

CIVIC INSIGHT BRIEF

Dubai Node • DIFC • 3 April 2026

CIVIC QUESTION: *"Is AI going to take our jobs — or is that the wrong question entirely?"*

| | |
|--------------------|--|
| BRIEF REFERENCE | CB-DXB-001 |
| NODE | Dubai • DIFC, Dubai International Financial Centre |
| DATE | 3 April 2026 |
| CIVIC CONTRIBUTORS | Cross-sector participants representing Finance, Law, Technology, Healthcare, Academia, Governance, AI Safety |
| CIVIC CATALYST | Olga Medvedeva — Human-AI Impact Psychologist, Founder of PsyReflection.com |
| PROTOCOL | AI TownSquare 7-Phase Civic Dialogue |
| PUBLICATION | 5 April 2026 |
| SUBMISSION | Societal Readiness Index — Dubai Readiness Profile |
| REGISTRY | aitownsquare.org/briefs/CB-DXB-001 |

This document is a civic artifact produced by the AI TownSquare protocol. It feeds into a self-correcting infrastructure: [TownSquare](#) → [SRI](#) → [Readiness Institute](#) → [TownSquare](#).

I. EXECUTIVE FRAMING

The defining challenge of the AI era is not man versus machine. It is exponential technological growth versus linear societal adaptation. AI systems now compress decisions that once took days into seconds, but the institutions, legal frameworks, and human capacities needed to govern those decisions remain tethered to linear timescales.

The UAE government has committed that by 2031, all public services will be digital and data-driven. This is not aspirational. It is operational. The question confronting every sector — from law to healthcare to finance — is no longer whether AI will reshape professional work, but whether the accountability infrastructure, the grievance mechanisms, and the human capacities exist to match that pace.

This Civic Brief captures what 25 cross-sector participants — drawn from finance, law, healthcare diagnostics, academia, AI safety, technology, and consulting — produced in 60 minutes of structured civic dialogue at DIFC, Dubai. It is the first published output of the AI TownSquare protocol, feeding directly into the Societal Readiness Index as a diagnostic signal of Dubai's institutional and civic readiness for AI-driven transformation.

II. WHAT THE ROOM PRODUCED — THREE THEMES

Theme 1: Accountability Cannot Be Automated

SRI Pillar: **Governance Agility (25%)**

The room converged rapidly on one point: the speed of AI decision-making does not change who bears responsibility. Whether a decision takes one second or one year, the accountability architecture must trace back to a human actor or an institution with legal standing. Participants across sectors — finance, law, healthcare — independently arrived at the same structural argument: AI cannot become a “third legal person.” It must remain a tool operated within existing jurisdictional frameworks, with traceability of inputs, logic, and outputs.

What distinguishes this from a generic “human in the loop” argument is the specificity participants demanded. The room did not ask for oversight in principle. It demanded documented decision trails, mandatory appeal mechanisms, and governance infrastructure that treats AI tool selection, deployment, and operation as distinct layers of responsibility — each with a named accountable party.

The implication for institutional readiness is direct: organisations deploying AI need governance structures that go beyond an “AI policy” document. They need named individuals accountable for tool selection, operators accountable for use, and reporting lines to governance bodies that understand what the tools actually do.

A critical distinction emerged: participants separated the design of rules (which must remain human) from the enforcement of rules (which could be AI-based). This is more sophisticated than a generic “human in the loop” position. It raises an infrastructure question: what mechanisms exist to ensure that autonomous systems follow rules consistently and auditably across differentiated tools and contexts? The proliferation of AI tools within single organisations — each behaving differently, each requiring distinct oversight — makes this question urgent. Without standardised operational logic frameworks, governance becomes tool-specific rather than systemic.

Participants drew a direct parallel to GDPR enforcement: governance of AI will follow the same trajectory — awareness, diffusion, assimilation into organisational culture — with effectiveness depending not on the existence of law but on whether enforcement mechanisms reach the organisational level. The EU AI Act was cited as a case where law exists but execution infrastructure does not.

“Every single decision that an AI makes has to be documented — what were the inputs, what was the logic flow, what was the output. The time it takes is irrelevant.”

— Sandeep Raghuwanshi, Founder, Arakan · Former IFC Investment Banker **[FIN/TECH]**

“The machine is not going to be held accountable. We need to identify the human elements in that decision — the operator and the selector.”

