

The AI-Enabled Finance Function:

Turning AI potential into value through trust, control and professional judgement

This article examines how finance functions can turn the potential of AI into sustainable value by building trust in AI-enabled information, applying proportionate controls and governance, and ensuring that professional judgement and accountability remain central.

Executive Summary

Artificial intelligence (AI), particularly generative AI, is reshaping how organizations operate, manage risk and make decisions. For finance functions, the opportunity is not limited to faster processing: AI can automate components of routine work, accelerate analysis and strengthen forward-looking insight, enabling finance teams to devote more attention to interpretation, challenge and business partnering. However, capturing the full potential of AI often requires organizations to reconsider processes rather than simply add AI to existing workflows, while preserving clear ownership of the conclusions and decisions that follow. Trust is essential to turning AI capability into sustainable value. Generative AI may produce inaccurate calculations, incomplete analysis or plausible but unsubstantiated content; it can also create risks involving confidentiality, privacy, cybersecurity, intellectual property and bias. These risks are most significant where AI affects financial or sustainability reporting, regulatory compliance, forecasts or material business decisions.

Organizations should therefore adopt proportionate, risk-based governance. Lower-risk applications may be supported by clear guidelines and proportionate review, while higher-risk applications require stronger validation against reliable evidence, appropriate documentation, competent human oversight and explicit accountability. A value-risk assessment can help finance leaders prioritize applications without either constraining useful experimentation or accepting unmanaged risk.

Accountants remain essential in an AI-enabled finance function because the technology changes where professional judgement is exercised rather than removing the need for it. Their roles increasingly converge: as users, they apply AI to improve finance activities; as guardians, they protect the integrity of information and controls; and as business partners, they help management evaluate the value, risks and implications of AI-enabled change.

Finance leaders should begin by understanding where AI is already being used and assessing applications according to their value and risk. They should then establish clear protocols for approved tools, data use and review, supported by defined ownership and accountability. Together with investment in AI literacy and professional capability, these actions can provide a practical foundation for responsible adoption and sustained value creation.

1. How AI is reshaping the finance function

Across organizations, AI is changing how work is performed, how products and services are developed and delivered, and how decisions are made. Within this broader transformation, the finance function is a natural focal point because it connects transaction data, performance analysis, reporting, planning and resource allocation.


For finance teams, the potential applications are wide-ranging. AI can support transaction processing, reconciliations, document review and the preparation of management reports. It can also assist with variance analysis, forecasting, scenario modelling, regulatory research, and financial and sustainability reporting by organizing information, identifying relevant requirements and reviewing large volumes of supporting documentation.

Not all of these applications rely on generative AI. Forecasting, anomaly detection and some forms of reconciliation often use predictive or rules-based systems, which raise different issues (such as model drift or weak training data) from the hallucination risk associated with generative models. The controls discussed later should be matched to the type of system in use.

The immediate attraction of these applications is often productivity. Tasks that previously required significant manual effort may be completed more quickly, allowing finance teams to respond faster to business needs. However, the potential value of AI extends beyond simply reducing processing time. A more significant opportunity may be the ability to shift finance capacity away from repetitive activities towards work that requires interpretation, judgement and engagement with the wider business. Transaction integrity, control design and the quality of reporting remain core finance responsibilities. AI may change the mix of work, but it does not move those duties out of finance.

This reflects a broader evolution in the role of finance. Traditionally, finance functions have been heavily involved in recording transactions, maintaining financial information and preparing reports. Over time, many organizations have sought to move finance further towards analysis and business partnering. AI may accelerate this development by taking on parts of the information-processing workload and enabling finance professionals to spend more time understanding what the information means and how it should influence decisions.

These opportunities can be viewed as a progression from lower- to higher-order uses of AI:



- **Efficiency (foundational use)** focuses on automating or streamlining repetitive, time-consuming activities, such as data extraction, reconciliation, document review and preparation of initial drafts.
- **Insight (analytical use)** uses AI to analyse information more rapidly and surface trends, relationships or anomalies for professional investigation, including variances and alternative assumptions or scenarios.
- **Decision support (higher-order use)** brings AI-generated analysis into management discussions. At this level, finance professionals add value by interpreting results, challenging assumptions and explaining the financial implications of alternative courses of action.

