Beyond Compliance:

The Positive-Alignment Frontier in Islamic Capital

By Amjad Hammad, MD, MBA, and Rushdi Siddiqui, JD
Capital Formation Institute in Islamic Markets

Two investments reach an allocator. Both pass Shariah screening. Both are permissible. Both sit within the eligible universe.

Yet their real-world effects diverge materially. One contributes to the preservation of life, the expansion of knowledge, the stability of households, and the productive circulation of wealth. The other does not.

Compliance establishes admissibility. It does not establish preference.

That distinction — long implicit — is becoming central to the future of Islamic capital. The timing is not coincidental.

What AI is about to do to compliance

Recent empirical work on large language models converges on a clear result: these systems perform well on retrieval, classification, and rule-based matching.

A 2025 rapid evidence review in *AI & Society* synthesizing 83 studies of legal use cases documents consistent AI competence at these tasks. Applied work on AI in Islamic fatwa and Shariah advisory reaches a parallel conclusion: mapping transactions to AAOIFI standards, surfacing analogous prior rulings, and triaging contracts for audit review are tractable applications, with measurable efficiency gains for Shariah supervisory boards.

These are not marginal functions. They are the operational core of eligibility-based screening — the same rule-based pattern matching that AI handles competently and will increasingly handle cheaply.

The implication is structural. The cost and scarcity that historically defined compliance infrastructure are likely to compress. The screening apparatus — built over decades by scholars, standard-setters, and supervisory boards — is becoming replicable at low marginal cost.

What was once a constraint is becoming a service layer.

Where AI fails — and why it matters

The same literature establishes a second, more important result. LLMs are reliable when asked to apply a single rule. They become unreliable when asked to synthesize across multiple competing principles, reconcile trade-offs, and reach an integrated judgment.

The evidence is consistent across three independent strands. A 2025 study of frontier reasoning models on hierarchical legal tasks found near-perfect performance on simple rule identification, and a collapse into the 11–34% range when the same models were asked to weigh competing principles. Research published in *Harvard Business Review* in March 2026, across more than 15,000 simulations of six leading models, found a persistent bias toward managerial buzzwords and consensus narratives over context-specific reasoning — what the authors term “trendslop.” Independent work from the Wharton group on idea generation has shown AI outputs clustering tightly around the centroid of plausible answers rather than expanding the option space.

These are not temporary limitations. They reflect the architecture of these systems.

LLMs are effective at rule application. They are unreliable at principled synthesis.

The layer Islamic capital has not yet formalized

Maqasid-based evaluation operates precisely at that second layer.

It is not a question of permissibility. It is a question of what an investment substantively advances — across life, knowledge, wealth, household integrity, and ethical coherence — and how those dimensions interact.

That evaluation requires trade-offs. It requires prioritization. It requires the reconciliation of competing goods under conditions of incomplete information.

It is, by construction, an integrated judgment problem.

And it is a problem for which neither conventional screening frameworks nor current AI systems provide a complete answer.

From avoidance to advancement

Islamic finance has, with considerable success, defined itself through a discipline of constraint. It has built a robust apparatus for determining what capital must avoid.

That achievement is foundational. It enabled scale, standardization, and institutional participation.

It is also incomplete.

A system defined only by exclusion cannot differentiate meaningfully within its own permissible universe. As Islamic capital moves into more complex asset classes — private markets, infrastructure, cross-border healthcare — that limitation becomes binding.

The next phase requires a complementary discipline: not only identifying what is admissible, but evaluating what is preferable.

There is a counter-version of this argument worth naming. The additive turn could itself collapse into a new form of box-ticking — maqasid evaluation reduced to a scoring rubric, the discipline of judgment displaced by the appearance of judgment. ESG offers the cautionary template. Whether positive alignment avoids that fate depends on whether the institutions building it resist the demand to make it scalable in ways that defeat its purpose.

Why the moat has moved

The competitive implication follows directly.

If compliance is becoming a commodity service, the analytical moat for Islamic capital is no longer the screening apparatus. It is the discipline of evaluating positive alignment — the layer that requires fiqh, scholarly tradition, institutional judgment, and the form of holistic synthesis at which current AI fails.

The institutions that build that discipline first will define the next decade of Islamic capital. The institutions that do not will find themselves competing on a screening function whose price is converging toward zero.

This shift is already visible. Sovereign wealth funds, pension boards, and institutional family offices are no longer satisfied with confirmation that their Islamic exposures avoid prohibited categories. They are asking what those exposures advance.

The screening apparatus delivers a white list — instruments and counterparties that pass eligibility tests. That is necessary. It is no longer sufficient. **The question that follows the white list — what does this investment substantively advance, and how does it compare to the alternatives also on the list — is where the next layer of institutional value will sit.** And it is a different question than the one ESG or sustainability screens are designed to answer. ESG asks whether an investment manages its externalities responsibly. Maqasid alignment asks something more demanding: whether the investment advances the substantive purposes the capital was meant to serve.

As preview, applying this lens at the sector level produces a structurally different picture than a Shariah screen alone. The figure below maps the economic sectors of public equity markets through a maqāṣid-alignment view. Conventional financial sectors — banks, insurance, financial intermediaries — appear in red, indicating exclusion from maqāṣid evaluation rather than a low alignment score. Methodology and metrics will follow.

