

GCC Arabic Language Models and the Value of Data Rights Adoption and Sovereign Options

A Decision Framework for Proprietary Corpora Workflow Evidence Deployment Control and Capital

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Arabic language models can create strategic and commercial value in the Gulf Cooperation Council when they combine lawful access to differentiated corpora with accepted performance in real workflows. Benchmark leadership, parameter counts and national association provide useful technical signals. They do not establish customer adoption, transferable rights, durable margins or enterprise value. The investment decision requires evidence linking data, dialect and domain capability to contracts, collections, deployment control and funded model renewal.

This paper develops a decision framework for investors, boards and public-sector sponsors assessing GCC Arabic-language model companies. It separates Modern Standard Arabic, national and regional dialects, code switching, transliteration, speech, optical character recognition and domain terminology. It then connects these language requirements to six evidence ledgers: corpus rights, evaluation, workflow adoption, deployment and security, compute economics, and cash. Sovereign option value is defined through specific rights and operating capabilities rather than geography alone.

Four hypothetical cases illustrate the framework: a broad Arabic application programming interface challenger, a regulated-workflow specialist, a private sovereign deployment platform and a sector-corpus network. The assumptions produce annual revenue from AED 520 million to AED 1,200 million, contribution from AED 155 million to AED 610 million and cash before financing from negative AED 180 million to positive AED 55 million. A separate probability-weighted illustration produces enterprise value of AED 5,795 million. These figures demonstrate the analytical method. They are not market observations, forecasts or valuation conclusions.

The analysis finds that defensible value depends on a documented chain of rights, task-specific evaluation, conversion from pilot to paid production, reproducible deployment and sufficient capital to maintain capability. Open-weight access can lower entry barriers while creating opportunities for controlled private deployment and adaptation. Transaction structures should release capital and consideration against accepted workflows, signed minimum commitments, verified rights, operational portability and collected cash.

JEL Classification: G12, G24, G32, G34, L11, L86, O31, O33

Keywords: Arabic language models, GCC artificial intelligence, data rights, sovereign AI, workflow adoption, foundation model valuation, dialect evaluation, private deployment

Introduction

Arabic is used across markets that differ in dialect, administrative terminology, legal systems, sector vocabulary and digital behaviour. A foundation model that performs well on Modern Standard Arabic may still fail in a Gulf contact-centre exchange, a bilingual banking instruction, a handwritten government

form or an engineering maintenance record. Commercial diligence therefore needs a workflow definition before it needs a model ranking.

The regional model landscape has developed quickly. The Technology Innovation Institute describes Falcon Arabic as a seven-billion-parameter model trained on 600 billion tokens and designed for Modern Standard Arabic and major dialects [7-9]. Core42 and its partners released Jais models with Arabic-English capabilities, while MBZUAI and Inception announced Jais 2 as a seventy-billion-parameter open-weight model built around a larger Arabic-first dataset [10-15]. HUMAIN describes ALLAM 34B as an Arabic-first model hosted and operated in Saudi Arabia [16-18]. Provider claims require independent and customer-specific testing before they support an investment case.

The decision for an investor or board is whether a company controls a repeatable system for turning Arabic language capability into accepted work and cash. That system includes rights, evaluation, product integration, customer access, deployment control, support and model renewal. Each component has a cost and evidence standard. This paper builds a valuation and transaction framework around those observable elements.

1 Define the investment decision and perimeter

The first document in the diligence room should state the decision. A growth investor may be deciding whether to fund a model owner, an application company or a sovereign deployment platform. A strategic buyer may seek data, talent, customer access or operating control. A public sponsor may want national capability, Arabic inclusion, resilience or productivity. These objectives create different value tests.

The perimeter should identify legal entities, shareholders, repositories, model weights, tokenisers, data pipelines, evaluation systems, deployment software, contracts, compute agreements and customer obligations. It should show which assets are owned, licensed, open-weight, partner-controlled or customer-specific. Shared branding and national association do not resolve legal ownership or transferability.

The operating model also matters. A company may train a general model, adapt third-party weights, orchestrate several models, sell an application or operate a private instance. Each position creates different capital needs and defensibility. The valuation should allocate revenue, cost and risk to the actual layer rather than to an undifferentiated Arabic AI narrative.

2 Segment Arabic capability before measuring performance

Arabic capability should be segmented into Modern Standard Arabic, country and regional dialects, code switching, Arabizi and other transliteration, speech, optical character recognition, handwriting, retrieval and domain terminology. Customer acceptance can depend on one narrow combination. A bank may require Gulf Arabic speech recognition with English product names. A court or ministry may require formal Arabic, citations and controlled document retrieval.

Evaluation should preserve the user population, source format and consequence of error. General question answering, cultural knowledge, mathematical reasoning and safety tests measure different capabilities. A single composite score can hide a material failure. The diligence team should inspect disaggregated results and the distribution of hard cases.

Provider benchmarks help form a technical hypothesis. The commercial test uses the customer's data, workflow, latency, security and review rules. Results should be reproduced on the version proposed for deployment. The company should retain prompts, retrieval sources, tools, human-review policy and acceptance decisions so that later model changes can be compared on the same basis.