The extent to which organizations capture these benefits will depend on how AI is integrated into existing processes. **Using AI simply to replicate an inefficient process more quickly may generate some productivity gains, but it may not deliver the full value of the technology.** Greater benefits may arise where organizations reconsider how work is performed, which activities require human involvement and where professional expertise adds the greatest value.

AI adoption is therefore not solely a technology initiative. Choices about where and how it is used involve business objectives, information needs, operating processes, controls and risk appetite. Finance professionals should bring their understanding of these matters to cross-functional decisions with technology, risk, legal and business teams, particularly where an application affects reporting, resource allocation or business performance.


Greater reliance on AI also raises a corresponding requirement for confidence in the results it produces. The central challenge is to realize the benefits of AI without eroding trust in financial analysis, reporting and organizational accountability. The next section examines the conditions that allow users to place appropriate reliance on AI-enabled financial information.

2. The trust challenge: when AI meets financial information

As AI becomes more integrated into finance processes, the question is no longer only what the technology can do, but also whether its outputs can be trusted. This is particularly important in finance, where information may influence management decisions, external reporting, resource allocation and other outcomes with significant consequences.

Generative AI, a widely used form of AI that creates content based on user prompts, can produce outputs that are inaccurate, incomplete, biased or unsubstantiated, and it can introduce data privacy, cybersecurity and intellectual property risks. **These risks do not necessarily preclude the use of AI, but they reinforce the need for controls suited to the purpose and significance of each application.**

To address these concerns in practice, finance functions can consider seven dimensions that help enhance the trustworthiness of AI-enabled financial information, drawing on established guidance including the NIST AI Risk Management Framework and OECD AI Principles:



- **Accuracy**

Generative AI can produce outputs that appear coherent and convincing but are not necessarily factually correct. This is commonly associated with the risk of “hallucination”, where an AI system generates information that is inaccurate or fabricated. In a finance context, this risk can arise in several ways. A general-purpose language model may perform arithmetic inconsistently, transpose figures or apply an incorrect formula. **Material calculations should therefore be performed or independently reproduced in a controlled calculation environment, rather than left to unverified generative output, and independently checked.** A model may also cite a non-existent requirement or misinterpret financial information. Because the output can appear fluent and confident, errors may be difficult to identify without comparison to underlying data and authoritative sources.

- **Completeness**

Trust is not only about whether the information presented is correct. It also depends on whether relevant information has been omitted. AI systems generate outputs based on the information available to them and the way in which questions are framed. They may therefore overlook relevant context, assumptions or factors that a finance professional would consider important. An answer may be technically accurate but still incomplete or misleading if material information has been excluded. **Finance professionals therefore need to consider not only “Is this correct?” but also “Is the information sufficiently complete for the purpose for which it will be used?”**

- **Traceability and explainability**

For financial information, users may need to understand how a conclusion was reached, including the sources, assumptions and reasoning that support it. This can be difficult when an AI system operates as a “black box”, with internal processes that are not readily interpretable, or when it cannot provide verifiable sources for a generated response. This is especially relevant for information that may be subject to review, challenge or assurance. **Finance functions should require sufficient traceability to authoritative source information and an explanation that enables a suitably skilled reviewer to understand and validate the conclusion.**

- **Confidentiality**

These risks can arise where users enter confidential information into AI tools without understanding how that information may be stored, processed, used or disclosed. They may also arise where contractual terms do not adequately address rights in input data, training data or generated content. **Organizations therefore need clear parameters for the information that may be provided to AI systems, the tools approved for use and the applicable intellectual property requirements.** For finance professionals, this means considering not only the usefulness of an AI tool but also whether the information being provided to it is appropriate and adequately protected. Organizations should also consider vendor and third-party arrangements (including training use of inputs, data residency and rights in outputs) and discourage the informal use of public or unapproved tools for finance information.

- **Fairness**

AI systems may reproduce or amplify bias in their data, design or use. In finance, this can affect how information is interpreted, which explanations are surfaced and how recommendations influence decisions. **Proportionate review should consider whether inputs, assumptions or outputs produce unjustified differences across relevant groups, transactions or decisions.** Material concerns should be investigated and escalated rather than treated as a purely technical issue.

