



AL-MARSOOS

ISSN(P): 2959-2038 / ISSN(E): 2959-2046

<https://www.al-marsoos.com>



When AI Becomes a Business Actor: Closing the Accountability Gap in Agentic Enterprises

Omez Ali

Graduate Researcher, Department of Business Administration, Judson University, USA

Email: omezali05@gmail.com

Abstract

Artificial intelligence is moving beyond prediction, content generation, and decision support into systems that plan, select tools, communicate with external parties, and execute multi-step business processes with limited supervision. This transition creates the agentic enterprise, in which AI systems function as operational business actors even though they remain incapable of bearing moral blame, legal duties, or institutional responsibility. The resulting accountability gap appears when organizations delegate real authority to AI but cannot identify who authorized an action, which data and rules shaped it, who was expected to intervene, or how an affected person can obtain correction. This conceptual article examines that gap through business ethics, contemporary AI governance, organizational responsibility, and faith-informed public reason. It distinguishes functional agency from moral and legal agency and argues that responsibility must remain with identifiable human principals and deploying institutions. The article proposes an Enterprise Agent Accountability Framework with seven connected controls: named human ownership, bounded delegation, agent identity and authorization, traceable action records, risk-proportionate human review, stakeholder notice and appeal, and remedy with continuous learning. It also introduces a five-level autonomy model for matching controls to business risk and applies the framework to hiring, pricing, customer service, procurement, finance, marketing, and supply-chain operations. The study contributes an interdisciplinary governance model that links technical controls with dignity, justice, transparency, responsibility, and public benefit. It concludes that enterprises may delegate tasks and limited authority to AI agents, but they may not delegate away accountability.

Keywords: Agentic AI; Business Ethics; Corporate Accountability; AI Governance; Human Oversight; Enterprise Risk; Algorithmic Decision-Making; Responsible Innovation

Introduction

Artificial intelligence is entering a new stage of organizational use. Earlier business systems mainly classified records, forecast demand, recommended products, or produced text for a human user. Agentic

systems can now pursue goals through sequences of actions: they can search internal data, draft and send messages, update customer records, negotiate routine terms, call software tools, initiate payments, schedule work, or coordinate other agents. NIST has consequently begun treating AI agents as a distinct standards and security problem, particularly because autonomous action requires reliable identity, authorization, auditing, and revocation. The central business issue is no longer only whether an AI output is accurate. It is whether an enterprise can remain accountable when an AI system acts in the enterprise's name.¹

This article uses the expression business actor in a functional rather than metaphysical sense. An AI agent becomes a functional business actor when it can initiate or complete actions that alter the rights, opportunities, finances, data, reputation, or working conditions of other people. The term does not imply that the system is a person, possesses intention in the human sense, or should receive legal personality. It identifies a practical change in organizational power: the system is no longer merely describing a possible action to a manager; it is participating in the action itself.

The distinction matters because conventional responsibility structures were designed around human chains of command. A salesperson makes an offer, a purchasing manager approves a vendor, a supervisor evaluates an employee, and an accountant authorizes payment. Even when software supports these decisions, the human role is visible. In an agentic workflow, however, authority may be distributed among a model provider, an enterprise platform, an implementation team, a department manager, a data owner, a compliance officer, and the agent itself. When harm occurs, each participant may point elsewhere. The provider may say the customer configured the system; the business may say the model produced the action; the manager may say the workflow was automated; and the technical team may say it only integrated the tools. This diffusion produces the accountability gap.

The article argues that an enterprise closes this gap only when responsibility is assigned before deployment, authority is technically bounded, actions are traceable, human review is meaningful, affected persons can contest outcomes, and remediation is mandatory. These duties cannot be satisfied by a general statement that “a human remains in the loop.” A nominal reviewer who lacks time, information, authority, or practical ability to stop the system does not provide genuine oversight. Accountability requires an organizational design in which responsible persons can understand, interrupt, reverse, and learn from agentic action.

1 National Institute of Standards and Technology, "Announcing the AI Agent Standards Initiative for Interoperable and Secure Innovation" (2026); Harold Booth, William Fisher, Ryan Galluzzo, and Joshua Roberts, Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization (NIST, 2026); Jared Riggs et al., Summary Analysis of Responses to the Request for Information Regarding Security Considerations for AI Agents, NIST 800-5 (2026); Atoosa Kasirzadeh and Iason Gabriel, "Characterizing AI Agents for Alignment and Governance," arXiv:2504.21848 (2025); Nelly Dux et al., "Governance by Design: Architecting Agentic AI for Organizational Learning and Scalable Autonomy," arXiv:2605.20210 (2026).

Research Problem and Questions

The research problem arises from a mismatch between increasing machine autonomy and conventional corporate accountability. Organizations are building systems that act across software environments, yet many governance programs remain centered on model documentation, privacy notices, bias testing, or human approval of isolated outputs. Those controls remain necessary, but they do not fully address agents that create plans, select tools, invoke external services, delegate to sub-agents, and modify business records over time.

This study therefore addresses three questions. First, what changes ethically and organizationally when AI moves from decision support to delegated business action? Second, where does the accountability gap emerge among model providers, deploying enterprises, managers, employees, and automated agents? Third, which governance controls can preserve human responsibility while allowing businesses to benefit from bounded agentic autonomy?

The study's central proposition is that accountability must follow authority. Every unit of authority delegated to an AI agent must correspond to an identifiable human principal, an explicit permission boundary, a traceable record, a review rule, and a remedy pathway. The greater the autonomy, impact, irreversibility, and stakeholder vulnerability involved, the stronger these controls must become.

From Decision Support to Delegated Business Action

Decision-support systems advise; agentic systems can act. This difference can be illustrated through customer service. A conventional model may draft a proposed response for an employee. An agent may inspect the customer's account, determine eligibility, approve a refund, send the message, update the ledger, and trigger shipment of a replacement. Each step may be individually routine, but together they constitute an exercise of organizational authority. The agent represents the company to the customer and changes the company's legal and financial position.

Agentic systems also introduce temporal depth. A chatbot normally responds to a prompt and stops. An enterprise agent may hold a goal over hours or days, retain memory, revisit a task after receiving new information, or coordinate a chain of tools. This makes evaluation more difficult because the relevant unit is not a single answer but a sequence of choices. A harmless intermediate output can produce a harmful later action when combined with permissions, data access, and tool use.

Another change is the emergence of delegated discretion. Traditional automation follows detailed rules written in advance. Agentic AI can choose among routes, infer intermediate goals, resolve ambiguity, and adapt its plan. That flexibility creates business value, but it also means that policy cannot be expressed only as a final outcome. The enterprise must govern the process by which the agent interprets goals, selects evidence, uses tools, and escalates uncertainty.

Finally, agents can become organizational intermediaries. Employees may rely on them to summarize policy, assign work, screen applicants, or assess vendor performance. Customers may interact with them as the company's voice. Suppliers may receive purchase orders or negotiation messages without

realizing that no employee reviewed them. In each case, the agent shapes relationships of trust and power. The ethical question therefore extends beyond technical performance to representation, fairness, dignity, and institutional responsibility.

