common-equity cheque, and for raising the return on the equity it does commit. The sponsor's discipline lies in sizing the layer to a leverage the project can survive under stress, in negotiating intercreditor terms that preserve its control of the project, and in matching the layer's maturity and deferral to the project's path to a take-out. The sponsor should treat the blended cost of capital, not the cost of any single layer, as the relevant hurdle, and should be wary of reaching for leverage that lifts the base-case return at the cost of the project's resilience.

6.2 For the Senior Lender

For the senior lender, the presence of a substantial intermediate layer beneath it is a benefit: it provides a cushion of junior capital that bears loss before the senior debt is touched, and it aligns a well-capitalised junior party with the project's success. The senior lender's concern is to ensure that the intermediate layer cannot disrupt the senior position, through a standstill that prevents the junior layers from enforcing while the senior debt is outstanding, and through controls on the junior layers' rights in a default. A senior lender that is comfortable with the intercreditor regime can often lend further or more cheaply than it would against common equity alone, because the junior cushion reduces its risk.

6.3 For the Intermediate-Capital Provider

For the intermediate-capital provider, the layer is an opportunity to earn a return above senior debt for a risk below common equity, provided the position is properly protected. The provider's discipline lies in

pricing the layer to its true position in the stack, in securing the rights, to cure, to consent, to step in, that protect it against a default it did not cause, and in ensuring that the project's path to a take-out is realistic, because the layer's repayment depends on it. The provider should be especially attentive to the deferred and payment-in-kind features, which determine how much of its return is at risk through the development phase, and to the exit, which determines whether and when it is repaid.

6.4 The Principal Negotiated Terms

Table 2. Indicative Intercreditor Terms by Layer

Illustrative summary of the principal negotiated terms by layer; actual terms are transaction-specific and jurisdiction-dependent. All figures and terms are illustrative.

Term	Mezzanine	Preferred equity
Form	Subordinated loan	Equity interest
Return	Coupon (current + PIK)	Preferred return (current or accreting)
Ranking	Behind senior debt	Behind all debt
Security	Often share/second-lien	Typically unsecured equity rights
Default rights	Standstill; cure; purchase option	Consent rights; step-in on shortfall
Maturity	Bridges to take-out	Bridges to take-out / exit event

Whatever the form of the intermediate layer, a recognisable set of terms is negotiated in every transaction. The coupon and its split between current pay and deferred or payment-in-kind determine the cash burden the layer places on the project during development. The maturity and any extension options determine the time the layer allows for the take-out. The intercreditor terms, the standstill, the cure rights, the purchase option, the consent rights, determine how the layer coexists with the senior debt and the common equity. The security, where the layer has any, and the controls on changes to the senior debt determine its protection in a default. The exit fee or minimum return determines the floor on the provider's return. And, in the Gulf, the structure used to achieve compatibility with Islamic finance, where required, shapes how all of these are documented. The art of structuring the layer lies in setting these terms so that each party's position is fair, clear and enforceable.

6.5 Gulf Market Practice in 2026

Several features of Gulf market practice in 2026 shape how the intermediate layer is used in the region, and they deserve note alongside the general considerations. The first is the prevalence of structures compatible with Islamic finance. Where a transaction is to be Shariah-compliant, the intermediate layer is commonly documented through profit-sharing or lease-based structures rather than as a conventional interest-bearing loan, so that the economics familiar from a mezzanine tranche are achieved through a different legal form. This affects documentation and enforceability more than it affects the underlying risk and return, but it is a first-order consideration for any sponsor or provider operating in the region.

The second feature is the depth and character of the regional capital base. Sovereign and family-office capital in the Gulf is increasingly willing to take structured, intermediate positions, both because such positions offer a return above senior debt for a risk below common equity and because they allow a large pool of capital to be deployed across many projects without taking outright development risk in any one. The growth of this appetite is part of what has turned the intermediate layer from a bilateral

arrangement into a market. The third feature is the maturing of the regional private-credit market, which has given sponsors a genuine set of counterparties to approach for a mezzanine or preferred tranche, and has begun to standardise the terms on which such capital is provided. The fourth is the continued discipline of senior lenders on development loan-to-cost, which keeps the gap wide enough that the intermediate layer remains genuinely useful rather than marginal. Together these features make the Gulf in 2026 a market in which the intermediate layer is a practical instrument, well suited to the region's deep pipeline and deep pools of capital.

