

Forward Contracts and Synthetic Pre-IPO Exposure: Structure and Counterparty Risk

How forwards and synthetic instruments deliver late-stage private exposure without owning the share, and the counterparty risk the buyer accepts in return.

Chennakeshav Adya

Independent Researcher

June 2026

ABSTRACT

Abstract. A buyer who wants exposure to a late-stage private company before it lists does not always have to buy the share. A growing share of pre-IPO secondary activity is transacted not as a transfer of stock on the issuer's register but as a contract that references the share: a forward, a profit-participation note, a total-return arrangement or a layered special-purpose vehicle. These synthetic structures solve real problems. They reach companies whose transfer restrictions make a direct purchase impossible, they let an intermediary aggregate small tickets into a single position, and they can be assembled quickly. They also introduce a risk that a direct purchase does not carry to the same degree, namely that the value the buyer ultimately receives depends on the solvency and the good faith of the party on the other side of the contract. This paper sets out the anatomy of forward and synthetic pre-IPO exposure, distinguishes the economic exposure from the legal claim, and offers a structured method for pricing and containing counterparty risk. It develops five propositions: that the buyer of a synthetic must first identify what it actually owns; that the headline economics are the same as a direct holding only when settlement performs; that counterparty risk is the price of the convenience and must be priced explicitly; that collateral, netting and legal venue convert a promise into a recoverable claim; and that some access is not worth having on synthetic terms. The framework is calibrated to Gulf Cooperation Council buyers and intermediaries operating in 2026 conditions. It is an educational and analytical document, not investment, legal or tax advice, and the figures are illustrative rather than sourced. The aim is to let a buyer or a broker reason clearly about an instrument that looks like a share, behaves like a share when all goes well, and behaves like an unsecured loan to a stranger when it does not.

JEL Classification: G12, G13, G23, G32, K12, D82

Keywords: pre-IPO secondaries, forward contracts, synthetic exposure, counterparty risk, total return swap, special-purpose vehicle, private markets, GCC, family office, collateral, netting

1. INTRODUCTION

A buyer who is offered exposure to a late-stage private company is usually offered a share. The transaction is presented as a transfer of stock from an existing holder, an employee or an early investor, to the buyer, with the company's register updated to reflect the new owner. This is the picture most buyers carry in their heads, and for a large part of the pre-IPO secondary market it is accurate. But a material and growing share of the market does not work this way at all. The buyer is offered not the share but a contract that tracks the share, an instrument whose value rises and falls with the company's fortunes but which leaves the share itself in someone else's hands, or in no one's hands at all until the moment of settlement.

These instruments go by several names. A forward contract fixes today the terms on which value will change hands at a future date tied to the company's listing or sale. A profit-participation note or a total-return arrangement promises the buyer the economic return on a reference share without conveying title to it. A special-purpose vehicle holds the share and sells the buyer an interest in the vehicle rather than in the company. Each of these is, in the language of markets, synthetic: the buyer holds the economics of the

share synthesised out of a contract, not the share itself. The distinction sounds technical. It is in fact the most important single fact about the position the buyer is taking, and it is the fact an offer document is most likely to blur.

This paper argues that synthetic pre-IPO exposure should be understood as two things bundled together. The first is the economic exposure to the underlying company, which is what the buyer wants and what the offer advertises. The second is a credit exposure to the counterparty that stands between the buyer and the company, which is what the buyer also receives whether or not the offer mentions it. When the counterparty is sound and settlement performs, the two are indistinguishable and the buyer has what it paid for. When the counterparty fails, the economic exposure evaporates and the buyer is left holding a claim against a party that cannot pay. The whole discipline of buying synthetic exposure consists of measuring the second exposure and refusing to let it be hidden by the first.

The paper is written with two readers in mind. The first is the intermediary, the broker, investment banker or platform who structures and syndicates these positions and who needs to be able to explain honestly what is being sold. The second is the family office or professional buyer who is shown a synthetic offer and must decide whether the access it provides is worth the counterparty risk it carries. The calibration reflects Gulf Cooperation Council conditions in 2026, where appetite for US late-stage technology exposure is strong, where direct access to the most sought-after names is scarce, and where synthetic structures are increasingly the route through which that appetite is met.

Three caveats frame the analysis. First, this is an educational and analytical document. It is not investment, legal, tax or accounting advice, and nothing in it is a recommendation to enter into or refrain from any transaction. Second, the figures, payoffs and risk gradings are illustrative. They are calibrated to make the structure legible, not to describe any real instrument, counterparty or company. Third, the law governing these contracts varies by jurisdiction, and the enforceability of a synthetic claim is a question for qualified counsel in the relevant venue, not for a whitepaper.

The paper makes four contributions. First, it sets out a clear anatomy of forward and synthetic pre-IPO structures and shows what the buyer owns in each. Second, it separates the economic payoff from the credit claim and shows why the two diverge precisely when it matters. Third, it offers a structured method for pricing counterparty risk and for converting an unsecured promise into a recoverable claim through collateral, netting and venue. Fourth, it provides a decision framework that tells a buyer when synthetic exposure is defensible and when the only honest answer is to decline.