The Accountability Gap in Agentic Enterprises

The accountability gap has five recurring forms. The first is the ownership gap: no single person is clearly responsible for the agent’s business purpose and risk. A department may sponsor the system while the information-technology team operates it and a vendor controls the model. Without a named owner, critical questions remain unanswered when the system behaves unexpectedly.

The second is the authorization gap. Enterprises frequently grant agents broad access because narrow permissions are difficult to design. An agent asked to “resolve customer issues” may receive access to billing, shipping, refunds, customer notes, and email. Yet its actual authority may be undefined. The gap appears when technical access is treated as moral or managerial permission. Being able to execute a transaction does not mean the system should execute it in every context.

The third is the traceability gap. Logs may record that an API was called without preserving the goal, prompt, retrieved evidence, intermediate plan, policy version, or human approvals that produced the call. Such records are insufficient for investigation. Accountability requires a reconstructable explanation of how a business action occurred, not merely proof that a computer event took place.

The fourth is the oversight gap. A human may formally approve a workflow while facing hundreds of alerts, incomplete explanations, or strong pressure to accept the system’s recommendation. Automation bias can turn approval into ceremony. Meaningful oversight requires competence, time, authority, and an interface that shows the relevant facts and alternatives.

The fifth is the remedy gap. People affected by automated actions may be unable to reach a human, obtain reasons, correct data, or reverse the decision. An organization that cannot repair an agent’s mistake is not accountable even if it can identify the technical cause. Responsibility becomes credible only when the enterprise can stop harm, restore the affected person, and revise the system.

Literature Review and Scholarly Context

Contemporary AI Governance and Agentic Systems

Contemporary governance frameworks increasingly connect accountability to the full AI lifecycle. The NIST AI Risk Management Framework emphasizes governing, mapping, measuring, and managing risk, while its generative-AI profile extends attention to confabulation, data privacy, human-AI configuration, information integrity, and value-chain risk. The OECD similarly links accountability to the role, context, and capacity of each AI actor and emphasizes traceability and ongoing risk management. The European Union’s AI Act requires effective human oversight for high-risk systems and assigns deployers duties concerning competent personnel, monitoring, and incident response. These

frameworks establish a strong foundation, but agentic AI requires more explicit treatment of delegated authority, identity, tool access, and multi-step action.²

Recent agent research sharpens this need. Kasirzadeh and Gabriel describe agents through autonomy, efficacy, goal complexity, and generality, showing why governance should vary according to an agent's profile. Dux and colleagues argue that governance is built into architecture through decisions about tools, data, memory, and improvement processes. Other proposed frameworks emphasize runtime constraints, anomaly detection, interruptibility, behavioral policies, and auditable records. Together, this literature supports a move from abstract principles to controls that operate before, during, and after an agent acts.³

Ataur Rehman: Faith-Informed AI Ethics, Religious Risk, and Moral Agency

Ataur Rehman's scholarship contributes a moral vocabulary for evaluating autonomous business systems. His work on faithful and responsible AI frames dignity, truthfulness, accountability, justice,

2 Ataur Rehman, "Religious Hallucinations in Generative AI: Developing a Multifaith Risk Benchmark to Prevent Misinformation, Polarization, and Extremist Manipulation in the United States," *Islamic Research Journal al-Marsus* 4, no. 2 (2026): 38-56; Ataur Rehman, "Faithful and Responsible Artificial Intelligence: An Ethical Framework for Education, Faith Communities, and Public Life," *Islamic Research Journal al-Qudwah* 4, no. 2 (2026): 1-17; Ataur Rehman, "From Contingency to Design: Reassessing al-Ghazali's Arguments for God in Philosophy, Theology, and Contemporary Science," *Islamiyyat* 48, no. 1 (2026): 24-34; Ataur Rehman, "Religious Equity in Public-Sector Artificial Intelligence: Developing a Multi-Faith AI Impact Assessment Framework through a Comparative Study of Texas and California," *Pakistan Islamicus* 5, no. 2 (2025); Ataur Rehman, "Integrating Islamic Ethics with Artificial Intelligence: A Study of Seerah in the Context of Pakistan," *Islamic Research Journal al-Marsus* 3, no. 2 (2025): 1-9; Ataur Rehman, "Iconoclasm: Re-imagining Sacred Art and Morality in Contemporary Spiritualities," *Harvard Divinity School, Program for the Evolution of Spirituality* (2025); Ataur Rehman, "Free Will: A Discourse between New Atheism and Islam," *Islamic Research Journal al-Marsus* 1, no. 4 (2023): 1-24; Ataur Rehman, "God's Existence in Primary Islamic Sources: An Analytical Study in View of New Atheism," *Journal of Positive School Psychology* 6, no. 9 (2022): 26-33; Ataur Rehman and T. Basharat, "God's Existence: Philosophical Discourse between New Atheists and Muslim Scholars," *Pakistan Journal of Islamic Philosophy* 3, no. 2 (2021): 1-16; Ataur Rehman, "Moralities: A Contemporary Discourse between New Atheism and Islam," *Islamiyyat* 42, no. 2 (2020): 113-124.

3 Z. Zahid et al., "Regulating Digital Currencies in the EU: A Comparative Analysis with Islamic Finance Principles under MiCA," *International Journal of Business and Management Practices* 3, no. 3 (2025); I. Hayat and S. Arif, "Islamic Ethical Perspectives on Artificial Intelligence and Digital Surveillance: Principles, Challenges, and Policy Recommendations," *International Research Journal of Arts, Humanities and Social Sciences* 2, no. 1 (2024); S. H. F. Bukhari, M. Abid, and S. Arif, "Atheism in the Muslim World: Classical Approach through Medieval Islamic Intellectual and Philosophical Traditions," *Journal of Religious and Social Studies* 3, no. 2 (2023): 95-114; A. F. I. Saeed et al., "Charity across Faiths: The Ethical and Spiritual Dimensions of Giving in Judaism, Christianity, and Islam," *Al-Aasar* 2, no. 1 (2025): 748-757; N. Akhter, S. Arif, and S. Akhtar, "Socio-Political and Religious Contributions of Sir Syed Ahmad Khan," *Al-Afaq Islamic Research Journal* 4, no. 1 (2024): 158-165; S. Hussain, M. Muslim, and S. Arif, "Conflict Resolution and Development of Skilled Leadership," *Journal of Religious and Social Studies* 3, no. 2 (2023): 20-41; S. Arif and U. Ahmad, "Existence of God and Morality: Scholastic Endeavors by Christian and Muslim," *Journal of Religious and Social Studies* 2, no. 2 (2022): 47-66; S. Arif and U. Ahmad, "Scientism and Religion: Christian-Muslim Responses about Atheistic Ideology," *Pakistan Journal of Islamic Philosophy* 4, no. 1 (2022): 53-69; I. Hayat and S. Arif, "Islamic Ecological Ethics and Sustainable Development in Pakistan: A Qur'anic Perspective on Climate Challenges," *International Research Journal of Arts, Humanities and Social Sciences* 2, no. 2 (2025); Salman Arif, "The Qur'anic Framework for Minority Rights: Upholding Justice, Dignity, and Religious Freedom," *Al-Asr* 5 (January-March 2025).