6.6 A Structuring Checklist

The considerations above can be reduced to a short checklist, set out in Appendix B, that a sponsor or capital provider can use to interrogate a proposed structure. The checklist asks, in sequence, whether the target leverage is survivable under stress, whether the layer is priced to its position, whether the intercreditor terms protect each party, whether the layer's maturity and deferral match the path to a take-out, and whether the documentation is enforceable in the relevant jurisdiction. A structure that answers each question well is more likely to convert leverage into value than one that does not.

7. CONCLUDING COMMENTS

The mezzanine and preferred layer is best understood not as expensive debt or as cheap equity but as a distinct position in the capital stack, with its own risk, its own price and its own protections. It exists to fill the gap between the level a senior lender will fund and the level of leverage a sponsor wishes to reach, and it does so at a blended cost between the senior and common-equity costs. Used well, it allows a sponsor to reach a target leverage without committing the full common-equity cheque, and lifts the return on the equity it does commit. Used carelessly, it converts a sound project into a fragile one, because the leverage that lifts the base-case return magnifies the loss under stress.

The Gulf setting in 2026 is one in which the layer has become a practical instrument. A deep development pipeline, a disciplined senior-debt market that leaves a genuine gap to be filled, deep regional pools of capital willing to take structured positions, and a maturing private-credit market have together turned the intermediate layer from a bilateral arrangement into a market that a sponsor can approach. The propositions this paper has examined, that required return rises with subordination, that the layer reaches target leverage at a blended cost, that leverage lifts and risks the equity return, that the stack shifts across the phase, and that the layers coexist only through their intercreditor terms, hold across the reasonable range of assumptions and describe the mechanics that any Gulf development must navigate.

The central lesson is one of discipline. The intermediate layer is a lever, and a lever must be sized to the load. The sponsor that sizes its leverage to the stressed case rather than the base case, that prices each layer to its position, and that negotiates intercreditor terms that protect every party, will find the layer a tool for creating value. The sponsor that reaches for the maximum attainable leverage will find that the same layer that lifted its base-case return has left its project unable to survive a disappointment. The difference between the two is not the instrument but the structuring, and it is the structuring, rather than the instrument, that this paper has sought to illuminate.

DISCLOSURES

This paper is a work of independent research and analysis. It is intended for general information and to contribute to professional and academic discussion. It does not constitute investment, financial, legal, tax or accounting advice, and it is not a recommendation, offer or solicitation to buy, sell or structure any security, instrument or transaction. The framework and figures are stylised, illustrative and forward-looking; they are chosen to expose mechanics rather than to forecast or to quote any market, asset or transaction, and they should not be relied upon as such. All figures in this paper are illustrative. Any party considering a financing structure of the kind discussed should take its own professional advice tailored to its circumstances and to the law of the relevant jurisdiction. The author holds no position dependent on the views expressed and received no compensation for the views expressed in this paper. The views are the author's own.

REFERENCES

- [1] F. Modigliani and M. H. Miller, "The Cost of Capital, Corporation Finance and the Theory of Investment," *The American Economic Review*, vol. 48, no. 3, pp. 261-297, 1958.
- [2] S. C. Myers, "The Capital Structure Puzzle," *The Journal of Finance*, vol. 39, no. 3, pp. 575-592, 1984.
- [3] M. C. Jensen, "Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers," *The American Economic Review*, vol. 76, no. 2, pp. 323-329, 1986.
- [4] S. C. Myers and N. S. Majluf, "Corporate Financing and Investment Decisions When Firms Have Information That Investors Do Not Have," *Journal of Financial Economics*, vol. 13, no. 2, pp. 187-221, 1984.
- [5] B. Esty, *Modern Project Finance: A Casebook*. Hoboken, NJ: John Wiley and Sons, 2004.
- [6] S. M. Geltner, N. G. Miller, J. Clayton and P. Eichholtz, *Commercial Real Estate Analysis and Investments*, 3rd ed. Mason, OH: OnCourse Learning, 2014.
- [7] A. Damodaran, *Applied Corporate Finance*, 4th ed. Hoboken, NJ: John Wiley and Sons, 2015.
- [8] G. Brealey, S. C. Myers and F. Allen, *Principles of Corporate Finance*, 13th ed. New York, NY: McGraw-Hill Education, 2020.
- [9] International Capital Market Association, *Sukuk and Islamic Finance: Market Practice and Structures*. Zurich: ICMA, 2024.
- [10] Bank for International Settlements, "Private credit and the financing of the real economy," *BIS Quarterly Review*, Basel, 2025.