3 Build the corpus rights ledger

Training and adaptation data should be recorded by source, owner, licence, purpose, jurisdiction, collection method, permitted use, retention, deletion and model version. Public availability is not a substitute for permission. Contracts should cover model development, commercial use, derivatives, sublicensing, transfer in a transaction and audit where relevant.

Arabic corpora can contain books, news, public records, social media, transcribed speech, call-centre recordings, educational material, technical manuals and customer documents. Each class carries different rights and privacy questions. Dialect and speech datasets may contain personal data or sensitive context. Customer-provided data may support retrieval or fine-tuning without creating a transferable company asset.

The ledger should distinguish raw data, cleaned datasets, annotations, synthetic data, feedback and derived evaluation sets. It should identify duplicate or contaminated benchmark material. A buyer should be able to trace a released model to the relevant dataset versions and approvals. Where provenance is incomplete, the model should include the cost and time needed to remove, replace, retrain and re-evaluate the affected capability.

4 Value proprietary corpora through replacement and revenue evidence

A dataset has strategic value when lawful access is scarce, the data improves a required capability and the company can use that improvement commercially. Volume alone is a weak measure. A smaller, current and carefully annotated corpus can be more useful than a larger body of duplicated or low-quality text. Replacement cost should include sourcing, rights, cleaning, annotation, quality assurance, storage, evaluation and elapsed time.

Income evidence requires a link from corpus to workflow. The company should show that a defined dataset improved acceptance, reduced review effort, increased conversion, supported a higher price or protected retention. The analysis should control for model, prompt, retrieval and product changes. Customer-specific improvements may be valuable to the contract while remaining unavailable for the broader platform.

The valuation should avoid counting the same advantage twice. Corpus benefits can appear in forecast revenue, contribution or option value. An additional asset premium is appropriate only when the method excludes those cash flows or recognises a distinct transferable right. IAS 38 and IFRS 13 provide relevant accounting and fair-value principles, while transaction analysis still requires company-specific legal and financial advice [1-5].

5 Design a dialect and domain evaluation system

The evaluation system should start with use-case acceptance criteria. For a regulated document workflow, the measures may include extraction accuracy, grounded answers, citation completeness, prohibited disclosure, latency, cost and human-review time. For speech, word error rate is useful, although task completion and escalation may be more important to the customer. For content generation, factuality, terminology, tone and approval effort can determine value.

Test populations should represent countries, age groups, accents, code switching, script variation and document quality where these factors affect use. Native reviewers need clear rubrics and adjudication procedures. Inter-reviewer agreement should be measured. Customer confidential examples require access controls and a permitted-use record.

The company should maintain versioned scorecards and failure taxonomies. Improvement on an average benchmark should not conceal regression in a high-value cohort. Release decisions should include

capability, safety, cost and contractual commitments. Evaluation becomes an operating asset when it shortens customer diligence and guides model routing, adaptation and human review.

6 Convert benchmarks into workflow evidence

The commercial evidence ladder begins with a reproducible benchmark and progresses through demonstration, controlled pilot, user acceptance, production, invoice, collection and renewal. Each stage has a different confidence level. Announced partnerships, memoranda and pilot pipelines should remain outside contracted revenue until the relevant obligations and consideration are enforceable.

The adoption ledger should record customer entity, sector, country, language variant, workflow, model version, data boundary, acceptance threshold, users, volume, price, direct cost, implementation effort and cash. Forecast conversion should follow observed cohorts. Delayed pilots and non-paying use should be visible rather than blended into an aggregate adoption claim.

Workflow evidence also needs a counterfactual. The company should compare quality, time, cost and risk with the prior process or an available competing model. A model can be technically superior and commercially unattractive if integration, review or private-deployment costs exceed the customer benefit. Accepted business outcomes provide stronger evidence than leaderboard position.

7 Measure customer adoption and retention

Usage should be measured at the level that customers buy. Relevant units can include accepted documents, resolved contacts, completed applications, analyst hours saved or transactions reviewed. Token volume may reconcile compute cost, while it rarely explains customer value on its own. Free usage and promotional credits should be separated from paid activity.

Retention analysis should use contracted cohorts and distinguish model, application and service revenue. A customer may renew because of data integration, support or switching cost even when the underlying model is substitutable. Conversely, an open-weight deployment may preserve customer control while reducing recurring licence revenue. The company should explain the source of retention and the cost of maintaining it.

Pricing evidence should include list price, discount, implementation, minimum commitment, capacity reservation, support, acceptance credits and collection. Revenue recognition and cash timing can differ materially. IFRS 15 provides a framework for identifying contracts and performance obligations; company-specific accounting conclusions require professional judgement [4].

8 Define sovereign option value precisely

Sovereign option value should be decomposed into rights and capabilities. Relevant elements can include local legal control, model and software rights, data location, key control, compute access, administrator access, release authority, security operations, continuity, portability and the ability to support customers during external disruption. A local server provides only one component.

The customer's requirement determines economic value. A ministry may require national hosting and continuity. A regulated institution may require audit, localisation, access restrictions and tested recovery. A commercial customer may prioritise Arabic quality and total cost. The company should map each requirement to evidence and to the revenue cohort for which it changes eligibility or willingness to pay.