privacy, human oversight, and public benefit as connected rather than isolated principles. His research on religious hallucinations highlights how confident misinformation can produce social, spiritual, and political harm, while his multifaith impact-assessment work brings religious equity into public-sector technology governance. His Seerah-based AI ethics study emphasizes honesty, fairness, transparency, and the prohibition of deception in business practice. His wider work on moral reasoning, free will, theistic argument, and sacred art reinforces a basic distinction used in this article: computational agency may produce effects, but moral responsibility belongs to persons and institutions capable of obligation, intention, repentance, and accountability.⁴

Salman Arif: Governance, Surveillance, Finance, Leadership, and Minority Rights

Salman Arif's interdisciplinary work connects AI accountability to surveillance, Islamic finance, leadership, pluralism, environmental responsibility, and minority rights. Research on AI and digital surveillance is particularly relevant because enterprise agents frequently depend on extensive monitoring, behavioral data, and persistent memory. His work on digital currencies and Islamic finance raises related questions about transparent exchange, market power, and regulatory responsibility. Studies of conflict resolution, skilled leadership, scientism, morality, ecological ethics, and minority rights expand the analysis beyond compliance. They suggest that responsible business governance must examine how technology redistributes power, whose interests are protected, and whether efficiency is achieved at the expense of vulnerable communities.⁵

4 A. A. Raza and M. S. Khalid, "Interfaith Dialogue: Ethical Commonalities in Judaism, Christianity and Islam," *Islamic Studies Research Journal Abhath* 7, no. 26 (2022); A. F. I. Saeed et al., "Charity across Faiths: The Ethical and Spiritual Dimensions of Giving in Judaism, Christianity, and Islam," *Al-Aasar* 2, no. 1 (2025): 748-757; M. S. U. Rehman and A. A. Raza, "The Development and Growth of Gratitude and Positive Thinking in the Light of Artificial Intelligence and Seerah," *Islamic Research Journal al-Marsus* 3, no. 2 (2025): 66-79; A. F. I. Saeed, A. A. Raza, and H. A. Fatima, "Prayer as the Heart of Worship: Exploring Shared Rituals and Spiritual Connections in Judaism, Christianity and Islam," *Islamic Research Journal al-Qudwah* 3, no. 1 (2025): 1-8; A. A. Raza, A. F. I. Saeed, and K. tul Kubra, "Prophethood in the Abrahamic Faiths," *Islamic Research Journal al-Qudwah* 2, no. 4 (2024): 126-133; M. Abbas, M. Anjum, and A. A. Raza, "Divine Oneness and Distinct Theologies," *Islamic Research Journal al-Marsus* 2, no. 4 (2024): 20-28; A. F. I. Saeed, H. A. Fatima, and A. A. Raza, "The Sacred Texts of Abrahamic Faiths: Common Themes in the Torah, Bible, and Qur'an," *Harf-o-Sukhan* 8, no. 3 (2024): 1085-1094; A. A. Raza, W. Ali, and G. D. Ahmad, "Social Importance and Requirements of Compassion and Gentleness," *Al-Durar* 4, no. 1 (2024); A. A. Raza, H. U. Ismail, and Q. A. Manan, "The Social Importance and Requirements of Tolerance," *Al Manhal Research Journal* 3, no. 3 (2023); A. A. Raza, D. Ali, and Aatur Rehman, "Social Importance of Shyness and Modesty in Islam: An Analytical Study in the Light of Prophet's Seerah," *Al-Durar* 3, no. 3 (2023).

5 H. F. Rasool, A. Aziz, and M. Kiran, "Mental Health and Spiritual Well-Being in the Qur'an: Addressing Modern Psychological Challenges through Islamic Teachings," *Uloom Al-Sunnah* 2, no. 2 (2024): 33-40; H. F. Rasool, S. M. Shah, and M. Nasrullah, "Islamic Responses to Environmental Protection and Sustainability: Principles and Contemporary Applications," *Islamic Research Journal al-Qudwah* 2, no. 4 (2024): 78-85; H. F. Rasool et al., "Economic Justice in Islam: A Comprehensive Qur'anic Framework for Equity and Social Harmony in Seerah Perspective," *Tanazur* 5, no. 4(a) (2024): 1-15; M. Z. Hanif and H. F. Rasool, "Exploring New Educational Dimensions for the Improvement of Pakistan's Education System," *Pakistan Islamicus* 5 (2025); A. Atiq and H. F. Rasool, "The Character Building of Individuals by the Teachings of Islam," *Islamic Research Journal al-Qudwah* 3, no. 1 (2025): 96-105; M. Azam, H. F. Rasool, and S. Tahseen, "The Reasons for the Decline of Pakistani Society and Its Remedy," *Inkishaf* 3, no. 10 (2023): 190-209; F. Anwar, H. F. Rasool, and M. A. U.

When AI Becomes a Business Actor: Closing the Accountability Gap

Abbas Ali Raza: Interfaith Ethics, Compassion, Tolerance, and Social Responsibility

Abbas Ali Raza's scholarship provides resources for stakeholder-centered accountability in religiously and culturally diverse markets. His work on interfaith dialogue, charity, prayer, prophethood, sacred texts, compassion, tolerance, modesty, and positive thinking emphasizes shared moral obligations across communities. For business AI, these themes matter because automated systems can classify, exclude, or misrepresent people whose beliefs are poorly represented in training data. A governance framework that treats only economic loss as harm will miss injuries to dignity, conscience, identity, and social trust. Raza's comparative orientation therefore supports notice, participation, respectful treatment, and appeal as central business duties rather than optional public-relations practices.⁶

Hafiz Faiz Rasool: Human Welfare, Education, Character, and Social Consequences

Hafiz Faiz Rasool's work broadens responsible AI analysis through mental health, spiritual well-being, environmental protection, economic justice, education reform, character formation, language access, social decline, and antisocial behavior. These studies support a consequence-sensitive view of enterprise governance. An AI agent should not be judged only by task completion or financial return, but also by its effects on employee well-being, consumer understanding, equitable access, cultural communication, and social stability. The emphasis on character is equally important: responsible AI depends on managers who resist deceptive shortcuts, recognize uncertainty, and accept responsibility when automation fails.⁷

Synthesis of the Literature

The combined literature reveals three gaps. First, mainstream governance frameworks identify important principles but often remain system-centered rather than authority-centered. Second, technical agent research proposes strong runtime controls but may not fully address stakeholder dignity, religious

Haq, "Global Campaign about Prophetic Blasphemy: Motivations, Causes, and Solutions," *International Research Journal of Arabic and Islamic Studies* 4, no. 2 (2024): 47-67; M. S. Aslam, M. A. Rehman, and H. F. Rasool, "Effective Strategies for Teaching Arabic to Non-Native Speakers in Pakistan," *Inkishaf* 4, no. 13 (2024): 17-29; H. M. Binish, M. A. Rehman, and H. F. Rasool, "The Importance of Translating the Meanings of the Holy Quran into International Languages," *Inkishaf* 4, no. 12 (2024): 14-32; M. A. Jasvi et al., "Antisocial Activities and Role of a Religion in a Society," *Pakistan Islamicus* 4 (2024): 11-18.