- **Security and resilience**

AI tools can expand the attack surface through new data flows, system integrations and dependencies on third parties. The risks are greater where a system connects to finance applications or can initiate actions rather than only

generate content. **Controls may include secure configuration, role-based access, segregation of duties, monitoring, incident response and contingency arrangements, calibrated to the use case and its potential consequences.**

- **Accountability**

Accountability underpins the other dimensions by making clear who is responsible for putting them into practice. AI may generate an output, but responsibility for how that output is reviewed, communicated and used remains with the organization and the people making the decision. For finance functions, this means that AI should not become a substitute for ownership of decisions. Where AI contributes to a financial analysis, forecast, disclosure or recommendation, **there should be clarity over who reviews the output, who determines whether it can be relied upon and who is accountable for the final decision.**

Addressing these dimensions enables finance functions to use AI with greater confidence and at greater scale. Trust should not, however, be treated as an inherent attribute of an AI tool. The degree of reliance placed on it should be calibrated to the purpose, inputs and context of the specific use, the evidence supporting its output, the controls in place and the consequences of error or misuse.

The next section considers how organizations can translate these principles into practical controls and governance while maintaining sufficient flexibility to support responsible innovation.

3. Controls and governance: matching oversight to risk

As organizations expand their use of AI, effective governance should support responsible use by providing a structured approach to determine how different applications should be approved, controlled and monitored. This does not mean subjecting every AI application to the same level of review. Instead, **oversight should be proportionate to the nature of the use case, the degree of system autonomy and the potential consequences of error, misuse or uncontrolled action.**

This is particularly relevant to finance functions, where significance varies according to how an output will be used. Improving the wording of a routine internal email using non-sensitive information is unlikely to require the same oversight as supporting a financial forecast, interpreting an accounting requirement or preparing information for external reporting. A proportionate approach can help provide a clear path for lower-risk applications while directing greater attention to applications that could have more significant financial, regulatory or reputational consequences.

Assessing AI use cases according to risk

A practical starting point is to assess an AI use case by considering its purpose, the information involved, the degree of autonomy or system access and the consequences of error or misuse.

Relevant considerations may include:

- the intended purpose and use of the AI application;
- the type of data or information provided to the system;
- the users of, and parties relying on, the output;
- the potential impact on financial reporting, regulatory compliance or material business decisions;
- the potential consequences of inaccurate, incomplete or misused outputs;
- the extent and appropriateness of human review; and
- the degree of autonomy and access to systems, data or transactions.

For finance leaders, the key principle is that the level of oversight should increase with the potential consequence of the AI use case.

Combining value and risk through practical tiers


A practical tiering approach can help organizations translate the assessment above into proportionate governance. This could be achieved by categorizing each use case in the tier illustrated below that best reflects the level of oversight it requires, also considering the organization's activities, risk appetite and regulatory obligations.

Classification should reflect the specific implementation and potential consequence rather than the activity label alone, and the same use case may fall into different tiers in different contexts.

- **Tier 1 – lower-risk applications.** Examples include improving routine internal communications, summarizing non-sensitive meeting notes or generating generic templates. Approved tools, data-use parameters and user review may provide an appropriate control baseline.

- **Tier 2 – moderate-risk applications.** Examples include management reporting, variance analysis or scenario modelling that informs internal discussion. These uses may require documented sources and assumptions, validation against underlying data, review by a suitably skilled person and periodic monitoring.
- **Tier 3 – higher-risk applications.** Examples include external reporting, regulatory interpretation, material forecasts or significant management decisions. These uses may require independent verification, an audit trail, senior approval and ongoing monitoring. For a specific higher-risk use case, AI may be unsuitable where adequate evidence, explainability or control cannot be achieved. Any such determination should be made on a case-by-case basis and should not be interpreted as a general restriction on the use of AI in reporting or forecasting.

Where AI contributes to external reporting, sustainability reporting, regulatory interpretation or material forecasts, management remains responsible for the information. The use of AI does not reduce the evidence, disclosure or assurance expectations that already apply.