The remainder of the paper is organised as follows. Section 2 reviews the relevant literature on forward pricing, credit and counterparty risk, collateral and the design of contracts under information asymmetry, and derives five propositions. Section 3 sets out the analytical framework and its assumptions. Section 4 develops the framework across structure, payoff, counterparty risk, mitigation and the comparison with a direct holding. Section 5 brings the analysis together in a decision matrix, a worked illustrative example and a calibration for the Gulf buyer. Section 6 discusses limitations. Section 7 concludes.

How synthetic exposure replaces ownership of the share




Figure 1. The Anatomy of a Synthetic Pre-IPO Position. *Illustrative schematic. The buyer holds a contract that references the underlying share; the share, the collateral and the settlement mechanism sit with or behind the counterparty. The buyer's economic exposure and its credit exposure travel along different paths and separate when the counterparty fails.*

2. LITERATURE REVIEW AND PROPOSITIONS

The framework draws on five bodies of work: the theory of forward and futures pricing; the literature on credit and counterparty risk; the analysis of collateral, netting and close-out; the economics of information asymmetry and contract design; and the specific literature on private-market and pre-IPO secondary transactions. Each contributes one element of the method developed in Sections 3 and 4.

2.1 Forward Pricing and Synthetic Replication

The starting point is the theory of forward contracts and replication. A forward is an agreement to exchange an asset, or its cash equivalent, at a fixed price on a future date. Classical no-arbitrage pricing, developed for traded assets, holds that the forward price equals the spot price compounded at the financing rate net of any income the asset yields. The deeper insight for this paper is the principle of synthetic replication: a payoff can be manufactured out of a contract that need never involve holding the underlying asset at all. Hull's standard treatment of derivatives makes the point that a long forward and a holding of the underlying financed by borrowing have the same payoff, which is why the two are priced off each other. The corollary, central here, is that the holder of a synthetic has the same economic payoff but a different balance sheet: the synthetic holder has a claim on a counterparty where the direct holder has an asset.

This literature also supplies the language of basis risk, the risk that the instrument used to obtain exposure does not move exactly with the thing the buyer cares about. In a pre-IPO synthetic the reference may be a specific class of share, a blended valuation or a vehicle whose own costs and frictions drive a wedge between its mark and the underlying. Basis risk is usually small relative to counterparty risk in this market, but it is real and it compounds over a long holding period.

2.2 Credit and Counterparty Risk

A second literature concerns credit risk and, specifically, counterparty credit risk, the risk that the party on the other side of a bilateral contract fails to perform. The financial crisis of 2008 turned this from a specialist concern into a central one, and the subsequent literature formalised the measurement of counterparty exposure, the credit valuation adjustment that prices it, and the wrong-way risk that arises when exposure to a counterparty grows precisely as that counterparty's own creditworthiness deteriorates. Wrong-way risk is acute in pre-IPO synthetics: a writer who has hedged poorly or over-committed is most likely to fail in exactly the buoyant market in which the buyer's position is most valuable, so the buyer's claim is largest at the moment it is least likely to be honoured.

The literature distinguishes default risk, the probability that the counterparty fails, from loss given default, the fraction of exposure that is not recovered when it does. The two are separate levers. A buyer facing a strong counterparty has a low probability of default and can tolerate weak recovery; a buyer facing a weak counterparty must insist on collateral and netting to raise recovery, because the probability of default cannot be wished away. This decomposition structures the mitigation analysis in Section 4.

2.3 Collateral, Netting and Close-Out

A third literature examines the mechanisms that convert an unsecured promise into a secured claim. Collateral, an asset posted by the counterparty that the buyer may seize on default, raises recovery directly. Netting, the right to set off mutual obligations so that only a single net amount changes hands, reduces the gross exposure that must be collateralised. Close-out, the right to terminate and value all outstanding obligations at once on a defined default event, prevents a failing counterparty from cherry-picking which contracts to honour. The architecture of the derivatives market, built on master agreements that combine these three mechanisms, exists precisely because bilateral credit exposure is otherwise unmanageable at scale. The pre-IPO synthetic market has no equivalent standard, which is why the buyer must assemble these protections contract by contract.

This literature also warns that the protections are only as good as their enforceability. A collateral right that cannot be perfected, or a close-out provision that a court in the counterparty's home jurisdiction will not honour, is decoration. Legal venue is therefore not a detail but a determinant of recovery, a theme developed in Section 4.4.

2.4 Information Asymmetry and Contract Design

A fourth literature concerns the design of contracts where one party knows more than the other. Akerlof's account of markets with hidden quality shows how asymmetry can drive good offers out and leave the buyer facing adverse selection. In the synthetic pre-IPO market the asymmetry is severe: the writer knows the true structure, the hedge, the chain of title behind the reference and its own financial condition, while the buyer sees a payoff diagram and a discount. The literature's prescription is signalling and screening, that the informed party demonstrate quality through verifiable commitments and the uninformed party demand them. A counterparty willing to post collateral, accept a strong governing law and submit to independent valuation is signalling; a counterparty that resists all three is screening itself out.