ABOUT THE AUTHOR

Chennakeshav (CK) Adya is an independent researcher writing on capital structure, private credit and the financing of real assets, with a particular focus on the markets of the Gulf Cooperation Council. His work examines how layered capital structures, the interaction of senior debt, mezzanine, preferred equity and common equity, are assembled and priced in development and infrastructure finance, and how those structures are adapted to the legal and market conditions of the region. He writes for a professional and academic readership and approaches each topic through transparent, clearly stated frameworks rather than proprietary models.

APPENDIX A. BASE-CASE ASSUMPTIONS

The table below records the indicative, internally consistent assumptions used to illustrate the capital-stack mechanics in this paper. They are forward-looking and illustrative, calibrated to a stylised Gulf development asset in 2026, and are not quotes, forecasts or advice.

Table 1. Base-Case Assumptions for the Capital-Stack Framework

Indicative, forward-looking assumptions used to illustrate capital-structure mechanics for a stylised GCC development asset in 2026; not quotes, forecasts or advice. All figures are illustrative.

Layer / parameter	Share of cost	Indicative cost / return	Position
Senior debt	55-60%	6-8% (single digit)	Repaid first; bears loss last
Mezzanine	10-14%	11-16%	Behind senior; ahead of equity
Preferred equity	11-15%	11-15%	Behind all debt; ahead of common
Common equity	15-20%	16-28%+	Paid last; bears first loss
Senior LTC ceiling	60%	n/a	Limit of senior funding
Target leverage	85% LTC	blended	Reached via intermediate layer

APPENDIX B. STRUCTURING CHECKLIST

The checklist below distils the implementation considerations of Section 6 into a sequence of questions a sponsor or capital provider can use to interrogate a proposed mezzanine or preferred-equity structure.

- **Survivability.** Is the target leverage survivable under a reasonable stress case, so that a disappointing outcome is a poor return rather than a default?
- **Pricing.** Is each layer priced to its true position in the stack, neither over-reaching for return that the position does not bear nor under-pricing risk that it does?
- **Intercreditor.** Do the intercreditor terms, the standstill, the cure and purchase rights, the consent rights, protect each party's position clearly and enforceably?
- **Maturity and deferral.** Do the layer's maturity, extension options and deferred or payment-in-kind features match the project's realistic path to a take-out?
- **Security and controls.** Does the layer have the security and the controls on changes to the senior debt that its position warrants?
- **Exit.** Is the exit, the refinancing or sale that repays the layer, realistic in timing and in pricing, and is the provider's minimum return defined?
- **Enforceability.** Is the documentation enforceable in the relevant jurisdiction, including any structure used to achieve compatibility with Islamic finance?

APPENDIX C. GLOSSARY OF KEY TERMS

Capital stack. The set of layers of capital that fund a project, ordered by seniority from senior debt at the base to common equity at the top.

Senior debt. The most senior and best-secured layer of capital, repaid first and bearing loss last, and so the cheapest layer.

Mezzanine debt. Subordinated debt that ranks behind senior debt but ahead of equity, commonly with a higher coupon and a deferred or payment-in-kind component.

Preferred equity. An equity interest that ranks behind all debt but ahead of common equity, commonly carrying a preferred return that must be satisfied before common equity is paid.

Common equity. The most junior layer of capital, which controls the project, bears the first loss and keeps the residual once all other layers are satisfied.

Loan-to-cost (LTC). The ratio of a layer or layers of financing to the total cost of a project, used to express leverage in development finance.

Target leverage. The level of total leverage, expressed as a loan-to-cost, that a sponsor wishes to reach across its debt and intermediate layers.

Intercreditor agreement. The agreement among the providers of capital that defines their relative rights, including the order of payment and the rights on default.

Standstill. A provision restricting a junior creditor's ability to accelerate or enforce its claim while the senior debt is outstanding.

Payment-in-kind (PIK). A feature under which interest accrues and is added to principal rather than paid in cash, common in development financing where the asset produces no income during construction.

Take-out. The refinancing or sale that repays development capital once an asset is built, leased and stabilised.

Waterfall. The defined order in which cash flow from a project is applied to the layers of capital.