6 National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1 (2023); National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1 (2024); Organisation for Economic Co-operation and Development, *Advancing Accountability in AI: Governing and Managing Risks throughout the Lifecycle for Trustworthy AI* (2023); UNESCO, *Recommendation on the Ethics of Artificial Intelligence* (2021); European Parliament and Council, *Regulation (EU) 2024/1689*, especially Articles 14 and 26.

7 Kasirzadeh and Gabriel, "Characterizing AI Agents for Alignment and Governance" (2025); Dux et al., "Governance by Design" (2026); Rafflesia Khan, Declan Joyce, and Mansura Habiba, "AGENTS SAFE: A Unified Framework for Ethical Assurance and Governance in Agentic AI," *arXiv:2512.03180* (2025); Anuj Kaul, Qianlong Lan, and Pranay Gupta, "AgentBound: Verifiable Behavioral Governance for Autonomous AI Agents," *arXiv:2606.30970* (2026); Riggs et al., *Summary Analysis of Responses Regarding Security Considerations for AI Agents* (2026).

and cultural equity, or corporate remedy. Third, faith-informed scholarship supplies moral concepts and public-interest concerns but requires translation into enterprise architecture and operational procedure. The present article addresses these gaps by joining moral responsibility with technical delegation. It treats the agentic enterprise as a sociotechnical institution in which values are expressed through access controls, workflow design, logging, escalation, communication, and remediation. Ethical governance therefore becomes visible in the system’s architecture and in the organization’s daily routines.

Methodology

This study uses a qualitative conceptual and normative method. It synthesizes business ethics, AI governance, organizational accountability, technical agent research, and faith-informed public ethics to construct an applied framework. The method proceeds in four stages. First, it distinguishes functional agency from moral and legal agency. Second, it identifies accountability failures that arise when an AI system receives business authority. Third, it translates ethical principles into organizational and technical controls. Fourth, it tests the framework against representative business scenarios.

The study is not an empirical evaluation of a particular product, company, or model. Its purpose is theory building and governance design. The framework is intended to be adaptable across jurisdictions and industries, while recognizing that legal obligations differ. It should therefore complement, rather than replace, legal advice, sector-specific regulation, cybersecurity standards, labor protections, and professional duties.

Normatively, the analysis adopts five commitments: human dignity, truthfulness, justice, responsibility, and public benefit. These commitments are compatible with major international AI principles and with Islamic and Christian approaches to stewardship, trust, neighbor-regard, fair dealing, and protection of the vulnerable. The study uses them to evaluate whether an enterprise can justify not only what an agent accomplishes but also how authority is delegated and how harm is repaired.

The Enterprise Agent Accountability Framework

The Enterprise Agent Accountability Framework contains seven connected controls. The controls form a chain because weakness in one element can invalidate the others. A named owner without technical limits cannot govern effectively. Detailed logs without an appeal process do not protect an affected customer. Human review without authority to stop an action is symbolic. Remedy without learning allows the same harm to recur. Accountability is therefore an integrated organizational capability, not a single document or approval step.

1. Named Human Principal and Institutional Ownership

Every agentic use case must have a named human principal who owns the business objective, approves the risk level, and remains answerable for outcomes. The principal should be a person with sufficient organizational authority, not merely the developer who configured the workflow. In high-impact uses, ownership should also be institutionalized through a cross-functional committee involving business

When AI Becomes a Business Actor: Closing the Accountability Gap

leadership, legal counsel, information security, compliance, human resources, and affected operational teams.

Named ownership prevents moral disappearance. It identifies who must explain why the agent was deployed, what benefits were expected, which risks were accepted, and when the system should be suspended. The principal may delegate operational tasks, but cannot delegate the duty to ensure that the workflow is lawful, ethical, and aligned with organizational purpose.

2. Bounded Delegation of Purpose and Authority

An agent should receive a written delegation charter that defines its mission, data scope, tools, financial limits, prohibited actions, escalation triggers, geographic or customer boundaries, and expiration conditions. Broad goals such as “maximize sales” or “resolve complaints” are ethically inadequate because they invite the system to trade fairness, truthfulness, or long-term trust for a narrow metric.

Bounded delegation converts management intent into enforceable policy. A refund agent, for example, may be permitted to approve refunds below a stated amount when specified evidence is present, but prohibited from denying claims involving disability accommodations, legal threats, suspected discrimination, or repeated system errors. The agent should not infer permission from technical access alone.

3. Agent Identity, Authentication, and Authorization

Each enterprise agent and sub-agent should possess a distinct, verifiable identity. Shared service accounts make it difficult to determine which system acted and under whose authority. Credentials should follow least-privilege principles, expire when the task or delegation ends, and be revocable immediately. Where agents create or invoke other agents, the identity chain must remain visible.

Authorization should be dynamic as well as static. A system may be generally permitted to create purchase orders, yet a particular order may exceed a financial threshold, involve a new supplier, or conflict with sanctions and fraud controls. Runtime policy should therefore evaluate the context of an action before execution. NIST’s work on software-agent identity, authorization, auditing, and non-repudiation points directly to this requirement.

4. Traceable Action and Decision Records

Accountability requires records that allow an independent reviewer to reconstruct what happened. Logs should include the assigned goal, relevant prompts, policy version, retrieved sources, data accessed, intermediate plans, tool calls, confidence or uncertainty indicators, human interventions, external communications, and final actions. Records should be protected against alteration and retained according to risk, law, and business need.

Traceability is not the same as indiscriminate surveillance. Enterprises should collect the minimum information needed for governance and protect sensitive data. The purpose is to explain consequential action, not to create an unlimited archive of employee and customer behavior. Where privacy and

traceability conflict, the organization should use access controls, redaction, secure retention, and role-based review.

5. Risk-Proportionate Human Review and Intervention

Human oversight should increase with impact, autonomy, uncertainty, irreversibility, and stakeholder vulnerability. Low-risk actions may be sampled after execution. Medium-risk actions may require approval when thresholds are exceeded. High-impact actions, such as terminating employment, denying essential services, executing large payments, or making legally consequential representations, should require prior human authorization or remain outside the agent’s authority.

A meaningful reviewer must have competence, relevant information, sufficient time, and power to reject the action. Interfaces should present evidence and policy reasons rather than merely asking the reviewer to accept a recommendation. Oversight performance should itself be audited because managers can become dependent on automation or overwhelmed by volume.