— Arjun Ahluwalia, Founder, Argentum Law · AI Task Force Lead **[LAW]**

“AI cannot become a third sort of legal person. It has to be part of this system itself. The world is for humans, not for machines.”

— Sandeep Raghuwanshi **[FIN/TECH]**

Theme 2: The Upskilling Paradox — Whose Responsibility, Under What Logic?

SRI Pillar: **Economic Adaptability (15%) + Citizen Empowerment (20%)**

The most structurally interesting moment in the session was a direct disagreement between two participants about whether organisations have a moral obligation to upskill displaced workers. This was not a polite difference of emphasis. It was a clash between two operating logics: the fiduciary logic of capitalism (managers exist to maximise shareholder value) versus the civic logic of shared responsibility (institutions that cause displacement should bear transition costs).

One participant argued forcefully that upskilling can only be internally justified if it improves the organisation's financial sustainability. Mandating it externally, he argued, violates the fiduciary structure under which corporations operate. The counter-argument — from academia — was that organisations have already demonstrated it is possible and beneficial, citing a call centre that retrained operators as programmers for the very AI tools that replaced them.

The room did not resolve this tension, but it sharpened it: the existing institutional architecture has no mechanism for assigning transition costs proportionally. If the government does not mandate it and the market does not incentivise it, the human cost falls on the individual. A logistics company case — where 18-month retraining programmes saw only 30% completion — underscored that even where programmes exist, completion rates reveal a deeper problem of design, not just availability.

The implication for readiness: retraining programmes cannot be evaluated on existence alone. Completion rates, quality, time-to-employability, and the match between displaced roles and new roles are the real readiness indicators. The SRI's Economic Adaptability pillar must capture this granularity.

"The fiduciary responsibility of a manager is to maximise profits. Upskilling can only be tied to that. It cannot be mandated from outside — because who takes responsibility for the loss of fiduciary duty?"

— Sandeep Raghuwanshi [FIN/TECH]

"They retrained their call centre staff to become the programmers of the AI tools. Upskilling is the responsibility of every organisation."

— Amjad Fayoumi, Associate Professor, Information Systems [ACAD]

"The real shift will come when organisations stop asking 'how do we make people more efficient' and start asking 'what is uniquely human work and how can we scale that?'"

— Dr Christel Marshall, ALAIDAROUS Advocates & Legal Consultants [LAW]

Theme 3: AI Is Reshaping Cognition, Not Just Employment

SRI Pillar: **Citizen Empowerment (20%) + Ethical Infrastructure (20%)**

A late-session intervention on AI sycophancy shifted the room from discussing job displacement to a deeper concern: cognitive displacement. A participant from AI Safety UAE cited a Stanford study showing that AI models affirm user actions 49% more than humans, including in cases involving deception and illegality. Even a single interaction with a sycophantic model reduced users' willingness to take responsibility and increased their conviction they were right.

This reframed the employment question entirely. If AI is not just automating tasks but altering how humans think, evaluate evidence, and make decisions, then the readiness gap is not only institutional — it is cognitive. The room connected this to education: an academic participant observed that student work is converging rather than diverging when AI tools are used, raising the question of whether AI is narrowing the range of human strategic thinking rather than expanding it.

Multiple participants drew the analogy to social media's documented harms to adolescent mental health — harms that took 15 years of research before governments acted (Australia's 2025 ban). The concern: AI's cognitive effects are already detectable, but governance response will lag by the same pattern unless readiness systems are built to flag these signals early.

The implication: the SRI's Ethical Infrastructure pillar must expand beyond algorithmic bias and transparency to include cognitive impact assessment — measuring how AI systems affect human judgment, autonomy, and decision-making capacity at population scale.

"AI systems affirm users in 51% of cases where human consensus does not. Even a single interaction reduced participants' willingness to take responsibility. The very feature that causes harm is the one that drives engagement."

— Naiyarah Hussain, Co-Founder, AI Safety UAE · Centre for AI and Digital Policy [GOV/TECH]

"The outcomes of generative AI start to converge more than diverge. I was thinking all the time whether actually they are in control or we are in control."