Business value, such as time saved, improved insight or better decision support, should then be considered within each tier. High-value, lower-risk applications may be suitable for faster experimentation; high-value, higher-risk applications may justify investment in stronger controls; and low-value, higher-risk applications may warrant reconsideration.

Controls across the AI-enabled process

Once a use case has been classified, controls should follow it throughout its lifecycle—from initial approval and deployment to ongoing monitoring and, where necessary, suspension or retirement. Clear ownership and documentation should connect these stages so that the control framework operates as a coherent system rather than a series of isolated checkpoints.

This lifecycle-based approach is consistent with two widely recognized sources of guidance: COSO’s guidance on applying internal control principles to generative AI

and the NIST Generative AI Profile on managing generative AI risks across the AI lifecycle. The following illustrate how this approach can be translated into practical controls that build on established financial controls while extending them to cover AI use cases, systems and data, output validation and professional review, documentation and accountability, and ongoing monitoring and change management.



- **Use-case inventory and ownership**

Organizations should maintain an inventory of material or approved AI use cases, including their purpose, owner, tool, data, risk tier and required review. The inventory should be updated when the model, integration, data or intended use changes.

A centralized inventory gives management visibility over the organisation's use of AI and helps identify applications that lack clear ownership or appropriate oversight.

- **AI systems and access**

Organizations should establish which AI tools are approved for use, who may access them and whether their capabilities, security, contractual terms, data-handling arrangements and change practices are suitable for the intended purpose.

For systems that can initiate actions or access finance applications, permissions should be granted on an as-needed basis for the approved use, with segregation of duties, transaction limits, approval gates and activity logging appropriate to the risk. Clear parameters can also reduce the use of unapproved applications and the entry of sensitive information into systems that have not been assessed.

- **Inputs**

Organizations should determine what information may be entered into AI systems. This is particularly important where financial, personal, commercially sensitive or confidential information is involved.

Finance professionals should understand whether the information provided to an AI tool is appropriate for that system and whether it is being handled in accordance with the organization's requirements.

- **Output validation and professional review**

AI-generated outputs should be validated according to their intended use.

Review may include reconciling figures to underlying data, reproducing calculations in an appropriate tool, checking cited requirements against authoritative sources, testing assumptions and considering whether relevant information has been omitted. The evidence retained should be sufficient for another suitably skilled person to understand the basis of a material conclusion.

Human review should not be a formality. For finance-related applications, the reviewer needs sufficient professional competence, context, independence of mind and authority to identify errors and prevent inappropriate reliance, consistent with the fundamental principles in the IESBA Code. The rigor and independence of review should increase with the potential risk-level consequence of the use case.

- **Documentation and accountability**

For moderate- and higher-risk applications, **organizations should document the use of AI, key sources and assumptions, validation procedures, material exceptions and approvals.** Documentation should be sufficient for another

suitably skilled person to understand how the conclusion was reached and challenged.

Clear accountability helps ensure that responsibility for decisions does not become obscured simply because AI was involved in producing the underlying analysis. The accountable individual should be clearly identified at an appropriate level of seniority, with key responsibilities and approvals documented and available for review where appropriate.

- **Monitoring and change management**

Approved applications should be monitored for performance, control failures and system changes and uses that extend beyond the original approval. Organizations should also define how errors, suspected data exposure, inappropriate outputs and unauthorized actions are reported, contained and escalated, including when an application should be suspended or reassessed.

The frequency and depth of monitoring should increase with the consequence of the use case, with more frequent or automated monitoring considered where appropriate. System enhancements and control adjustments should be made in a timely manner as risks, technology or use cases evolve.

Taken together, these controls can provide a clearer framework for responsible AI use by clarifying which tools and uses are approved, when additional review is required, and who owns important decisions, while ensuring that stronger oversight is applied where the potential consequences are more significant.

Controls are effective only when the people applying them have the competence and ability to exercise judgement. For professional accountants, this human element is reinforced by ethical responsibilities that continue to apply when technology is used. The next section therefore considers how professional judgement may evolve in an AI-enabled finance function.