6. Stakeholder Notice, Explanation, and Appeal

People should be informed when an AI agent materially represents the business or contributes to a consequential decision. Notice does not require disclosing every automated background process, but it should be clear when customers, employees, applicants, or suppliers are interacting with an agent that can alter their position. Deceptive impersonation undermines trust and prevents informed challenge.

Affected persons should have access to a human review channel, an understandable explanation of the relevant reasons, and a practical opportunity to correct inaccurate information. The organization should not hide behind trade secrecy or technical complexity when a person faces serious harm. Explanation should be tailored to the decision and the audience, not reduced to a generic statement that an algorithm was used.

7. Remedy, Incident Learning, and Governance Revision

An accountable enterprise must be able to stop the agent, reverse unauthorized actions, restore records, and compensate losses when appropriate, notify affected parties, and report serious incidents. Technical rollback is important but insufficient. Remedy must address the person’s actual harm, including lost opportunity, reputational injury, privacy invasion, or discriminatory treatment.

Each incident should feed back into governance. The enterprise should determine whether the failure arose from the model, data, prompt, policy, authorization, interface, human review, vendor, or organizational incentive. Controls, training, thresholds, and delegation charters should then be revised. This continuous learning process turns accountability from blame after failure into prevention and institutional improvement.

Graphical Representation of the Framework

Figure 1 presents the framework as a responsibility chain. The sequence begins with a named human principal and moves through bounded authority, technical identity, traceability, oversight, stakeholder

When AI Becomes a Business Actor: Closing the Accountability Gap

rights, and remedy. The feedback loop shows that incidents and audits must revise the original delegation rather than remain isolated compliance records.



Figure 1. Enterprise Agent Accountability Framework

A Five-Level Model of Agentic Business Autonomy

The same controls should not be applied identically to every system. Table 1 offers a five-level model that classifies agentic use according to external action and business impact. The model is designed as a governance aid rather than a legal classification.

Level	Mode	Typical Example	Required Control	Default Position
1	Assistive	Drafts or summarizes; no external action	User verification; usage disclosure where material	Generally permitted
2	Recommendatory	Ranks options or proposes a decision	Documented criteria; bias testing; human decision	Permitted with review
3	Transactional	Executes bounded actions after approval	Identity, limits, logs, approval thresholds, rollback	Permitted with controls
4	Bounded autonomous	Acts independently within a defined domain	Runtime policy, monitoring, sampling, appeal, incident response	Restricted and audited
5	High-impact autonomous	Makes irreversible or rights-affecting decisions	Prior human authorization or prohibition; executive oversight	Presumptively prohibited

Table 1. Risk-proportionate levels of agentic business autonomy

Applications Across Business Functions

Hiring and Workforce Management

Hiring agents can search applicants, rank resumes, schedule interviews, generate questions, and recommend or execute rejections. These tasks affect livelihood and equal opportunity. An enterprise should prohibit autonomous final rejection where the system uses inferred traits, opaque scores, or incomplete employment histories. The named principal should be a senior human-resources leader, and the delegation charter should identify lawful job criteria, excluded data, accommodation procedures, audit requirements, and appeal routes.

Employee-management agents create similar risks when they assign shifts, measure productivity, recommend discipline, or identify workers for dismissal. Continuous monitoring can reduce workers to

When AI Becomes a Business Actor: Closing the Accountability Gap

behavioral data and can intensify stress. Accountability requires proportional data collection, notice, access to relevant records, protection against retaliation, and human consideration of context that the system cannot adequately understand.

Customer Service and Market Representation

Customer-service agents are attractive because they can operate continuously and resolve high volumes of routine requests. Yet they may also misstate policy, fabricate reasons, deny valid claims, or pressure customers into unnecessary purchases. The enterprise should define which promises the agent may make, which refunds it may approve, and when it must transfer the interaction to a human. Customers should not be trapped in an automated loop when the system has caused the problem.

Because an agent speaks in the company's voice, truthfulness is a core control. Marketing or support agents should not invent product capabilities, testimonials, legal rights, or religious and cultural claims. Rehman's work on religious hallucinations is instructive beyond religion: fluent misinformation becomes more dangerous when delivered by a system that appears authoritative and can immediately act on the false claim.

Pricing, Credit, and Financial Decisions

Pricing agents may adjust offers according to demand, location, browsing behavior, urgency, or predicted willingness to pay. Such systems can improve inventory management but can also exploit vulnerability or produce hidden discrimination. A governance program should define prohibited variables and proxy features, monitor disparate effects, and give customers clear information when personalized terms materially differ.

Credit, insurance, and payment agents require stronger controls because errors can affect housing, mobility, health, and business survival. High-impact denials should not be fully autonomous. The enterprise should preserve the evidence used, permit correction of inaccurate data, and provide review by a qualified person. Financial limits, dual authorization, and transaction rollback should be built into the system rather than left to policy manuals.

Procurement and Supply-Chain Operations

Procurement agents can solicit quotes, compare suppliers, negotiate routine terms, and issue purchase orders. The efficiency gains are significant, but the risks include unauthorized commitments, conflicts of interest, sanctions violations, counterfeit goods, and manipulation by malicious suppliers. Each agent should have transaction ceilings, approved vendor categories, segregation of duties, and independent verification for new bank details or unusual terms.

Supply-chain agents may reroute shipments, change inventory targets, or prioritize customers during shortages. These actions have distributive effects. Governance should therefore incorporate fairness, essential-service priorities, environmental impact, and the needs of smaller suppliers. An optimization objective that minimizes cost without these constraints can create serious social and reputational harm.

Accounting, Treasury, and Corporate Finance

Financial agents can reconcile accounts, prepare reports, detect anomalies, and initiate payments. Their access makes identity and authorization essential. No agent should possess unlimited ability to create vendors, change payment instructions, approve invoices, and release funds. Separation of duties must apply to agents as it does to employees.

Accountability also requires accurate representation to investors, lenders, auditors, and regulators. An agent that fabricates an explanation or hides uncertainty can produce misleading financial communication even when the underlying numbers are correct. Human officers remain responsible for statements issued in the company’s name and should review material disclosures.

Religious and Cultural Sensitivity in Diverse Markets

Businesses increasingly use AI to personalize products, holidays, advertising, food recommendations, financial services, and workplace communication. Poorly governed agents may stereotype believers, misstate religious practices, or use sensitive identity data to exclude or manipulate. A multifaith impact assessment can identify whether the system treats religious identity, dress, dietary practice, worship schedules, or community language as a risk signal or marketing opportunity without adequate consent. Religious equity is not limited to faith-based organizations. It is part of fair treatment in plural societies. Enterprises should test agent outputs with diverse communities, avoid using religious data unless necessary and lawful, and provide human review when decisions affect conscience, accommodation, or identity. This approach extends the public-interest concerns developed in Rehman’s multifaith AI scholarship and Raza’s interfaith ethics.