The asymmetry has a temporal dimension as well. At the moment of sale the writer's incentive to disclose is weakest, because disclosure can only reduce the price it receives, while after the sale the buyer's ability to compel disclosure depends entirely on rights it secured in the contract. The lesson the literature draws is that the buyer's leverage is highest before it commits and falls to almost nothing afterwards, so the

screening must be done up front. A buyer that signs first and asks questions later has surrendered the only moment at which the informed party had a reason to answer them.

2.5 Private-Market and Pre-IPO Secondary Transactions

A fifth literature addresses private-market secondaries directly. It documents the discounts at which late-stage private shares change hands, the friction created by issuer transfer restrictions and rights of first refusal, and the role of intermediaries and vehicles in bridging buyers to positions they could not access directly. This literature explains why synthetic structures exist at all: the binding constraint in much of the market is not price but access, and a synthetic is often the only route to a name whose shares cannot lawfully be transferred to the buyer. It also documents the cost of that route, in fees, in layered vehicles and in the counterparty risk that is this paper's central concern.

2.6 Synthesis and Propositions

Taken together, these literatures imply that a synthetic pre-IPO position should be analysed not as a cheaper or faster way to own a share but as a different instrument with a different risk that happens to share a payoff. They yield five propositions that the rest of the paper develops and illustrates.

- Proposition 1. Identify what you own before you price it. The first question is not the discount but the legal nature of the claim: a transferred share, a secured interest in a vehicle, or an unsecured promise. Everything else follows from the answer.
- Proposition 2. The economics equal a direct holding only when settlement performs. The payoff diagram is honest about the upside and silent about the condition attached to it, namely that the counterparty is there to pay at the end.
- Proposition 3. Counterparty risk is the price of the convenience and must be priced explicitly. The access, the speed and the aggregation a synthetic provides are paid for in credit exposure, and that exposure should be deducted from the value of the position, not ignored.
- Proposition 4. Collateral, netting and venue convert a promise into a recoverable claim. The difference between a synthetic worth holding and one worth declining is usually the strength of these three mechanisms, not the identity of the underlying company.
- Proposition 5. Some access is not worth having synthetically. When the counterparty is weak and the protections are thin, no discount compensates for the risk, and the disciplined answer is to take the share directly or not at all.

3. DATA AND METHODOLOGY

This study is analytical and descriptive rather than empirical. It builds a structured method for understanding and pricing synthetic pre-IPO exposure and illustrates that method with stylised figures and a worked example. It does not estimate parameters from a dataset, and the numbers used in the illustrations are chosen to make the structure clear, not to measure any real market.

3.1 The Analytical Framework

The framework has four components, corresponding to the four questions a buyer of a synthetic must answer. The first is structural: what instrument is this, and what does the buyer legally hold? The second is economic: what is the payoff, and how does it compare with a direct holding when all goes well? The third is credit: how large is the counterparty exposure, and how should it be priced? The fourth is mitigation: what collateral, netting and legal architecture turn the promise into a claim the buyer could actually enforce? The four are sequential. A buyer who cannot answer the first should not reach the fourth.

3.1a What the Framework Does and Does Not Do

It is important to be clear about scope. The framework is a tool for reasoning about a class of instrument, not a valuation engine and not a substitute for legal opinion. It does not tell a buyer the right price for a specific position, because that depends on facts the framework deliberately abstracts away: the identity and balance sheet of the counterparty, the exact terms of the contract and the law that governs it. What the framework does is impose an order on the analysis so that no element is skipped and the counterparty risk, which offer documents tend to leave implicit, is made explicit and priced. The choice of an analytical, illustrative method rather than an empirical one is deliberate. The synthetic pre-IPO market is private and heterogeneous; reliable cross-sectional data on counterparty failures and recoveries does not exist in a form that would support estimation, and a spuriously precise model would mislead more than a transparent framework.

3.2 Assumptions and Their Justification

The stylised inputs used in the figures, the financing rate embedded in the forward price, the size of the haircut applied to collateral, the recovery fractions in the counterparty waterfall and the gradings in the risk grid, are calibrations chosen to make the method legible. They are not forecasts, sourced data or the parameters of any actual instrument. Table 1 sets out the base-case assumptions and the reasoning behind each. A reader who disagrees with a calibration can substitute another; the structure of the analysis does not depend on the particular figures.

Parameter	Base-case value	Rationale
Reference last-round value	Indexed to 100	Anchors the payoff and basis figures to the most recent priced round.
Forward financing rate	5% per annum	Reflects the cost of carry the writer embeds in the forward price in 2026 USD conditions.
Collateral haircut	20%	Buffer against fall in collateral value between default and realisation.
Netting benefit	Up to 25% of gross	Set-off of mutual obligations under a single master agreement.
Base recovery, weak counterparty	30 to 40%	Unsecured claim against a thinly capitalised writer.
Base recovery, collateralised	75 to 90%	Secured claim with perfected collateral and enforceable close-out.
Holding period to exit	12 to 24 months	Typical interval from synthetic purchase to listing or trade sale.