Option value also depends on exercise cost. A company that can move a workload only after extensive re-engineering has a weaker option than one with tested portability. The model should include reserved capacity, duplicate tooling, specialist staff and migration time. Sovereign positioning supports value when the right is usable, funded and relevant to customer decisions.

9 Test private deployment and portability

Private deployment can take the form of customer cloud, sovereign cloud, dedicated hosted capacity, on-premise infrastructure or a hybrid system. The diligence team should reproduce a representative workload in the proposed environment. It should measure quality, latency, throughput, security, monitoring, recovery, update process and cost.

The portability test should identify model licences, container images, proprietary dependencies, tokenisers, retrieval services, moderation, observability, accelerators and network requirements. Control of weights does not establish control of the complete service. The customer and supplier should agree who operates each layer and who can restore it after an incident.

Financial analysis should separate licence, implementation, capacity, managed service and support. Private deployments can create larger contracts and longer retention. They can also require bespoke engineering, inventory, parallel model versions and working capital before acceptance. Minimum prices and maintenance reserves should reflect the actual configuration.

10 Reconstruct compute and inference economics

Compute analysis should reconcile training, adaptation, evaluation and inference. Evidence includes provider invoices, accelerator reservations, cloud credits, utilisation, storage, networking, colocation, energy, support and termination rights. A nominal accelerator allocation does not show available productive capacity or priority during demand peaks.

Inference cost should be measured per accepted workflow unit. The schedule should include input and output tokens, caching, retrieval, tools, guardrails, retries, human review, idle capacity and service-level credits. Model routing, quantisation and smaller task-specific models can change economics. Their effect should be demonstrated on customer workloads.

Training economics require a programme ledger. Each run should record objective, data, compute, elapsed time, evaluation, release decision, customer relevance and cash effect. Research programmes can be necessary for renewal and option value, while repeated spending without adoption should reduce confidence in the forecast. Capital planning should state the financing required for the next capability threshold.

11 Separate open weight access from owned advantage

Open-weight models can lower development cost, support private deployment and widen the ecosystem. They can also make baseline capability easier to reproduce. The diligence team should read the applicable licence and verify the exact version. Open access to weights does not imply unrestricted rights to every dataset, tool, trademark or hosted service.

A company using open weights can still build durable value through lawful domain data, evaluation, workflow integration, distribution, support and controlled deployment. Each advantage should be tested for transferability and customer relevance. Adaptation artefacts and customer fine-tunes require clear ownership and confidentiality terms.

The substitution test should ask how long a competent competitor would need to reproduce the accepted workflow with another model. It should include data access, evaluation, integration, security approval, customer migration and support. A short technical replacement can still involve a long commercial transition. The valuation should model both.

12 Reconcile privacy security and human accountability

Arabic model workflows can process personal, confidential and regulated information. The UAE Personal Data Protection Law establishes a federal framework for personal-data protection, while Saudi Arabia's Personal Data Protection Law and implementing materials govern processing and transfer in that jurisdiction [19-23]. Applicable duties depend on entities, data, purpose and location. Legal advice is required for a specific deployment.

Controls should cover data classification, lawful basis, minimisation, access, encryption, retention, deletion, transfer, incident response and supplier management. Training, retrieval, logging and feedback are separate processing activities and should not be conflated. Customer prompts should not become training data without an authorised basis and clear contract.

NIST's AI Risk Management Framework and Generative AI Profile provide voluntary structures for governing, mapping, measuring and managing risk [24-27]. Human accountability remains necessary for consequential decisions. The company should define review, escalation and override, measure reviewer performance and avoid presenting a language model as the accountable decision maker.

13 Assess procurement and distribution power

The route to market can include direct enterprise sales, government frameworks, sovereign cloud, telecommunications providers, systems integrators, cloud marketplaces and software partners. The revenue forecast should show which party controls the customer, billing, data, support and renewal. Partner logos and technical integrations do not establish distribution volume.

Concentration should be measured across leads, contracts, collections, capacity and service delivery. A company can have many customers while depending on one marketplace or government programme for access. The downside case should model changed commercial terms, reduced visibility, delayed awards or partner failure.

Direct customer evidence, reproducible deployments and alternative channels can support resilience. Distribution value should be attributed only where contracts, observed conversion and economics justify it. Revenue sharing, integration cost, account ownership and termination rights belong in the contribution model.

14 Value talent and the Arabic research operating system

Arabic foundation-model development depends on researchers, data engineers, linguists, evaluators, safety specialists, infrastructure engineers and product teams. Critical knowledge can sit with a small group. The company should map roles, employment and invention rights, access privileges, succession, retention and the time needed to replace each capability.

The research operating system should preserve experiment records, data versions, code, model artefacts, evaluation, release decisions and incident learning. A buyer should be able to reproduce a defined build and explain changes between releases. Documentation that exists only in personal accounts or informal messages weakens transferability.

Talent value should be reflected through the cost of retaining and rebuilding capability, not through a separate premium that ignores future compensation. The forecast should include competitive pay, visas where relevant, compute support and management capacity. A strong system reduces key-person exposure and converts individual expertise into organisational capability.