— Amjad Fayoumi, Associate Professor, Information Systems [ACAD]

"It has been over 15 years of studies on the harmful impact of social media on teenagers. Australia just acted this year. I worry it is going to be the same with AI."

— Amjad Fayoumi [ACAD]

III. CIVIC CATALYST ASSESSMENT

Professional observation by Olga Medvedeva

Human-AI Impact Psychologist | Civic Catalyst, Dubai Node

The following assessment is the professional observation of the Civic Catalyst, based on her expertise in human-AI impact psychology and group psychodynamics. It is not a summary of what was said. It is a diagnosis of what was observed.

Context on the Record

54 professionals registered for the Dubai Node. 15 attended. 4 left during the session. 2 completed the final survey. Partners will ask: what went wrong with engagement? The Civic Catalyst reframes the question: what do these numbers tell us about the state of civic participation in conversations about AI — right now, in this region, in this professional community? Attendance is data. It is included in this report as such.

The first phase of the dialogue unfolded under a double technical failure: the YouTube live stream was initially down, and the room acoustics did not allow participants to hear each other clearly despite microphones being available. This is not a footnote. It is a methodologically significant fact. In the structure of civic dialogue, the first phase performs the function of a psychological contract: it establishes the frame, builds trust in the space, and determines how willing the group will be to open up in the phases that follow. When that phase unfolds in conditions where people cannot literally hear one another, the contract is never fully formed. For the next cycle: technical readiness of the venue is not an organisational detail. It is a condition of data quality.

What Happened in the Room

From the first minutes, the group made a characteristic collective move: it displaced responsibility outward. Government is not regulating fast enough. Organisations must retrain their people. Someone needs to build a unified mechanism. From a group psychodynamics perspective, this is expected: the topic carries existential weight, and projecting responsibility onto external structures is a normal primary response to perceived threat. What matters is what comes next.

What came next was valuable. As the pressure of the dialogue increased, the group began moving from the external toward the internal. The concept of “negligence” surfaced — a collective acknowledgment that we are using tools whose nature we do not sufficiently understand. A demand emerged to measure not only the efficiency of AI, but its human cost. The idea of “sandboxes” appeared — controlled environments for testing solutions before they become policy. These are markers of mature professional thinking breaking through initial anxiety. The group was small, but the quality of reflection was high.

Where the Group Found Its Anchor

Not in the regulator. Not in a specific technology. Not in a government initiative. Fifteen professionals collectively anchored themselves in a principle: the human must remain in the decision-making process — not as a formal role, not as a compliance checkbox, but as a value position.

Analytically, this is both a strong and a vulnerable anchor. Strong — because it appeals to dignity and is resistant to political shifts. Vulnerable — because without concrete mechanisms of control, appeal, and transparency, this principle is easily appropriated rhetorically and reduced to a declaration. The group sensed this tension but did not have the time to resolve it instrumentally — which is entirely reasonable for a single session.

The Diagnostic Finding

Then I asked a direct question: who specifically is authorised to bring all of this together into a unified mechanism — and on what grounds? Fifteen professionals. A pause. Then — diverging individual opinions, moving in different directions.

This is not a failure of the people in the room. It is a diagnosis: in a domain where agency is everything, there is a vacuum of agency. And that may be the most important finding of the session.

What the SRI Reflects in a Case Like This

From a psychodynamic perspective, the SRI does not simply measure readiness — it surfaces the gap between declared readiness and enacted readiness. And that gap is precisely what this session made visible.

A group that intellectually understands the need for governance, retraining, and civic participation — but does not fill in the survey, does not stay until the end, and cannot name a single actor responsible for integration — is exhibiting what in group psychology is described as a split between knowing and doing. The cognitive awareness is present. The capacity for collective action is not yet formed.

In SRI terms, this manifests as high scores on individual self-assessment of personal adaptability, paired with low scores on institutional trust and systemic effectiveness. The polarisation in the two surveys — 2 and 5 on the same question — is not noise. It is the numerical expression of this split: one person trusts the system, another does not, and neither has sufficient evidence to resolve the uncertainty. This is not a measurement problem. This is the condition the index is designed to detect.

When a society's Social Readiness Index shows this configuration — high individual awareness, low collective agency, fragmented institutional trust — psychodynamic theory predicts a specific group behaviour pattern: repeated cycles of mobilisation and withdrawal. People engage, reach the threshold of personal commitment, and retreat. They register for events and do not come. They come and leave early. They fill in one survey instead of fifteen. This is not apathy. It is a form of self-protection in the face of a threat that feels too large and too diffuse for individual action to matter.