4. Professional judgement in an AI-enabled finance function

As AI becomes more capable of generating analysis, explanations and recommendations, professional judgement becomes increasingly important. AI can support the work of finance professionals, but it cannot remove the need to assess whether information is appropriate, sufficiently supported and suitable for the purpose for which it will be used.

For accountants, this means that AI should be viewed as a tool that can inform judgement rather than replace it.

Ethical responsibilities in a technology-enabled environment

The IESBA International Code of Ethics for Professional Accountants is directly relevant to the use of AI. Its recent technology-related revisions clarify that the use of technology does not diminish the need for professional judgement, competence, due care, confidentiality and ethical decision-making. Professional accountants therefore remain responsible for applying the Code's fundamental principles and conceptual framework when using, evaluating or relying on technology-enabled outputs, including those generated by AI. In practice, these principles have several practical implications:

- **Professional competence and due care.** Accountants need sufficient understanding of an AI tool's capabilities, limitations, assumptions and data requirements to evaluate whether its output is fit for purpose. Competence also requires keeping pace with relevant technological developments and involving specialists where necessary.
- **Objectivity and an inquiring mind.** Accountants should guard against automation bias and undue reliance on technology. Applying an inquiring mind means assessing the source and relevance of inputs, the sufficiency of outputs and whether contradictory evidence or alternative explanations have been considered. In audit and assurance contexts, this also supports the exercise of professional skepticism required by the relevant standards
- **Confidentiality.** The use of AI does not relax the duty to protect confidential information. Accountants should not place proprietary, client, customer, employee or other protected data into public or unapproved tools without proper authority and appropriate safeguards.
- **Integrity and professional behaviour.** AI-assisted work should not obscure errors, unsupported assertions or the extent of human involvement. Accountants remain responsible for communicating information honestly, complying with applicable requirements and avoiding association with misleading information.

These principles turn human oversight from a procedural checkpoint into professional responsibility: While the same professional responsibilities apply to non-AI processes, the apparent authority and limited transparency of AI-generated outputs can increase the risk of undue reliance. These principles therefore reinforce that human oversight should involve informed challenge and accountability, rather than serve merely as a procedural checkpoint.


From human-in-the-loop to professional-in-the-loop

Many approaches to responsible AI emphasize the importance of keeping a “human in the loop”. Human review is an important safeguard, but for consequential finance activities, the mere presence of a reviewer may not be sufficient.

What matters is whether the person reviewing the output has the appropriate professional competence, knowledge and ability to challenge it. For finance information, that will often mean a suitably skilled finance professional; it does not mean that only accountants can review AI output in every function.

For finance-related applications, effective oversight may therefore require a **professional-in-the-loop**: someone who can assess the output not only for obvious errors, but also in the context of relevant accounting requirements, business circumstances, underlying evidence and the intended use of the information.

Beyond the ethical principles above, professional-in-the-loop oversight adds three practical requirements:




- **Contextual understanding.** A reviewer needs sufficient accounting, business and AI-related knowledge to determine whether the output makes sense in the relevant commercial, organizational and regulatory circumstances.
- **Ability to challenge.** Effective oversight requires testing assumptions, considering contradictory evidence, seeking additional support and communicating limitations clearly, together with the ability to prevent inappropriate reliance.

- **Explicit accountability.** Responsibility for the conclusion or decision should remain clear; the use of AI should not blur ownership of the resulting analysis, communication or action.

The evolving role of the accountant

These developments reinforce a broader evolution in the role of accountants. As AI takes on more information-processing activities, accountants may increasingly contribute in three interconnected ways.



As users, accountants can apply AI to improve efficiency, accelerate analysis and support the delivery of finance activities; **As guardians**, they can help maintain the integrity of information, challenge unreliable outputs and contribute to effective internal controls over AI-enabled processes; **As business partners**, they can help management evaluate where AI can create value, how risks should be managed and how AI-enabled insights should inform business decisions.

These roles are not separate. In practice, an accountant may use AI to generate an initial analysis, act as a guardian by challenging the reliability of the result, and then serve as a business partner by helping management interpret its implications.

This combination of technological capability and professional judgement may become increasingly important as AI moves beyond routine productivity applications into areas that influence financial reporting, planning and strategic decision-making.