15 Construct four GCC Arabic model cases

The analytical cases represent different operating strategies. The broad API challenger competes on Arabic quality, developer access and price. The regulated-workflow specialist focuses on a smaller number of high-consequence use cases with stronger evaluation and documentation. The private sovereign platform sells controlled deployment and continuity. The sector-corpus network builds domain-specific datasets and workflows with participating customers.

Each case should be modelled by customer cohort. Assumptions include conversion, accepted volume, price, compute, implementation, support, research, data, compliance, working capital and financing. Correlations matter. Slower adoption can reduce revenue while increasing the period over which the company funds research and capacity.

The cases are decision tools rather than forecasts. Management should replace every assumption with company evidence and assign an owner. The board should define triggers that move the plan between cases, such as accepted deployment, signed minimum commitments, corpus-rights exceptions, capacity agreements or delayed collections.

16 Build the integrated financial model

The model should connect technical and commercial ledgers to revenue, contribution, operating expense, capital expenditure, working capital and financing. Customer revenue should reconcile to contract, acceptance, invoice and cash. Compute should reconcile to usage and provider terms. Research and corpus spending should reconcile to the programme plan.

Cash timing can determine feasibility. A public or regulated deployment may require evaluation, security, integration and dedicated capacity before customer acceptance. The model should show milestone billing, receivables, supplier terms and liquidity by programme. A profitable contract can create a financing gap when cost precedes collection.

The base case should use observed cohort performance. Upside should remain conditional on evidence and funding. Downside should combine weaker conversion, lower price, higher inference and support cost, rights remediation, capacity delay and a slower financing round. Management responses should be specific and costed.

17 Apply valuation methods to the evidence states

Discounted cash flow can value established cohorts when revenue, cost and reinvestment are supportable. Market approaches can provide context, although differences in model ownership, data rights, deployment, growth, margin and capital intensity limit simple revenue multiples. A replacement-cost analysis can help assess corpora and engineering assets without proving income value.

The option method is useful when value depends on future technical and commercial gates. States can include accepted pilots, repeatable paid deployments, scaled cohorts and strategic distribution. Each state needs a probability, timing, capital requirement and enterprise value based on its economics. Probabilities should change when observable events occur.

The methods should reconcile. The valuation date, currency, debt, leases, commitments, grants, credits and non-operating assets should be explicit. Public support or compute access should be included only under its terms and available duration. IFRS 13 and International Valuation Standards provide relevant principles for fair-value and valuation work [1-3].

18 Illustrative Arabic model economics

The hypothetical cases show how strategy changes the relationship between revenue and cash. The API challenger has the largest assumed revenue and contribution, supported by scale and a standard service. The sector-corpus network has lower revenue and a smaller absolute research programme, while customer-specific delivery and data work constrain contribution.

The regulated specialist and private platform can support stronger contract value but carry documentation, implementation and capacity costs. Negative cash before financing does not itself show weak enterprise value. It shows that the capability and adoption plan require funding. The investor must decide whether future evidence is reachable with the available capital.

The numbers in the figures and tables are management assumptions created for illustration. They are not observations about a named company or market forecast. A real valuation must replace them with contracts, workload telemetry, invoices, payroll, compute agreements, rights evidence and a funded plan.

19 Value data rights adoption and sovereign options separately

The framework separates three value sources. Data rights support reproducible capability and remediation control. Adoption evidence supports revenue quality and customer utility. Sovereign options support access, continuity and strategic flexibility for defined buyers. The same fact should not be credited in more than one category without a distinct economic effect.

A rights score should measure chain of title, permitted use, transferability, traceability and remediation. An adoption score should measure acceptance, production, contribution, collection and renewal. A sovereign-options score should measure control, portability, capacity, security and customer eligibility. Weightings depend on the investment thesis.

The scores are governance aids rather than valuation formulas. They help identify missing evidence and conditions. Cash-flow and option models translate verified effects into value. The committee should document which evidence changes revenue, cost, risk, probability or capital.

20 Structure investment and consideration around evidence

Investment tranches can align funding with corpus clearance, reproducible evaluation, accepted deployment, minimum customer commitments and capacity. Acquisition consideration can combine base value, escrow, holdback and contingent payments. Milestones should use objective data sources and defined measurement periods.

Representations should address ownership, licences, data provenance, privacy, model versions, security, customers, compute and employment rights. Indemnities allocate financial exposure while leaving operational remediation to the company. A rights exception that threatens a core capability may require a closing condition or price adjustment rather than a general warranty.

The capital plan should protect customers if a later tranche is not released. Boards should avoid structures that create an unfunded service obligation. Instruments, tax and accounting treatment require current professional advice in the relevant jurisdictions.

21 Govern value after investment

Post-investment governance should use the same evidence architecture as diligence. The board pack should report corpus-rights exceptions, language and domain evaluation, accepted workflows, contracted cohorts, contribution, collections, compute, incidents, runway and financing. Model releases and material customer configurations should have accountable approval.

The company should refresh the valuation case after material changes in model, data, deployment, customer use, regulation, capacity or funding. Historical scores can become stale after a new release or supplier change. A dated exception register should show owner, remediation and effect on customer commitments.