Table 1. Base-Case Assumptions for the Synthetic-Exposure Framework. *Indicative, illustrative calibrations used to make the method legible; not forecasts, sourced data or the parameters of any real instrument. See Appendix A for the full set.*

4. THE FRAMEWORK: STRUCTURE, PAYOFF, COUNTERPARTY RISK AND MITIGATION

This section develops the framework in the order in which a buyer should work through a synthetic offer. It begins with the structure and what the buyer owns, turns to the payoff and its hidden condition, measures and prices the counterparty risk, sets out the mitigation that converts a promise into a claim, and closes by comparing the synthetic with the direct holding it imitates.

4.1 Structure: What the Buyer Actually Owns

Proposition 1 holds that the first question is the legal nature of the claim. Synthetic pre-IPO offers come in a small number of recurring shapes, and the buyer should be able to place any offer into one of them. In a forward, the buyer and a counterparty agree today that, at a defined future event, value will change hands based on a reference price; until then the buyer owns a contractual right and obligation, not a share. In a profit-participation or total-return structure, the buyer is promised the economic return on a reference share, positive or negative, in exchange for a payment, again without owning the share. In a vehicle structure, a special-purpose entity holds the share and the buyer owns an interest in the entity, which interposes the vehicle's own governance, costs and creditors between the buyer and the company. Each shape delivers a similar economic exposure and a different legal claim, and the difference determines what happens in distress.

The practical test is to ask, for any offer, three questions. Whose name is on the company's register against the reference share? What does the buyer hold that it could enforce, and against whom? And what stands between the buyer and the share if the counterparty fails? A buyer who can answer these has identified the instrument; a buyer who cannot has not understood the offer, whatever the discount.

The three shapes also differ in how they fail, which is the property that matters most. If a forward writer defaults, the buyer holds a contractual claim for damages and joins the writer's other unsecured creditors, recovering whatever the writer's estate can pay. If a total-return counterparty defaults, the position is similar, an unsecured claim for the return promised. But if a vehicle structure is used and the vehicle is bankruptcy-remote and actually holds the share, the buyer's interest in the vehicle may survive the failure of the party that sold it, because the asset sits behind a wall the seller's creditors cannot cross. This is the strongest form of synthetic exposure, and it is also the form most often imitated in name without being delivered in substance. A buyer told the structure is bankruptcy-remote should verify that the vehicle truly holds the share, that the security is perfected and that the remoteness has been opined on by counsel, because the label is worth nothing without the substance behind it.

4.2 Payoff: The Economics and Their Hidden Condition

Proposition 2 holds that the synthetic equals a direct holding only when settlement performs. The payoff of a long forward or a total-return position is linear and symmetric: the buyer gains one for one as the reference rises above the agreed price and loses one for one as it falls below, with no premium paid and no optionality to cushion the downside. Figure 2 sets out this payoff against the writer's mirror-image position. The diagram is accurate as far as it goes, and it is what an offer document shows. What it does not show is the condition attached to the upper half of the line: the buyer realises the gain only if the counterparty is solvent and willing to pay at settlement. The payoff diagram is, in effect, drawn on the assumption that counterparty risk is zero, which is precisely the assumption the buyer must not make.




Figure 2. The Payoff of a Synthetic Long Position. Illustrative payoff indexed to the agreed forward price. The synthetic long is linear and symmetric with no premium and full downside, the mirror of the writer's short. The diagram assumes settlement performs; the counterparty risk that conditions the upside is shown separately in Figures 3 and 4.

The linearity matters for a second reason. Because there is no premium and no embedded option, the buyer cannot treat the position as a capped bet. The full notional is at risk to the underlying, and the full notional is also the measure of the credit exposure to the counterparty. A buyer sizing a synthetic position should think of the notional twice: once as market risk to the company and once as credit risk to the writer.

4.3 Counterparty Risk: Measuring and Pricing It

Proposition 3 holds that counterparty risk is the price of the convenience and must be priced explicitly. The buyer's loss if the counterparty fails is not the whole notional; it is the notional reduced by whatever can be recovered. Figure 3 sets out this build-up as a waterfall, from the gross notional the buyer is exposed to, through the collateral it can seize, the obligations it can net, and the recovery it can expect on the residual unsecured claim, to the net loss it would actually bear. The discipline the figure enforces is that the buyer should never quote its counterparty exposure as the headline notional, which overstates it, nor ignore it, which is worse, but should work down the waterfall to the net figure and price that.