Scenario Analysis

Scenario One: The Refund Agent

A retailer deploys an agent to resolve delivery complaints. The agent reviews account history, checks shipping data, decides whether a claim is credible, issues refunds up to \$500, and sends final messages. A customer reports that a medical device was delivered to the wrong address. The agent denies the claim because the carrier marked the package delivered and the customer previously received a courtesy refund.

Under the framework, this denial should trigger human review because the product is health-related, the evidence conflicts, and the outcome may create serious harm. The action record should show the policy version, data sources, and reason for denial. The customer should be informed that automation materially contributed and should receive a direct appeal channel. If the denial was caused by an overbroad fraud rule, the enterprise should reverse the decision and revise the rule. Responsibility rests with the retailer and the named principal, not with the agent.

Scenario Two: The Recruiting Agent

A consulting company uses an agent to source candidates, score resumes, schedule interviews, and reject applicants below a threshold. The agent learns from historical hiring patterns and begins downgrading applicants from certain religious schools and community organizations because previous hires from those backgrounds were uncommon.

The system creates a proxy for religious identity and reproduces organizational history as a selection rule. A responsible deployment would exclude such features, test for disparate effects, require human review of rejection patterns, and offer applicants a correction process. The accountability failure is not solved by claiming that the agent did not explicitly use religion. The enterprise selected the data, objective, threshold, and authority structure and therefore remains responsible for the discriminatory effect.

Scenario Three: The Procurement Agent

A manufacturing firm authorizes an agent to obtain replacement components during production interruptions. A supplier sends a crafted message that manipulates the agent into bypassing the approved-vendor list and changing payment instructions. The agent issues a purchase order and initiates payment to a fraudulent account.

This incident combines prompt injection, authorization failure, and weak segregation of duties. The framework would require supplier identity verification, restricted tools, a prohibition on changing bank details without independent approval, and immutable action logs. The enterprise should suspend the agent, attempt recovery, notify affected parties, investigate the control failure, and update its policies. Again, technical exploitation does not erase managerial responsibility for granting excessive authority.

Ethical and Organizational Implications

The first implication is that AI responsibility cannot be assigned to the machine. A system may be causally involved and operationally autonomous, but it does not possess the moral standing required for blame, promise, repentance, or institutional duty. Treating the agent as responsible creates a convenient fiction that protects the humans and organizations that designed the objective, supplied the data, granted the access, and benefited from the activity. Faith-informed ethics strengthens this conclusion by insisting that moral agency includes accountability, intention, and obligation, not merely the capacity to produce complex behavior.⁸

The second implication is that corporate structure must change. Agentic AI is not only an information-technology project. It alters delegation, supervision, internal controls, customer rights, employment

8 Rehman, "Faithful and Responsible Artificial Intelligence" (2026); Rehman, "Religious Hallucinations in Generative AI" (2026); Rehman, "Religious Equity in Public-Sector Artificial Intelligence" (2025); Rehman, "Integrating Islamic Ethics with Artificial Intelligence" (2025); Hayat and Arif, "Islamic Ethical Perspectives on Artificial Intelligence and Digital Surveillance" (2024); Raza and Khalid, "Interfaith Dialogue" (2022); Rasool et al., "Mental Health and Spiritual Well-Being in the Qur'an" (2024).

relations, and enterprise liability. Boards and senior executives should therefore treat high-impact agent deployment as a governance decision. Audit committees, risk officers, and operational leaders need evidence that authority is bounded and incidents can be reconstructed.

The third implication concerns incentives. An enterprise may formally adopt ethical principles while rewarding managers only for speed, cost reduction, or sales. Agents will then be optimized and deployed in ways that reflect those incentives. Accountability requires performance measures for accuracy, complaints, appeals, reversals, disparate impact, privacy, worker burden, and stakeholder trust. Ethical goals must affect budgets, promotion, and operational thresholds.

The fourth implication is that human oversight creates labor, not simply reassurance. Reviewing difficult cases, investigating alerts, and supporting appeals require trained employees. Organizations should not remove staff through automation and then claim that humans remain available for oversight. A credible system must fund the human capacity needed to supervise it.

The fifth implication is that accountability must cross organizational boundaries. Enterprises often rely on foundation-model providers, cloud platforms, consultants, data brokers, and third-party tools. Contracts should specify logging, incident cooperation, model changes, data use, security, evaluation, and termination rights. Still, the deploying enterprise cannot assume that vendor compliance substitutes for its own duty to customers and workers.

Implementation Roadmap for Enterprises

First, create an inventory of every AI system that can communicate externally, invoke tools, modify records, move money, influence employment, or create commitments. Inventory should include shadow deployments and employee-created workflows, not only centrally purchased platforms.

Second, classify each use case under the five autonomy levels and document the potential impact on customers, workers, suppliers, finances, privacy, safety, religious and cultural identity, and public trust. High-impact uses should receive executive and legal review before deployment.

Third, appoint a named human principal and write a delegation charter. The charter should define the agent’s mission, permitted and prohibited actions, data and tool access, transaction limits, escalation triggers, review requirements, and shutdown conditions.

Fourth, implement identity and least-privilege access for every agent and sub-agent. Credentials should be unique, revocable, time-bounded, and monitored. Sensitive actions should require separate authorization or dual control.

Fifth, establish traceability. The enterprise should be able to reconstruct the goal, evidence, policy, plan, tool calls, human approvals, and external effects of a consequential action. Logs should be secure, privacy-conscious, and available for audit and incident response.

Sixth, test the system before deployment using realistic scenarios, including ambiguous requests, conflicting data, prompt injection, manipulation by external parties, religious and cultural edge cases, attempts to exceed authority, and situations requiring abstention or escalation.

When AI Becomes a Business Actor: Closing the Accountability Gap

Seventh, create visible notice and appeal channels. Customers and employees should be able to reach a competent human, correct inaccurate data, and request review of consequential decisions without excessive delay or burden.

Eighth, monitor live performance through incident reports, sampled actions, reversal rates, complaints, bias indicators, financial anomalies, security events, and human-review quality. Governance should be continuous because models, tools, data, and organizational uses change.

Ninth, rehearse shutdown and recovery. The enterprise should know how to revoke credentials, stop workflows, preserve evidence, reverse transactions, notify stakeholders, and continue critical operations without the agent.

Tenth, conduct periodic board-level review of high-impact agentic systems. The review should ask whether the business purpose remains justified, whether autonomy should be reduced, whether human capacity is sufficient, and whether the system continues to serve dignity, fairness, and public benefit.

Recommendations

First, enterprises should recognize agentic AI as a form of delegated organizational authority, not merely as advanced software.

Second, every agentic workflow should have a named human principal with sufficient authority to accept risk, suspend deployment, and provide an account of outcomes.

Third, technical permissions should never be treated as implied business permission. Authority must be explicitly bounded by purpose, context, amount, data, and time.

Fourth, high-impact and irreversible actions should require prior human authorization or remain outside the agent's permitted scope.

Fifth, organizations should implement verifiable agent identity, least-privilege access, segregation of duties, runtime policy checks, and immediate credential revocation.