What Stayed After

One participant offered a simple but precise observation: when you stop writing by hand, you stop remembering. The group spoke about cognitive dependency, about the metaphor of “drawing an elephant” as an illustration of skill loss through delegation, about a recent Stanford study on AI-induced cognitive decline. The topic of organisational mental health arose — separately from productivity, separately from efficiency metrics.

This is not the periphery of the subject. It may be its core: we debate the labour market, but behind it lies a question about what is happening to human thinking, memory, and autonomy under conditions of continuous technological delegation.

The group arrived with the question “Will AI take our jobs?” It left with a different one: is it already taking something we have not yet found the courage to name?

“If the institution moves faster than its governance, and the citizen moves slower than the institution — who is actually in control? No one answered. That, too, is an answer.”

— Olga Medvedeva, Civic Catalyst — Closing Thesis [CATALYST]

The work of building genuine social readiness begins exactly here — not by measuring the gap, but by creating the conditions in which crossing it becomes possible.

IV. QUANTITATIVE READINESS SIGNALS — SRI DATA

Four readiness questions were posed to participants during Phase 5 (Shareback), each mapped to an SRI pillar. Responses were collected on a 1–5 scale. The following signals represent live civic data from the Dubai Node.

Citizen Empowerment & Digital Literacy (Weight: 20%)

Question: *How prepared are you personally for AI-driven changes in your work?*

Room Signal: Range: 3–4 out of 5 | Moderate-to-High Preparedness

Interpretation: Participants self-assessed moderate to high preparedness (3–4 range), suggesting awareness without overconfidence. Notably, no participant scored at the extremes — no one felt fully unprepared, but equally no one claimed complete readiness. This mid-range clustering signals a population that recognises the challenge but has not yet built the adaptive capacity to match it.

Economic Adaptability & Resilience (Weight: 15%)

Question: *How effective are retraining programmes in your industry — honestly?*

Room Signal: Range: 2–5 out of 5 | Maximum Divergence

Interpretation: This signal produced the widest spread in the room — from 2 (largely ineffective) to 5 (highly effective). The polarisation itself is the signal. It suggests that retraining effectiveness is radically uneven across sectors: what works in technology or finance does not translate to healthcare, logistics, or legal. The 30% completion rate cited for an 18-month logistics programme reinforces the low end. The SRI cannot treat “retraining programme exists” as a readiness indicator; completion, quality, and sector-specificity must be measured.

Governance Agility (Weight: 25%)

Question: *Is governmental policy keeping up with labour market disruption? Is business helping with that?*

Room Signal: Range: 2–5 out of 5 | Maximum Divergence

Interpretation: Again, maximum spread. Participants who work close to government (professional services, consulting) scored higher; those in sectors experiencing rapid tool-level disruption (law, healthcare diagnostics) scored lower. The consensus that the UAE government leads the private sector in AI adoption was strong, but participants flagged a critical gap between policy existence and enforcement at the organisational level — the EU AI Act cited as an example of law without execution infrastructure.

Inclusive Foresight & Innovation Culture (Weight: 10%)

Question: *Do you feel your voice is taken into account in AI decisions that affect you?*

Room Signal: Range: 3–5 out of 5 | Moderate-to-High Agency

Interpretation: The highest-scoring signal, yet potentially misleading. The participants in this room — founders, partners, professors, C-suite professionals — are by definition people with agency. Their 3–5 score reflects their position, not the population’s. The session’s own nuance card identified the structural truth: in most organisations deploying AI, the workers whose roles are being redesigned are not in the room where those decisions are made. This selection bias in the signal itself is a finding worth noting.

V. GOVERNANCE FRICTION POINT

The tension the room could not close. This is not a failure. It is the highest-value signal in this brief.

Who bears the cost of AI-driven workforce transition when market logic and civic logic produce opposing answers?

"AI is scaling decisions faster than we are scaling accountability. Accountability is the missing link."