Figure 3. From Gross Notional to Net Loss-Given-Default. *Illustrative waterfall indexed to gross notional of 100. Collateral, netting and expected recovery reduce the exposure step by step to the net loss the buyer would bear if the counterparty defaulted. The size of each step is a property of the contract, not of the underlying company.*

Pricing the risk means deducting its expected cost from the value of the position. The expected cost is the probability of counterparty default over the holding period multiplied by the net loss-given-default from the waterfall. A position offered at an apparently attractive discount to the reference value may, once this deduction is made, be no cheaper than a direct holding, or dearer. The deduction is the synthetic buyer's equivalent of the credit valuation adjustment that a derivatives desk computes as a matter of course, and the fact that no one computes it for the buyer is exactly why the buyer must compute it itself.

Counterparty risk is not uniform across the kinds of party that write these contracts. Figure 4 grades the principal risk dimensions across a spectrum of counterparties, from a regulated bank to an offshore platform. The grid is illustrative, but its shape carries the message of the section: risk rises as the chain between the buyer and the underlying lengthens and as the counterparty moves further from regulation, capital and enforceable law. A buyer facing the left of the grid can rely more on the counterparty's standing and less on contractual protection; a buyer facing the right must insist that every protection is in place and documented, because nothing else stands behind the promise.

Counterparty risk rises as the chain lengthens

Concentration	Low	Med	Med	High
Operational / custody	Low	Med	Med	High
Legal enforceability	Low	Med	High	High
Collateral quality	Low	Med	Med	High
Default risk	Med	Med	High	High
	Regulated bank	Specialist broker	SPV / feeder	Offshore platform

Figure 4. Counterparty Risk Across Types of Writer. *Illustrative grading of risk dimensions across counterparty types. Risk rises from left to right as the writer moves away from regulation, capital and enforceable law, and as the chain between the buyer and the underlying share lengthens. Gradings are schematic and not specific to any institution.*

A subtler risk, basis risk, is shown in Figure 5. Even when the counterparty performs, the synthetic mark may drift from the true value of the underlying over the holding period, because the reference is imperfect, the vehicle carries its own costs, or the valuation is updated on a different cadence than the company's own rounds. Over a short hold the drift is small; over the eighteen to twenty-four months that often separate a synthetic purchase from a listing, it can become material. Basis risk is second-order relative to counterparty risk, but a buyer who has priced the credit risk and ignored the basis has still not finished.




Figure 5. Basis Risk Over the Holding Period. *Illustrative paths for the reference value and the synthetic mark. Even with a performing counterparty, the synthetic can drift from the underlying as costs, imperfect references and valuation cadence drive a wedge between them. The gap widens with the length of the hold.*

4.4 Mitigation: Turning a Promise Into a Claim

Proposition 4 holds that collateral, netting and venue convert a promise into a recoverable claim. This is where the buyer has agency. The counterparty's creditworthiness is largely given, but the recovery the buyer would achieve on default is not; it is built out of contractual protections the buyer can demand before committing. Three protections do most of the work.

The first is collateral. A counterparty that posts an asset the buyer can seize on default, held by an independent custodian and subject to a perfected security interest, transforms an unsecured claim into a secured one and lifts expected recovery sharply. The quality of the collateral and the enforceability of the security interest matter as much as its existence; collateral the buyer cannot reach, or cannot value, is not collateral. The second is netting and close-out, the right to terminate all obligations and settle a single net amount on a defined default event, which prevents a failing counterparty from honouring the contracts that suit it and defaulting on the rest. The third is legal venue, the choice of a governing law and a forum whose courts will enforce the collateral and the close-out as written. A strong contract under a weak law is a weak contract.

These protections are not free. A counterparty that posts collateral and accepts a strong governing law will price that into the terms, and the buyer should expect to pay, in a smaller discount or an explicit fee, for the recovery it is buying. That is the correct trade. The buyer is converting an uncertain, low recovery into a high, more certain one, and paying for the difference is paying for exactly the thing that distinguishes a synthetic worth holding from one worth declining.

A practical sequencing follows from this. The buyer should settle the protections before it settles the price, because the price is only interpretable once the recovery profile is known. An offer at a 15% discount with full collateral and an enforceable close-out is a different instrument from the same 15% discount unsecured, and treating them as the same number is the error the framework exists to prevent. In

negotiation, therefore, the buyer should establish what mitigation the counterparty will accept, derive the recovery profile that mitigation implies, price the residual counterparty cost, and only then judge whether the discount on offer is adequate. A counterparty that wants to agree the headline first and the protections later is reversing the order for a reason, and the buyer should resist it.

4.5 Synthetic Versus Direct: What Is Traded Away

It is worth setting the synthetic against the direct holding it imitates, because the comparison clarifies what the buyer gains and gives up. Figure 6 sets out the two side by side across the dimensions that separate them. The direct holding puts the buyer's name on the register, carries no counterparty risk beyond settlement of the original purchase, and typically conveys whatever information and consent rights attach to the share. It is, however, often unavailable, blocked by transfer restrictions, and it ties up the full capital from day one. The synthetic reaches names the direct route cannot, can be assembled quickly and in size, and may require less capital upfront, but it substitutes a credit claim for an asset, adds legal complexity, and usually strips out the information rights. Neither is strictly better. The synthetic is the right choice when access is the binding constraint and the counterparty is strong; the direct holding is the right choice when it is available and the counterparty is not worth the risk.