Governance should remain decision oriented. The board needs the variables that change capital allocation, customer obligations and enterprise value. Detailed ledgers provide auditability. The principal report should explain the economic consequence and the required action.

Conclusion

GCC Arabic-language model value rests on a connected operating system. Lawful and transferable data supports capability. Task-specific evaluation demonstrates fitness. Customer acceptance, invoice, collection and renewal demonstrate adoption. Tested deployment, security and portability support control. Compute and funding determine whether the system can continue.

Benchmarks, national association and parameter scale remain useful signals. Their economic meaning depends on the rights, workflows and cash they support. An investor should value verified customer cohorts and controlled options while funding technical progress through explicit evidence gates.

The proposed framework gives boards a practical way to separate corpus value, adoption value and sovereign option value. It also keeps hypothetical scenarios distinct from observed evidence. The resulting transaction case can be updated as rights are cleared, deployments are accepted and cash is collected.

Arabic model rights and provenance register

The register should contain source, owner, licence, jurisdiction, collection method, personal-data class, permitted training and adaptation, commercial-use rights, sublicensing, transfer, retention, deletion, dataset version, model version and approval. It should include raw, cleaned, annotated, synthetic and feedback data. Every exception should have an owner, customer effect and remediation plan.

The company should test the register by tracing a selected model release back to dataset evidence. A static policy is insufficient when source records cannot be reproduced. The test should also cover data added after the principal training run and customer-specific fine-tunes. Access and confidentiality should reflect contractual duties.

For material gaps, management should estimate replacement data, labelling, compute, evaluation, delay and customer consequences. The schedule should distinguish a removable source from a dataset central to Arabic or domain capability. That distinction affects warranties, reserves, price and closing conditions.

Dialect domain and workflow evaluation file

The file should define user population, country, dialect, register, script, domain, task, consequence, model, prompt, retrieval, tools, human review, acceptance threshold and test date. It should retain examples, reviewer qualifications, agreement, failures and release decisions. Benchmark contamination and overlap with training data should be checked.

Customer tests should remain distinct from public benchmarks. Public suites help compare technical capability under published conditions. Customer acceptance measures performance inside the contracted workflow. Both are useful when their scope and limitations are stated.

The file should connect each accepted evaluation to product configuration, contract and support obligation. A later release should be compared with the accepted baseline. Material regression should trigger a release decision and customer communication under the applicable contract.

Adoption contract and collection ledger

The ledger should trace each customer from qualified opportunity through diligence, pilot, acceptance, production, invoice, collection and renewal. It should contain entity, sector, country, workflow, model, language variant, volume, price, discount, direct cost, implementation effort, partner share, receivable and accountable owner.

Forecast accuracy should be measured by cohort. Management should compare original probability, award date, production date, revenue, contribution and collection with actual results. Persistent delay should change pipeline conversion and working-capital assumptions.

Free, grant-funded, strategic and commercial workloads should be identified separately. A high-usage free service can create technical learning without proving pricing. A government programme can create procurement access while depending on budget and acceptance. The board should see the economic status of each cohort.

Sovereign deployment and continuity dossier

The dossier should state legal ownership, hosting, model and software rights, administrator access, cryptographic key control, compute priority, data location, subcontractors, support, incident response, recovery, portability and exit. Each assertion should link to a contract, test or controlled operating record.

A representative workload should be restored in the proposed environment and timed. The exercise should cover model artefacts, retrieval, moderation, monitoring, identity, secrets, network and data. Exceptions should show the customer and service consequence.

The dossier should distinguish control from location. Infrastructure in the GCC can remain dependent on a foreign supplier for capacity, software or administration. A distributed component can still be governed through enforceable rights and tested alternatives. The commercial claim should match the actual control position.

Compute capital and transaction file

The compute file should reconcile reservations, available accelerators, utilisation, invoices, credits, storage, networking, energy, support, location, service level, renewal and termination. Forecast capacity should match customer volumes and research programmes. Constrained and downside cases should include delayed procurement and lower utilisation.

The capital file should show programme funding, working capital, debt and equity by evidence gate. It should identify obligations that survive if a grant, allocation or financing round is delayed. The board should preserve sufficient liquidity for contracted customer service and orderly remediation.

The transaction room should connect title, data rights, model versions, evaluation, customers, compute, security, employees, financing and forecast vintages. Every committee claim should link to a dated source and accountable owner. This linkage allows the valuation case to change when the evidence changes.

The financial model should also preserve a bridge between technical events and cash. A new model release can increase research cost before it changes customer revenue. A rights exception can require retraining, evaluation and customer communication. A capacity reservation can improve continuity while creating a fixed commitment. Each event should have a dated assumption, an accountable owner and a defined line in the forecast. This bridge lets the board distinguish a technical milestone from a financially funded operating plan.

Management should maintain a forecast-vintage file. The file should record the customer, workload, volume, price, direct cost, implementation, collection and financing assumptions used at each decision date. Actual results should be compared with the original case and the latest revision. Persistent optimism


in conversion, delivery or collection should reduce the weight given to management estimates. Forecast discipline matters because a model company can consume substantial cash before the difference between technical interest and paid adoption becomes visible.