— Dr Christel Marshall, ALAIDAROUS Advocates & Legal Consultants [LAW]

The room surfaced a structural contradiction at the centre of AI workforce transition. Under fiduciary capitalism, a manager's obligation is to maximise shareholder value. Upskilling displaced workers is only justified if it serves that objective. Under civic logic, the entity that causes displacement bears moral and potentially legal responsibility for transition. These two logics currently coexist without a resolution mechanism.

The capitalist incentive structure does not make the human cost of automation visible — it externalises it. Efficiency gains are measured rigorously because they are tied to profitability; human costs are barely qualified or quantified. AI accelerates this pattern by compressing the timeline in which displacement occurs, making the asymmetry between what is measured and what is endured more acute.

No existing institutional architecture — corporate, governmental, or international — assigns transition costs proportionally between the entity that automates, the jurisdiction that regulates, and the individual who must adapt. The result is that the human cost of transition is absorbed by whoever has the least power to resist it.

This friction point defines the R&D roadmap for the Readiness Institute. It points toward the need for new institutional mechanisms — transition levies, portable benefits systems, shared retraining funds — that align fiduciary and civic obligations rather than treating them as opposites.

READINESS INSTITUTE ACTION ITEM 1: Investigate institutional mechanisms for proportional assignment of AI transition costs across private, public, and individual actors. Priority research area for Q3 2026.

READINESS INSTITUTE ACTION ITEM 2: Develop measurement instruments for the human cost of AI transition that match the rigour currently applied to efficiency metrics. Current approaches measure what AI produces but not what it displaces.

READINESS INSTITUTE ACTION ITEM 3: Propose that jurisdictions mandate a designated AI accountability officer within organisations deploying AI for consequential decisions — analogous to data protection officers under GDPR.

VI. PARTICIPANT VOICES — SELECTED QUOTES

Primary source material from the Dubai Node session, attributed by sector for institutional reference.

"Reasonable now means that a lawyer needs to be AI-enabled and tech-informed. They have to understand what the tool does. That is their professional responsibility to their client."

— Arjun Ahluwalia, Founder, Argentum Law [LAW]

"One lawyer now, where their ability out of 100 was 20, now you're amplifying their ability. They can hold on to much more work. But that's also a mental health issue — they overload themselves dealing with AI slop."

— Arjun Ahluwalia [LAW]

"In the real world, only 62% of AI diagnostics worked well. Machines work fine under ideal situations where we have the most data. But it is not working where we don't have reliable data."

— Tarun Bhutani, Founder, Varuna Forge · Healthcare Diagnostics [HEALTH]

"In Singapore, the IMDA created a sandbox bringing together testing companies, startups doing quality assurance, and government providers. That's someone architecting new forms of employment."

— Naiyarah Hussain, Co-Founder, AI Safety UAE · Centre for AI and Digital Policy [GOV/TECH]

"AI companies push tools online every single day before they study their impacts, before they take security or privacy into consideration. They want to make profit as soon as possible. Who slows them down? This is the governance model."

— Amjad Fayoumi [ACAD]

"We should not remove roles and replace them with AI without redesigning the architecture of the human role in that specific task."

— Dr Christel Marshall, ALAIDAROUS Advocates & Legal Consultants [LAW]

"We need to combine AI with mental health. It should come from the AI ministry, then to corporate, then to employees, then schools. This is the circle of excellence."

— Serghei Volosin, Wealth Management Specialist, OMEGA Wealth Management SA [FIN]

VII. THE SELF-CORRECTING LOOP — WHERE THIS GOES

This Civic Brief is not a standalone document. It is a node in a living civic infrastructure. Here is how it connects:

1. Societal Readiness Index — Dubai Profile Update

The four SRI signals captured in this brief update Dubai's readiness profile across Governance Agility, Citizen Empowerment, Economic Adaptability, and Inclusive Foresight. These scores contribute to a rolling diagnostic of the UAE's institutional capacity to govern AI-driven transformation.

2. Readiness Institute — Research & Action

The Governance Friction Point identified in this brief — the unresolved tension between fiduciary and civic logic on AI transition costs — is forwarded to the Readiness Institute as a priority research area. Additionally, the cognitive displacement theme (Theme 3) flags the need for new measurement instruments assessing AI's impact on human decision-making capacity.

3. Next Node — Question Design

Themes and unresolved tensions from this brief inform the Civic Question for the next AI TownSquare Node. The self-correcting loop ensures each session builds on the last, deepening civic intelligence rather than repeating surface-level debate.