Figure 6. Direct Holding Versus Synthetic Exposure. *Illustrative comparison across the dimensions that separate the two routes. The synthetic buys access and speed at the cost of counterparty exposure, legal complexity and lost information rights. The right choice depends on which constraint binds.*

5. BRINGING THE ANALYSIS TOGETHER

5.1 The Decision Matrix

The four components combine into a single decision. Two axes capture most of what matters: the strength of the counterparty, which governs how much the buyer must rely on contractual protection, and the quality of the access the synthetic provides, which governs how much the buyer gains by taking the synthetic route at all. Figure 7 arranges the resulting choices into four quadrants. Where access is scarce

and the counterparty is strong, the synthetic is viable and the discount can be justified. Where access is ample and the counterparty is strong, the buyer should negotiate the convenience premium down or take the share directly. Where the counterparty is weak, the matrix points away from the synthetic regardless of access, towards a direct purchase if one can be had and towards declining if it cannot.

The buyer's decision matrix for synthetic exposure




Figure 7. The Buyer's Decision Matrix for Synthetic Exposure. *Illustrative two-axis matrix. Counterparty strength governs reliance on protection; quality of access governs the gain from going synthetic. The quadrants map to viable, negotiate, direct-only and decline. The matrix orders the judgement; it does not replace it.*

5.2 A Worked Illustrative Example

Consider a stylised case. A buyer is offered synthetic exposure to a sought-after late-stage technology company through a forward written by a specialist broker, at a price that implies a 15% discount to the company's last priced round. The direct share is unavailable because the issuer's transfer restrictions block a sale to a new holder, so the synthetic is the only route. Working through the framework, the buyer first identifies the instrument: it holds a contractual right against the broker, not a share, and the broker holds or hedges the underlying. The payoff is the linear long shown in Figure 2, conditioned on the broker performing at settlement.

The buyer then prices the counterparty risk. The broker is unregulated and thinly capitalised, so the probability of default over the eighteen-month expected hold is not negligible, and the unsecured recovery would be low, of the order of the 30 to 40% in Table 1. Applying the waterfall in Figure 3 to the raw offer, the expected counterparty cost consumes a large part of the headline 15% discount, leaving the position only marginally cheaper than a direct holding would be, and dearer once basis risk is allowed for. The buyer's response is not to walk away but to demand mitigation: collateral posted with an independent custodian, a perfected security interest, a close-out right and a governing law whose courts will enforce both. With those in place, expected recovery rises towards the 75 to 90% range, the expected counterparty cost falls, and the 15% discount once again represents real value. The example shows the framework

working as intended: the same offer is a poor deal unmitigated and a sound one mitigated, and the difference is the contractual architecture, not the underlying company.

5.3 Calibrating the Framework for the Gulf Buyer in 2026

The framework is general, but its application to a Gulf buyer in 2026 has specific features. Appetite among GCC family offices and intermediaries for US late-stage technology exposure is strong, and the most sought-after names are precisely those whose transfer restrictions make direct purchase hardest, so the synthetic route is more often the only route here than in markets with deeper direct access. That makes the discipline of this paper more important, not less: the buyer who must go synthetic to get the exposure it wants has the most need to price the counterparty risk that comes with it.

Two further features matter. First, the counterparties writing these contracts for Gulf buyers range widely, from regulated institutions in established financial centres to lightly regulated platforms offshore, so the grid in Figure 4 spans its full width in practice and the choice of counterparty is a live decision rather than a formality. Second, the question of legal venue is sharper for a cross-border buyer, because the collateral may sit in one jurisdiction, the counterparty in another and the buyer in a third, and the enforceability of the close-out depends on how those jurisdictions interact. A Gulf buyer should treat the governing-law clause and the location of the collateral as first-order terms, to be settled with qualified counsel before the economics are agreed, not afterwards.

There is also a reputational dimension specific to the region. The Gulf pre-IPO market is relationship-driven and the community of serious buyers and intermediaries is small, so a counterparty's conduct on one transaction is quickly known across the market. This works in the buyer's favour: a writer that expects to remain active has a strong incentive to honour its contracts and to be seen to do so, and that reputational collateral, while it cannot be seized, genuinely lowers default risk for established counterparties. It is not, however, a substitute for the contractual protections, because reputation protects against an unwillingness to pay and not against an inability to pay, and it is precisely the second that destroys a synthetic position. The Gulf buyer should weigh reputation as a real but partial comfort and insist on collateral and venue regardless, especially where the counterparty is newer to the market or its balance sheet is opaque.