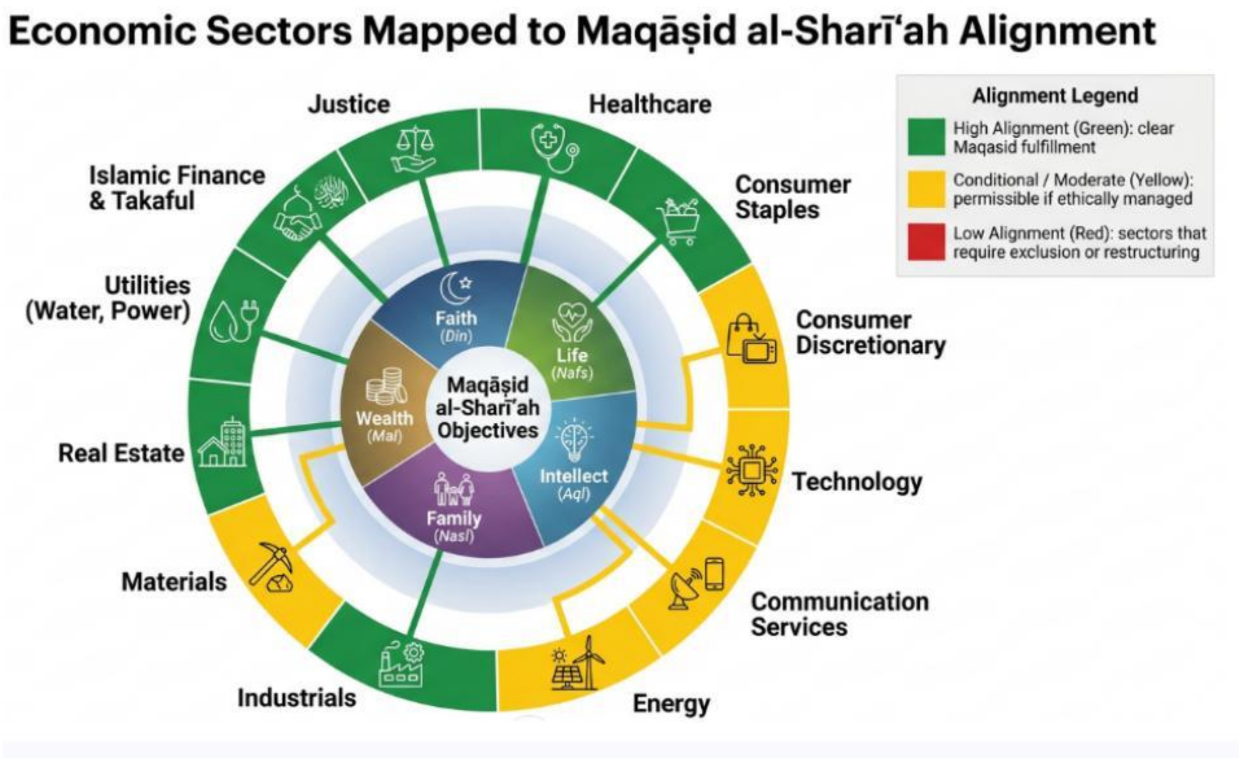

Existing frameworks were not designed to answer that question.

The frontier

At Capital Formation Institute in Islamic Markets, our work is focused on this transition: from eligibility-based screening to disciplined evaluation of positive alignment.

This requires a different type of architecture — one grounded in the maqasid tradition, but structured for institutional use; one that recognizes the strengths of AI at the compliance layer while remaining clear about its limits in evaluative judgment.

Operationally, the discipline produces three things: an articulation of what an investment advances, a documented record of the trade-offs accepted, and a defensible institutional view that an allocator can present to a board, a beneficiary, or a sovereign LP. The screen still operates. The discipline sits above it.

The objective is not to replace existing frameworks. It is to extend them into the domain they do not cover.

We will share more on this work in due course.

Compliance establishes the boundary. It does not determine direction.

As the cost of enforcing that boundary approaches zero, direction becomes the primary source of differentiation.

Compliance can be coded.

Maqasid alignment cannot.

That is the frontier Islamic capital is now approaching.

References

Billah, M. M., & Zahid, Z. (2026). Artificial Intelligence in Islamic Fatwa and Shariah Advisory: Opportunities, Limitations, and the Imperative of Human-in-the-Loop Shariah Governance. *International Journal of Innovative Inventions in Social Science and Humanities*, 3(1), 103–115.

Girotra, K., Meincke, L., Terwiesch, C., & Ulrich, K. T. (2023). Ideas are Dimes a Dozen: Large Language Models for Idea Generation in Innovation. Wharton School Research Paper. SSRN 4526071.

Kelsall, J., Tan, X., Bergin, A., et al. (2025). A rapid evidence review of evaluation techniques for large language models in legal use cases. *AI & Society*.

Meincke, L., Mollick, E. R., & Terwiesch, C. (2024). Prompting Diverse Ideas: Increasing AI Idea Variance. Wharton School Research Paper. SSRN 4708466.

Romasanta, A., Thomas, L. D. W., & Levina, N. (2026, March 16). Researchers Asked LLMs for Strategic Advice. They Got "Trendslop" in Return. *Harvard Business Review*.

Zhang, L., Grabmair, M., Gray, M. A., & Ashley, K. D. (2025). Thinking Longer, Not Always Smarter: Evaluating LLM Capabilities in Hierarchical Legal Reasoning. *arXiv* 2510.08710. Forthcoming, CSLAW '26.