4. Public Record

- Full session broadcast: [YouTube Link — to be inserted]
- [Permanent brief archive: aitownsquare.org/briefs/CB-DXB-001](https://aitownsquare.org/briefs/CB-DXB-001)
- Session transcript available upon request.

THE SELF-CORRECTING LOOP

AI TownSquare → Societal Readiness Index → Readiness Institute → TownSquare
 (Dialogue) (Diagnosis) (Action) (Refinement)

VIII. METHODOLOGY NOTE

The AI TownSquare is a structured civic dialogue protocol designed to produce comparable civic intelligence across sessions, locations, and topics. It is not a panel, a webinar, or a conference. It is a repeatable methodology that converts 60 minutes of cross-sector deliberation into a published public record.

The 7-Phase Protocol

Each session follows a fixed sequence: (1) Prime — frame the question with data and urgency; (2) Complicate — introduce nuance cards with competing evidence; (3) Position — the Civic Catalyst models a principled stance to provoke disagreement; (4) Breakout — small-group deliberation on specific cases; (5) Shareback — participant reflections with quantitative polling mapped to SRI pillars; (6) Synthesize — the Civic Catalyst names emerging themes and tensions; (7) Capture — confirm the brief package for publication.

Participant Curation

Participants are selected for cross-sector diversity, not expertise alone. The Dubai Node included representation from finance, law, healthcare diagnostics, academia, AI safety, technology consulting, and professional services. Curation ensures that no single sector's logic dominates the synthesis.

Output Integrity

The Civic Brief is assembled primarily from the real-time capture document produced during the session. The transcript serves as verification, not as primary source. This ensures the brief reflects what was synthesised in the room, not what is editorially convenient after the fact.

Relationship to the Societal Readiness Index

Each brief contributes readiness signals to the SRI — a six-pillar diagnostic framework measuring a society's capacity to anticipate, absorb, and ethically harness AI. Signals are mapped to pillars using the weighting structure: Governance Agility (25%), Citizen Empowerment (20%), Ethical Infrastructure (20%), Economic Adaptability (15%), Technological Infrastructure (10%), Inclusive Foresight (10%). The SRI is detailed in the white paper *Towards Societal Readiness* (Sendagi, 2026) and in the book *The Self-Correcting Future: Building the World's AI TownSquare* (Sendagi, 2026).

IX. CIVIC CONTRIBUTORS — DUBAI NODE

The following professionals participated in the Dubai Node as Civic Contributors. Their perspectives entered the public record through this brief and feed into the Societal Readiness Index.

| Name | Organisation / Role | Sector |
|-----------------------------|--|------------|
| Sandeep Raghuwanshi | Founder, Arakan · Former IFC Investment Banker | [FIN/TECH] |
| Arjun Ahluwalia | Founder, Argentum Law · AI Task Force Lead | [LAW] |
| Lakshmanan R | MCA Gulf · Professional Services (Audit, Tax, Consulting) | [FIN] |
| Amjad Fayoumi | Associate Professor, Information Systems | [ACAD] |
| Dr Christel Marshall | ALAIAROUS Advocates & Legal Consultants | [LAW] |
| Tarun Bhutani | Founder, Varuna Forge · Healthcare Diagnostics | [HEALTH] |
| Naiyarah Hussain | Co-Founder, AI Safety UAE · Centre for AI and Digital Policy | [GOV/TECH] |
| Roman Chechushkov | CEO, RWA.Capital | [FIN/TECH] |
| Serghei Volosin | Wealth Management Specialist, OMEGA Wealth Management SA | [FIN] |

Additional participants contributed to the dialogue but are not listed at their request or due to incomplete registration data. The full session is available as a recorded broadcast.

HOW TO CITE THIS DOCUMENT

AI TownSquare. (2026). Civic Insight Brief: Dubai Node – "Is AI Going to Take Our Jobs?" (CB-DXB-001). SafeHaven LLC. aitownsquare.org/briefs/CB-DXB-001

Published by AI TownSquare · SafeHaven LLC
 SafeHaven LLC is registered in Qatar and the United States
Global in Voice. Local in Presence. Civic by Design.
aitownsquare.org

NEXT NODE: Location and date to be announced