6. LIMITATIONS AND DISCUSSION

Several limitations qualify the analysis. First, the framework is analytical and its figures are illustrative; it orders a judgement rather than computing a price, and a buyer facing a real instrument needs the specific terms, the counterparty's actual financial condition and qualified legal advice that the framework deliberately abstracts away. Second, the counterparty-risk parameters, default probabilities and recovery fractions, are not estimated from data, because reliable cross-sectional data on synthetic pre-IPO failures does not exist in a usable form; the figures are reasoned calibrations and a reader may substitute others. Third, the legal analysis is generic. The enforceability of collateral and close-out is a jurisdiction-specific question, and the paper can flag its importance but cannot answer it. Fourth, the framework treats the underlying company's value as given and concentrates on the wrapper; a buyer must still do the separate work of diligencing the company and its price, which is the subject of a companion analysis. Finally, the market is evolving, and structures, counterparties and the regulatory treatment of synthetic private-market exposure may change in ways that alter the calibration even as the underlying logic holds.

Within these limits the paper's contribution is a way of thinking. The central claim, that a synthetic is an economic exposure and a credit claim bundled together and that the buyer's job is to unbundle and price both, is robust to the particular figures. A buyer who internalises that claim, who refuses to read a payoff

diagram as if counterparty risk were zero, and who demands the collateral, netting and venue that convert a promise into a claim, will make better decisions than one who treats a synthetic as simply a cheaper share, whatever calibration is used.

7. CONCLUSION

Forward contracts and synthetic instruments have become a significant route through which buyers obtain exposure to late-stage private companies, particularly where transfer restrictions make a direct purchase impossible and where intermediaries aggregate demand into a single position. They are genuinely useful. They reach names that the direct market cannot, they can be assembled quickly and in size, and they let a buyer express a view it could not otherwise express. But they are not shares, and the difference is the whole point. The buyer of a synthetic holds an economic exposure to a company wrapped around a credit claim on a counterparty, and the two separate precisely in distress, when the exposure is most valuable and the counterparty least able to honour it.

This paper has argued that the buyer's task is to unbundle the two and to price the credit claim that offer documents leave implicit. It has set out the anatomy of the principal structures, shown that the advertised payoff is conditional on settlement, provided a waterfall and a grid for measuring counterparty risk, and identified the collateral, netting and legal venue that convert an unsecured promise into a recoverable claim. The decision matrix and the worked example show the framework discriminating between offers that look alike on their economics and differ entirely on their risk. For the Gulf buyer and intermediary in 2026, for whom the synthetic route is often the only route to the exposure they want, the discipline is not optional. The convenience of a synthetic is real, and so is its price; the buyer who pays the price knowingly, and secures the claim properly, can use these instruments well, while the buyer who reads only the discount is lending to a stranger and calling it an investment.

There is a final point that bears on how an intermediary should use this framework. A broker or platform that structures these positions has an interest in the buyer's confidence, but it has a stronger long-term interest in the buyer's survival as a repeat counterparty, and the two align around honesty about counterparty risk. An intermediary that explains the credit claim, that helps the buyer secure collateral and venue, and that declines to place a position it could not defend, builds the kind of relationship that outlasts any single trade. The discipline this paper sets out is therefore not only a buyer's protection but an intermediary's standard, and the two readers the paper addresses are, in the end, served by the same rule: price the promise, secure the claim, and never let the discount stand in for either.

REFERENCES

1. Akerlof, G. A. (1970). The Market for Lemons: Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics*, 84(3), 488-500.
2. Amihud, Y. and Mendelson, H. (1986). Asset Pricing and the Bid-Ask Spread. *Journal of Financial Economics*, 17(2), 223-249.
3. Bharath, S. T. and Shumway, T. (2008). Forecasting Default with the Merton Distance to Default Model. *Review of Financial Studies*, 21(3), 1339-1369.
4. Black, F. and Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81(3), 637-654.
5. Brigo, D., Morini, M. and Pallavicini, A. (2013). *Counterparty Credit Risk, Collateral and Funding*. Chichester: Wiley.
6. Cochrane, J. H. (2005). The Risk and Return of Venture Capital. *Journal of Financial Economics*, 75(1), 3-52.
7. Cox, J. C., Ingersoll, J. E. and Ross, S. A. (1981). The Relation between Forward Prices and Futures Prices. *Journal of Financial Economics*, 9(4), 321-346.

8. Duffie, D. and Singleton, K. J. (2003). Credit Risk: Pricing, Measurement and Management. Princeton: Princeton University Press.
9. Gregory, J. (2020). The xVA Challenge: Counterparty Risk, Funding, Collateral, Capital and Initial Margin. 4th ed. Chichester: Wiley.
10. Hull, J. C. (2018). Options, Futures and Other Derivatives. 10th ed. Harlow: Pearson.
11. International Swaps and Derivatives Association (2010). ISDA Master Agreement and Credit Support Documentation: User's Guide. New York: ISDA.
12. Jensen, M. C. and Meckling, W. H. (1976). Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure. Journal of Financial Economics, 3(4), 305-360.
13. Longstaff, F. A. (1995). How Much Can Marketability Affect Security Values? Journal of Finance, 50(5), 1767-1774.
14. Merton, R. C. (1974). On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. Journal of Finance, 29(2), 449-470.
15. Metrick, A. and Yasuda, A. (2010). Venture Capital and the Finance of Innovation. 2nd ed. Hoboken: Wiley.
16. Pulvino, T. C. (1998). Do Asset Fire Sales Exist? An Empirical Investigation of Commercial Aircraft Transactions. Journal of Finance, 53(3), 939-978.
17. Spence, M. (1973). Job Market Signaling. Quarterly Journal of Economics, 87(3), 355-374.
18. Stulz, R. M. (2010). Credit Default Swaps and the Credit Crisis. Journal of Economic Perspectives, 24(1), 73-92.