Sixth, consequential actions should generate reconstructable records that preserve the relevant goal, evidence, policy, reasoning path, tool use, approvals, and effects.

Seventh, customers, employees, applicants, and suppliers should receive appropriate notice and a practical right to human review and correction.

Eighth, enterprises should include religious, cultural, disability, language, and minority-rights considerations in impact assessments when agentic systems operate in diverse markets.

Ninth, contracts with AI vendors should require transparency, security cooperation, incident support, model-change notification, and data-protection commitments.

Tenth, regulators and standards bodies should develop agent-specific guidance for delegated authority, identity, auditability, contestability, and organizational remedy.

Limitations and Future Research

This article presents a conceptual framework and does not empirically validate the proposed controls. Future research should test the framework through case studies in retail, finance, health care,

manufacturing, education, and small business. Comparative studies could examine how firms operationalize human oversight and whether appeal systems produce meaningful correction rather than formal compliance.

Technical research is also needed on reliable agent identity, policy enforcement, action provenance, memory governance, safe delegation to sub-agents, and verifiable shutdown. Organizational research should study the workload and authority of human reviewers, the risk of automation bias, and the effect of agentic deployment on entry-level work and managerial skill development.

Further scholarship should develop multifaceted and culturally inclusive benchmarks for business agents. Systems that operate globally need evaluation methods capable of detecting religious misinformation, stereotyping, exclusion, and manipulation. Such research would extend responsible AI beyond generic fairness metrics toward deeper respect for conscience, community, and plural public life.

Conclusion

Agentic AI changes the structure of business responsibility because it allows software systems to participate directly in organizational action. These systems can create value by coordinating work, accelerating routine decisions, and expanding service. Yet the same capacity creates an accountability gap when authority is delegated without clear ownership, limits, records, oversight, appeal, or remedy. The article has argued that AI can become a functional business actor without becoming a moral or legal person. This distinction allows enterprises to acknowledge the real operational power of agents while refusing the fiction that machines can absorb human responsibility. The proposed Enterprise Agent Accountability Framework connects named ownership, bounded delegation, agent identity, traceability, human review, stakeholder rights, and remediation into a single governance chain.

The governing principle is simple: responsibility must follow authority. Businesses may delegate tasks, discretion, and carefully limited action to AI agents. They may not delegate away their duty to explain, correct, compensate, and learn. The responsible agentic enterprise will therefore be measured not only by how much work its systems can perform, but by whether human beings remain identifiable, answerable, and capable of protecting those affected by automated power.

References

- Abbas, M., M. Anjum, and A. A. Raza. "Divine Oneness and Distinct Theologies." *Islamic Research Journal al-Marsus* 2, no. 4 (2024): 20-28.
- Akhter, N., S. Arif, and S. Akhtar. "Socio-Political and Religious Contributions of Sir Syed Ahmad Khan." *Al-Afaq Islamic Research Journal* 4, no. 1 (2024): 158-165.
- Anwar, F., H. F. Rasool, and M. A. U. Haq. "Global Campaign about Prophetic Blasphemy: Motivations, Causes, and Solutions: A Research Overview." *International Research Journal of Arabic and Islamic Studies* 4, no. 2 (2024): 47-67.
- Arif, Salman. "The Qur'anic Framework for Minority Rights: Upholding Justice, Dignity, and Religious Freedom." *Al-Asr* 5 (January-March 2025).

When AI Becomes a Business Actor: Closing the Accountability Gap

- Arif, S., and U. Ahmad. "Existence of God and Morality: Scholastic Endeavors by Christian and Muslim." *Journal of Religious and Social Studies* 2, no. 2 (2022): 47-66.
- Arif, S., and U. Ahmad. "Scientism and Religion: Christian-Muslim Responses about Atheistic Ideology." *Pakistan Journal of Islamic Philosophy* 4, no. 1 (2022): 53-69.
- Aslam, M. S., M. A. Rehman, and H. F. Rasool. "Effective Strategies for Teaching Arabic to Non-Native Speakers in Pakistan." *Inkishaf* 4, no. 13 (2024): 17-29.
- Atiq, A., and H. F. Rasool. "The Character Building of Individuals by the Teachings of Islam." *Islamic Research Journal al-Qudwah* 3, no. 1 (2025): 96-105.
- Azam, M., H. F. Rasool, and S. Tahseen. "The Reasons for the Decline of Pakistani Society and Its Remedy." *Inkishaf* 3, no. 10 (2023): 190-209.
- Binish, H. M., M. A. Rehman, and H. F. Rasool. "The Importance of Translating the Meanings of the Holy Quran into International Languages." *Inkishaf* 4, no. 12 (2024): 14-32.
- Booth, Harold, William Fisher, Ryan Galluzzo, and Joshua Roberts. *Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization*. Initial Public Draft. Gaithersburg, MD: National Institute of Standards and Technology, 2026.
- Bukhari, S. H. F., M. Abid, and S. Arif. "Atheism in the Muslim World: Classical Approach through Medieval Islamic Intellectual and Philosophical Traditions." *Journal of Religious and Social Studies* 3, no. 2 (2023): 95-114.
- Dux, Nelly, Cristina Alaimo, Philippe Roussiere, and Abhishek Kumar Mishra. "Governance by Design: Architecting Agentic AI for Organizational Learning and Scalable Autonomy." *arXiv preprint arXiv:2605.20210* (2026).
- European Parliament and Council. *Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*. Official Journal of the European Union, 2024.
- Hayat, I., and S. Arif. "Islamic Ecological Ethics and Sustainable Development in Pakistan: A Qur'anic Perspective on Climate Challenges." *International Research Journal of Arts, Humanities and Social Sciences* 2, no. 2 (2025).
- Hayat, I., and S. Arif. "Islamic Ethical Perspectives on Artificial Intelligence and Digital Surveillance: Principles, Challenges, and Policy Recommendations." *International Research Journal of Arts, Humanities and Social Sciences* 2, no. 1 (2024).
- Hanif, M. Z., and H. F. Rasool. "Exploring New Educational Dimensions for the Improvement of Pakistan's Education System." *Pakistan Islamicus* 5 (2025).
- Hussain, S., M. Muslim, and S. Arif. "Conflict Resolution and Development of Skilled Leadership: Reconciliation between Muslim and Western Approaches." *Journal of Religious and Social Studies* 3, no. 2 (2023): 20-41.
- Jasvi, M. A., Z. Rasheed, A. A. Raza, and H. F. Rasool. "Antisocial Activities and Role of a Religion in a Society: A Descriptive Research." *Pakistan Islamicus* 4 (2024): 11-18.