The transaction file should include a severe but plausible continuity case. The case can combine loss of a corpus licence, delay in a major customer, temporary capacity constraint, required security remediation and slower financing. The company should identify the services it would preserve, the research it would defer, the cash needed for customers and the decisions available to the board. This exercise tests whether sovereign positioning is supported by operational resilience and whether the proposed capital structure can carry the obligations created by customer contracts.

The committee should close the evidence loop after investment. Conditions, indemnities, holdbacks and milestones should be assigned to owners and reviewed against the same records used in diligence. If a customer acceptance gate releases consideration, the acceptance evidence should reconcile to the product version, contract and cash schedule. If a rights remediation closes an exception, the company should show the affected data, model and evaluation path. Governance that retains this traceability makes later financing, acquisition and regulatory review more reliable.


Finally, the evidence room should separate observations from scenarios. Contracts, invoices, collections, model artefacts, licences and test results are observations. Conversion rates, renewal, future price, research success and enterprise value remain assumptions until events occur. Each scenario should state its basis, date and sensitivity. This separation helps the investment committee update probabilities without rewriting the historical record or presenting a hypothetical outcome as a current market fact.

Figure 1 GCC Arabic language model value architecture




Proposed framework; conclusions require company-specific rights technical commercial and financial evidence.

Figure 2 Hypothetical annual revenue and contribution by GCC Arabic model case




Management assumptions in AED millions; figures are not market observations forecasts or valuation conclusions.

Figure 3 Hypothetical research data compliance capital and cash before financing



Management assumptions in AED millions; negative cash indicates illustrative financing need.


Figure 4 Hypothetical probability weighted enterprise value by evidence state



Management assumptions in AED millions; total probability weighted value is AED 5795 million.

Figure 5 First hundred days evidence and control sequence

Convert Arabic capability into accepted workflows controlled deployment and cash



Proposed sequence; timing should follow transaction customer regulatory and financing constraints.

Table 1 GCC Arabic model investment perimeter

Component	Required evidence	Valuation question	Principal risk
Models and software	repositories weights licences assignments release records	what capability is owned and transferable	open or restricted dependencies
Corpora and evaluation	provenance licences versions dialect and domain tests	can capability be reproduced and defended	missing rights or unsupported performance
Workflow adoption	acceptance contracts invoices collections renewals	which capabilities create paid demand	pilots mistaken for recurring economics
Deployment and compute	agreements utilisation hosting keys continuity exit tests	can service continue under required controls	capacity or supplier concentration

Component	Required evidence	Valuation question	Principal risk
Governance and security	privacy role map tests incidents approvals	can obligations be evidenced at scale	remediation interruption or liability

Proposed diligence structure; requirements depend on the company customer and transaction.

Table 2 Arabic capability segmentation

Segment	Example evidence	Commercial use	Diligence limitation
Modern Standard Arabic	controlled reasoning extraction and generation tests	government media education formal documents	average score can hide domain failure
Gulf and national dialects	native reviewer speech and dialogue evaluation	contact centres assistants field service	dialect labels can be too broad
Code switching and transliteration	mixed Arabic English and Arabizi task set	digital commerce support and social channels	tokenisation and spelling vary
Domain Arabic	customer corpus terminology and acceptance	banking legal health engineering	confidential tests may not transfer
Speech OCR and multimodal	audio document image and task completion evidence	public services operations archives	component accuracy differs from workflow utility

Each segment requires a task and customer-specific evidence standard.

Table 3 Adoption and customer economics ledger

Field	Required evidence	Control	Decision use
Validation and acceptance	protocol outputs review sign off	versioned model workload language	conversion and utility
Contract and billing	signed order price term invoice service level	entity and term reconciliation	revenue and renewal
Direct delivery cost	compute data implementation evaluation support partner	invoice and allocation reconciliation	contribution and cash
Collection	receivable settlement bank receipt	aged balance and exception owner	liquidity and customer quality
Retention	renewal usage accepted outcome and support	cohort and configuration continuity	durability of workflow value

Proposed minimum evidence fields for customer-level analysis.

Table 4 Hypothetical annual operating cases

Case	Revenue	Contribution	Research data and compliance	Cash before financing
Arabic API challenger	1200	610	430	55
Regulated workflow specialist	960	430	390	-70
Private sovereign platform	760	270	350	-135
Sector corpus network	520	155	290	-180

Management assumptions in AED millions; figures are not market observations forecasts or valuation conclusions.

Table 5 Hypothetical regulated workflow contribution bridge

Item	AED millions	Evidence required
Revenue	960	customer contract acceptance invoice and cash ledger
Compute and model delivery	-260	workload telemetry provider invoice capacity allocation
Corpus evaluation and compliance	-105	people supplier rights and programme records
Implementation partner and support	-165	project time partner settlement and service evidence
Contribution before central cost	430	reconciled customer cohort schedule

Management assumptions in AED millions; excludes central research sales administration financing and tax.