APPENDIX A. BASE-CASE ASSUMPTIONS

The following table records the full set of stylised inputs used in the figures and the worked example. The values are illustrative calibrations chosen to make the framework legible and are not forecasts, sourced data or the parameters of any real instrument.

Parameter	Value	Used in
Reference last-round value	100 (indexed)	Figures 2, 5
Forward financing rate	5.0% p.a.	Forward price, Figure 2
Expected holding period	12 to 24 months	Figures 5, worked example
Collateral haircut	20%	Figure 3, mitigation
Netting benefit	Up to 25% of gross	Figure 3
Recovery, unsecured weak writer	30 to 40%	Figure 3, worked example
Recovery, collateralised claim	75 to 90%	Figure 3, worked example
Illustrative headline discount	15% to last round	Worked example
Basis drift, annualised	1 to 4% of value	Figure 5

Table A1. Full Base-Case Assumptions. *Indicative, illustrative calibrations. Substituting other values changes the figures but not the logic of the framework.*

APPENDIX B. SYNTHETIC-EXPOSURE DILIGENCE CHECKLIST

The checklist condenses the framework into questions a buyer should be able to answer before committing to a synthetic position. It is illustrative and not exhaustive.

- Structure. What instrument is this, and what does the buyer legally hold: a transferred share, a secured vehicle interest, or an unsecured contractual promise?
- Register. Whose name stands against the reference share on the issuer's register, and what stands between the buyer and that share if the counterparty fails?

- **Payoff.** Is the payoff linear and unconditional in the offer's presentation, and on what condition, namely counterparty performance, does the advertised upside actually depend?
- **Counterparty.** Who is the writer, how is it regulated and capitalised, and where does it sit on the risk grid in Figure 4?
- **Exposure.** What is the gross notional, and what is the net loss-given-default after collateral, netting and expected recovery from the waterfall in Figure 3?
- **Pricing.** Has the expected counterparty cost been deducted from the headline discount, and does the position still represent value once it has?
- **Collateral.** Is collateral posted, of what quality, with an independent custodian, and is the security interest perfected and enforceable?
- **Netting and close-out.** Is there an enforceable right to terminate and net all obligations on a defined default event?
- **Venue.** What law governs the contract, where would enforcement take place, and have qualified counsel confirmed the collateral and close-out are enforceable there?
- **Basis.** Over the expected hold, how far might the synthetic mark drift from the underlying, and is that drift acceptable?

APPENDIX C. COUNTERPARTY RED-FLAG LIST

The following are illustrative warning signs that a synthetic offer carries more counterparty risk than its presentation suggests. Any one warrants closer scrutiny; several together warrant declining.

- The counterparty resists posting collateral or proposes collateral the buyer cannot independently value or reach.
- The offer presents the payoff diagram but does not name the writer or describe its financial condition.
- The governing law is an unfamiliar offshore jurisdiction chosen by the counterparty and resistant to independent enforcement.
- There is no close-out or netting provision, so the buyer could be cherry-picked in the counterparty's default.
- The chain between the offer and the underlying share runs through several undisclosed vehicles or intermediaries.
- The discount is unusually large relative to comparable direct positions, which may compensate for risk the buyer has not yet identified.
- The counterparty is unwilling to submit the position to independent valuation during the hold.
- Pressure to commit quickly, before the structure and the legal terms can be verified by counsel.

APPENDIX D. GLOSSARY OF KEY TERMS

Basis risk. The risk that the instrument used to obtain exposure does not move exactly with the underlying it references.

Close-out. The contractual right to terminate and value all outstanding obligations at once on a defined default event.

Collateral. An asset posted by a counterparty that the buyer may seize on default, raising expected recovery.

Counterparty risk. The risk that the party on the other side of a contract fails to perform its obligations.

Credit valuation adjustment. The deduction from a position's value that prices the expected cost of counterparty default.

Forward contract. An agreement to exchange an asset or its cash equivalent at a fixed price on a future date.

Loss given default. The fraction of an exposure that is not recovered when a counterparty defaults.

Netting. The right to set off mutual obligations so that only a single net amount changes hands.

Notional. The reference amount on which a synthetic position's payoff and exposure are calculated.

Synthetic exposure. Economic exposure to an asset obtained through a contract rather than by owning the asset itself.

Total-return arrangement. A contract conveying the economic return on a reference asset without conveying title to it.

Wrong-way risk. The risk that exposure to a counterparty grows as that counterparty's creditworthiness deteriorates.