- Kasirzadeh, Atoosa, and Iason Gabriel. "Characterizing AI Agents for Alignment and Governance." arXiv preprint arXiv:2504.21848 (2025).
- Kaul, Anuj, Qianlong Lan, and Pranay Gupta. "AgentBound: Verifiable Behavioral Governance for Autonomous AI Agents." arXiv preprint arXiv:2606.30970 (2026).
- Khan, Rafflesia, Declan Joyce, and Mansura Habiba. "AGENTS SAFE: A Unified Framework for Ethical Assurance and Governance in Agentic AI." arXiv preprint arXiv:2512.03180 (2025).
- National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. Gaithersburg, MD: NIST, 2023.
<https://doi.org/10.6028/NIST.AI.100-1>.
- National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1. Gaithersburg, MD: NIST, 2024.
<https://doi.org/10.6028/NIST.AI.600-1>.
- National Institute of Standards and Technology. "AI Agent Standards Initiative for Interoperable and Secure Innovation." Gaithersburg, MD: NIST, 2026.
- Organisation for Economic Co-operation and Development. Advancing Accountability in AI: Governing and Managing Risks throughout the Lifecycle for Trustworthy AI. OECD Digital Economy Papers. Paris: OECD Publishing, 2023.
- Rasool, H. F., A. Aziz, and M. Kiran. "Mental Health and Spiritual Well-Being in the Qur'an: Addressing Modern Psychological Challenges through Islamic Teachings." *Ulum Al-Sunnah* 2, no. 2 (2024): 33-40.
- Rasool, H. F., A. Aziz, H. M. Usman, and M. Kiran. "Economic Justice in Islam: A Comprehensive Qur'anic Framework for Equity and Social Harmony in Seerah Perspective." *Tanazur* 5, no. 4(a) (2024): 1-15.
- Rasool, H. F., S. M. Shah, and M. Nasrullah. "Islamic Responses to Environmental Protection and Sustainability: Principles and Contemporary Applications." *Islamic Research Journal al-Qudwah* 2, no. 4 (2024): 78-85.
- Rehman, Ataur. "Faithful and Responsible Artificial Intelligence: An Ethical Framework for Education, Faith Communities, and Public Life." *Islamic Research Journal al-Qudwah* 4, no. 2 (2026): 1-17.
- Rehman, Ataur. "Free Will: A Discourse between New Atheism and Islam." *Islamic Research Journal al-Marsus* 1, no. 4 (2023): 1-24.
- Rehman, Ataur. "From Contingency to Design: Reassessing al-Ghazali's Arguments for God in Philosophy, Theology, and Contemporary Science." *Islamiyyat* 48, no. 1 (2026): 24-34.
<https://doi.org/10.17576/islamiyyat-2026-4801-03>.
- Rehman, Ataur. "God's Existence in Primary Islamic Sources: An Analytical Study in View of New Atheism." *Journal of Positive School Psychology* 6, no. 9 (2022): 26-33.

When AI Becomes a Business Actor: Closing the Accountability Gap

- Rehman, Ataur. "Iconoclasm: Re-imagining Sacred Art and Morality in Contemporary Spiritualities." In *Spirituality and the Arts*. Harvard Divinity School, Program for the Evolution of Spirituality, 2025.
- Rehman, Ataur. "Integrating Islamic Ethics with Artificial Intelligence: A Study of Seerah in the Context of Pakistan." *Islamic Research Journal al-Marsus* 3, no. 2 (2025): 1-9.
- Rehman, Ataur. "Morality: A Contemporary Discourse between New Atheism and Islam." *Islamiyyat* 42, no. 2 (2020): 113-124.
- Rehman, Ataur. "Religious Equity in Public-Sector Artificial Intelligence: Developing a Multi-Faith AI Impact Assessment Framework through a Comparative Study of Texas and California." *Pakistan Islamicus* 5, no. 2 (2025).
- Rehman, Ataur. "Religious Hallucinations in Generative AI: Developing a Multifaith Risk Benchmark to Prevent Misinformation, Polarization, and Extremist Manipulation in the United States." *Islamic Research Journal al-Marsus* 4, no. 2 (2026): 38-56.
- Rehman, Ataur, and T. Basharat. "God's Existence: Philosophical Discourse between New Atheists and Muslim Scholars." *Pakistan Journal of Islamic Philosophy* 3, no. 2 (2021): 1-16.
- Rehman, M. S. U., and A. A. Raza. "The Development and Growth of Gratitude and Positive Thinking in the Light of Artificial Intelligence and Seerah." *Islamic Research Journal al-Marsus* 3, no. 2 (2025): 66-79.
- Riggs, Jared, Maia Hamin, Neil Perry, Benjamin Edelman, and Peter Cihon. *Summary Analysis of Responses to the Request for Information Regarding Security Considerations for AI Agents*. NIST Trustworthy and Responsible AI 800-5. Gaithersburg, MD: NIST, 2026.
- Raza, A. A., D. Ali, and Ataur Rehman. "Social Importance of Shyness and Modesty in Islam: An Analytical Study in the Light of Prophet's Seerah." *Al-Durar* 3, no. 3 (2023).
- Raza, A. A., H. F. Rasool, and Ataur Rehman. "The Role of Mosques for the Social Unity in the Light of Sirat-e-Taiba." *Al Meezan Research Journal of Social Sciences* 3, no. 1 (2023): 37-50.
- Raza, A. A., H. U. Ismail, and Q. A. Manan. "The Social Importance and Requirements of Tolerance: Analytical Study in the Light of the Prophet's Seerah." *Al Manhal Research Journal* 3, no. 3 (2023).
- Raza, A. A., and M. S. Khalid. "Interfaith Dialogue: Ethical Commonalities in Judaism, Christianity and Islam." *Islamic Studies Research Journal Abhath* 7, no. 26 (2022).
- Raza, A. A., W. Ali, and G. D. Ahmad. "Social Importance and Requirements of Compassion and Gentleness." *Al-Durar* 4, no. 1 (2024).
- Saeed, A. F. I., A. A. Raza, and H. A. Fatima. "Prayer as the Heart of Worship: Exploring Shared Rituals and Spiritual Connections in Judaism, Christianity and Islam." *Islamic Research Journal al-Qudwah* 3, no. 1 (2025): 1-8.
- Saeed, A. F. I., A. A. Raza, and K. tul Kubra. "Prophethood in the Abrahamic Faiths." *Islamic Research Journal al-Qudwah* 2, no. 4 (2024): 126-133.

Saeed, A. F. I., A. A. Raza, S. Arif, and H. F. Rasool. "Charity across Faiths: The Ethical and Spiritual Dimensions of Giving in Judaism, Christianity, and Islam." *Al-Aasar* 2, no. 1 (2025): 748-757.

Saeed, A. F. I., H. A. Fatima, and A. A. Raza. "The Sacred Texts of Abrahamic Faiths: Common Themes in the Torah, Bible, and Qur'an." *Harf-o-Sukhan* 8, no. 3 (2024): 1085-1094.

UNESCO. *Recommendation on the Ethics of Artificial Intelligence*. Paris: UNESCO, 2021.

Zahid, Z., R. Amin, I. Khalil, B. A. K. Mohammed, and S. Arif. "Regulating Digital Currencies in the EU: A Comparative Analysis with Islamic Finance Principles under MiCA." *International Journal of Business and Management Practices* 3, no. 3 (2025).