Table 6 Hypothetical valuation evidence states

Evidence state	Enterprise value	Probability	Weighted value
Rights mapped and accepted pilots	1600	25%	400
Repeatable paid deployments	4200	35%	1470
Scaled GCC workflow cohorts	8500	25%	2125
Strategic distribution and sovereign access	12000	15%	1800
Total		100%	5795

Management assumptions in AED millions; this is not a valuation conclusion.

Table 7 Investment and transaction evidence gates

Gate	Required evidence	Capital or consideration response	Post decision measure
Rights and control	assignments licences corpus provenance reproduced build	closing condition or protected tranche	exceptions and controlled release
Workflow adoption	accepted pilot signed order production invoice collection	base value after verification	conversion contribution and cash
Dialect and domain quality	controlled evaluation and customer acceptance	milestone funding	accepted outcome and review cost
Privacy security and continuity	role map tests incidents recovery and remediation	reserve holdback or staged release	exceptions and closure
Compute and portability	capacity terms utilisation cost and exit test	funding reserve or price adjustment	accepted task cost and continuity

Proposed framework; actual instruments require current legal tax accounting and financial advice.

Frequently asked questions

What creates value in an Arabic language model company

Value follows lawful and transferable data rights, accepted performance in customer workflows, repeatable deployment, durable contribution and sufficient capital to maintain capability. Benchmarks help assess technology and should be connected to commercial evidence.

How should proprietary Arabic corpora be valued

The analysis should test rights, scarcity, replacement cost, capability improvement and customer economics. A corpus should not receive a separate premium when its benefit is already included fully in forecast cash flow.

Do benchmark scores support a valuation premium

They can support a technical hypothesis. A premium requires reproducible results on relevant dialects domains and workflows, followed by customer acceptance, contribution and retention evidence.

What does sovereign AI mean for a GCC customer

It should be defined through specific control rights and operating capabilities across ownership, data, model, compute, keys, administration, security, continuity and portability. The required combination depends on the customer and use case.

How should open weight Arabic models be treated

Open weights can reduce development cost and support private deployment. The investor should verify the licence, complete service dependencies and the company's additional value in data, evaluation, workflow, distribution and support.

What is the principal downside case

The useful downside combines slower customer conversion, weaker dialect or domain acceptance, rights remediation, higher inference and support cost, capacity delay, longer collections and constrained financing. Company evidence determines the assumptions.

Which evidence should release investment capital

Useful gates include cleared rights, reproducible evaluation, accepted deployment, signed minimum commitments, contracted capacity, invoicing and collected cash. The transaction documents should define measurement and dispute procedures.

Which Matchpoint practice is relevant to this work

This research connects to Matchpoint Partners' strategy and execution, valuation, capital raising, transaction diligence and mergers and acquisitions work for technology companies, investors, public-sector sponsors and strategic partners.

References

- [1] IFRS Foundation. IFRS 13 Fair Value Measurement. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/>
- [2] IFRS Foundation. Educational material on measuring unquoted equity instruments within IFRS 13. <https://www.ifrs.org/content/dam/ifrs/supporting-implementation/ifrs-13/education-ifrs-13-eng.pdf>
- [3] International Valuation Standards Council. International Valuation Standards. <https://www.ivsc.org/standards/>

- [4] IFRS Foundation. IFRS 15 Revenue from Contracts with Customers. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-15-revenue-from-contracts-with-customers/>
- [5] IFRS Foundation. IAS 38 Intangible Assets. <https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/>
- [6] IFRS Foundation. IAS 36 Impairment of Assets. <https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/>
- [7] Technology Innovation Institute. Falcon Arabic. <https://falconllm.tii.ae/falcon-arabic.html>
- [8] Technology Innovation Institute. TII launches Falcon Arabic and Falcon H1. 21 May 2025. <https://www.tii.ae/news/middle-easts-leading-ai-powerhouse-tii-launches-two-new-ai-models-falcon-arabic-first-arabic>
- [9] Technology Innovation Institute. Falcon H1 Arabic launch. 5 January 2026. <https://falconllm.tii.ae/abudhabi-tii-launches-falcon-h1-arabic-establishing-the-worlds-leading-arabic-ai-model.html>
- [10] Core42. Core42 model organisation and Jais model index. <https://huggingface.co/core42>
- [11] Core42. Jais 30B v3 model card. <https://huggingface.co/inception42/jais-30b-v3>
- [12] Core42. Jais family white paper. <https://core42.ai/files/Whitepaper-Jais-30B.pdf>
- [13] Core42. Compass API frequently asked questions. <https://www.core42.ai/compass/documentation/compass-api-faqs>
- [14] Mohamed bin Zayed University of Artificial Intelligence. Meet Jais. 30 August 2023. <https://mbzuai.ac.ae/news/meet-jais-the-worlds-most-advanced-arabic-large-language-model-open-sourced-by-g42s-inception/>
- [15] Mohamed bin Zayed University of Artificial Intelligence. Inception Cerebras and MBZUAI release Jais 2. 9 December 2025. <https://mbzuai.ac.ae/news/inception-cerebras-and-mbzuai-release-jais-2-the-next-generation-of-the-worlds-leading-arabic-open-weight-llm/>
- [16] HUMAIN. HUMAIN Chat powered by ALLAM launch. 25 August 2025. <https://www.humain.ai/en/news/humain-chat-launch/>
- [17] HUMAIN. HUMAIN IQ and ALLAM model information. <https://www.humain.ai/en/humain-iq/>
- [18] Saudi Data and AI Authority. ALLAM large language model. <https://sdaia.gov.sa/en/SDAIA/about/Files/AI/ALLaM.pdf>
- [19] UAE Government. Data protection laws. <https://u.ae/en/about-the-uae/digital-uae/data/data-protection-laws>
- [20] UAE Government. Federal Decree Law Number 45 of 2021 on the Protection of Personal Data. <https://u.ae/en/information-and-services/justice-safety-and-the-law/cyber-safety-and-digital-security/data-protection-laws>
- [21] Saudi Data and AI Authority. Personal Data Protection Law. <https://dgp.sdaia.gov.sa/wps/wcm/connect/ce84f976-2fdf-4a76-a8c4-638b7ca9e3ba/Personal%2BData%2BEnglish.pdf>
- [22] Saudi Data and AI Authority. Guide to the Saudi Personal Data Protection Law. https://dgp.sdaia.gov.sa/wps/wcm/connect/85a664fd-6e4e-4edb-bab4-1cddb02ba6d5/Guide%2Bto%2Bthe%2Bsaudi%2BPDP%2Blaw%2BHANDBOOK__%D9%90EN.pdf
- [23] Saudi Data and AI Authority. Data governance platform knowledge centre. <https://dgp.sdaia.gov.sa/wps/portal/pdp/knowledgecenter/>
- [24] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework. <https://www.nist.gov/itl/ai-risk-management-framework>
- [25] National Institute of Standards and Technology. AI Risk Management Framework resource centre. <https://airc.nist.gov/airmf-resources/airmf/>

- [26] National Institute of Standards and Technology. Generative AI Profile NIST AI 600-1. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- [27] National Institute of Standards and Technology. Secure Software Development Practices for Generative AI. <https://www.nist.gov/publications/secure-software-development-practices-generative-ai-and-dual-use-foundation-models-ssdf>
- [28] UNESCO. Recommendation on the Ethics of Artificial Intelligence. <https://www.unesco.org/en/articles/recommendation-ethics-artificial-intelligence>
- [29] UNESCO. Arabic language Recommendation on the Ethics of Artificial Intelligence. <https://www.unesco.org/ar/artificial-intelligence/recommendation-ethics>
- [30] UNESCO. Coalition for Linguistic Diversity in Artificial Intelligence. <https://www.unesco.org/ethics-ai/en>
- [31] World Intellectual Property Organization. Artificial Intelligence and intellectual property. <https://www.wipo.int/en/web/frontier-technologies/artificial-intelligence/index>
- [32] OECD. OECD AI Principles. <https://oecd.ai/en/ai-principles>
- [33] OECD. OECD AI Incidents Monitor. <https://oecd.ai/en/incidents>
- [34] International Organization for Standardization. ISO IEC 42001 artificial intelligence management systems. <https://www.iso.org/standard/81230.html>
- [35] International Organization for Standardization. ISO IEC 23894 artificial intelligence risk management. <https://www.iso.org/standard/77304.html>
- [36] International Organization for Standardization. ISO IEC 5259 data quality for analytics and machine learning. <https://www.iso.org/standard/81088.html>
- [37] Inception. Jais family model card. <https://huggingface.co/inceptionai/jais-family-590m>
- [38] Technology Innovation Institute. Falcon model catalogue. <https://huggingface.co/tiiuae>
- [39] MLCommons. MLPerf Inference benchmarks. <https://mlcommons.org/benchmarks/inference-datacenter/>
- [40] ArabicMMLU. Arabic Massive Multitask Language Understanding benchmark. <https://aclanthology.org/2024.findings-acl.334/>
- [41] ACL Anthology. Jais and Jais chat Arabic centric foundation and instruction tuned models. <https://aclanthology.org/2024.acl-long.518/>
- [42] ACL Anthology. Arabic natural language processing benchmark research. <https://aclanthology.org/venues/arabicnlp/>
- [43] Qatar Computing Research Institute. Fanar Arabic foundation model. <https://www.qcri.org/research/fanar>
- [44] Qatar Ministry of Communications and Information Technology. Fanar platform. <https://www.mcit.gov.qa/en/fanar/>
- [45] Google Research. MADLAD 400 multilingual model research. <https://research.google/blog/madlad-400-a-multilingual-and-document-level-large-audited-dataset/>
- [46] Common Voice. Arabic speech dataset. <https://commonvoice.mozilla.org/ar/datasets>
- [47] Unicode Consortium. Arabic script and encoding resources. <https://www.unicode.org/charts/PDF/U0600.pdf>
- [48] International Telecommunication Union. AI for Good and multilingual inclusion. <https://aiforgood.itu.int/>
- [49] World Bank. Digital development overview. <https://www.worldbank.org/en/topic/digitaldevelopment/overview>
- [50] FinOps Foundation. FinOps Framework 2025. <https://www.finops.org/wp-content/uploads/2025/05/English-FinOps-Framework-2025.pdf>

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.