Capabilities for an AI-enabled finance function

The profession's comparative advantage will increasingly lie not in producing information manually, but in interpreting AI-enabled information by combining technological fluency with accounting expertise, ethical reasoning and understanding of the business. Finance leaders should therefore develop a balanced capability set: AI and data literacy; technical accounting and reporting knowledge; commercial understanding; the ability to design and evaluate controls; and communication skills that make assumptions, uncertainty and limitations clear.

Capability building should be practical and role-based. Professionals need experience in selecting suitable use cases, deciding what information may be entered into a tool, validating outputs, escalating concerns and documenting how important conclusions were reached—not the skills to build AI models themselves.

As organizations move further towards AI-enabled finance, building these capabilities will be an important part of responsible adoption. The final section considers the practical questions finance leaders need to ask as they seek to capture the benefits of AI while maintaining appropriate trust and oversight.

5. Moving forward: questions for finance leaders

As organizations move from experimenting with AI towards more widespread adoption, finance leaders will need to balance two objectives: capturing the potential benefits of AI and maintaining appropriate trust, control and accountability.

The following questions may help finance leaders assess whether their organizations are positioned to adopt AI responsibly and effectively.

- **Where is AI already being used across the organization and within the finance function?**

AI adoption may develop through both formal initiatives and individual use of widely available tools. Finance leaders should therefore understand which applications are being used, for what purposes and what types of information are being provided to or generated by these systems.

This can help identify areas where additional guidance, controls or oversight may be required.

- **Which AI applications have the greatest potential to create value?**

Organizations should look beyond adopting AI simply because the technology is available. The focus should be on use cases that address genuine business needs, whether through greater efficiency, improved insight or better decision support.

The potential value should also be considered alongside the investment, organizational change and governance required to implement the application effectively.

- **Which uses of AI could materially affect financial information or business decisions?**

Classify applications by their potential effect on reporting, compliance, forecasts, resource allocation and significant decisions. This classification should determine the required evidence, validation, documentation, approval and monitoring, rather than relying on a generic distinction between drafting and decision support.

- **Are appropriate controls and accountability in place?**

Finance leaders should consider whether there are clear expectations around approved AI tools, the information that may be entered into them, how outputs should be validated, how changes and incidents will be monitored and when additional review or approval is required.

For more consequential applications, there should also be clarity over who is responsible for reviewing the output and who remains accountable for the resulting decision.

- **Do finance professionals have the capability to both use and challenge AI?**

Effective adoption requires more than access to technology. Finance professionals need sufficient AI literacy to understand the capabilities and

limitations of the tools they use, together with the professional competence and business understanding necessary to evaluate their outputs.

As AI increasingly supports the production and analysis of information, the ability to challenge assumptions, identify gaps and determine whether an output can appropriately be relied upon will become increasingly important.

Looking ahead, the development of AI will continue to reshape how finance activities are performed and how organizations use information to make decisions. There is unlikely to be a single model for adoption that is appropriate for every organization. The nature of the opportunities, risks and controls required will depend on the organization, the technology used and the significance of each application.

Finance's role extends beyond simply adopting AI tools. Finance professionals can help determine where AI creates genuine value, what evidence is needed to rely on its outputs and who remains accountable when AI influences a decision. Accountants can help organizations navigate the next stage of AI-enabled transformation by applying professional judgement, proportionate governance and clear accountability to the pursuit of efficiency and innovation.

Selected references

Aldasoro, I., Gambacorta, L., Korinek, A., Shreeti, V. and Stein, M. (2025), “[Intelligent financial system: How AI is transforming finance](#),” *Journal of Financial Stability*, vol. 81, article 101472.

Chartered Accountants Worldwide (CAW) (2025), [AI and the Future of the Global Chartered Accountancy Profession](#).

Committee of Sponsoring Organizations of the Treadway Commission (COSO) (2026), [Achieving Effective Internal Control Over Generative AI](#).

International Ethics Standards Board for Accountants (IESBA) (2023), [Final Pronouncement: Technology-related Revisions to the Code](#)

International Ethics Standards Board for Accountants (IESBA) (2025), [2025 Handbook of the International Code of Ethics for Professional Accountants \(including International Independence Standards\)](#).

International Organization for Standardization and International Electrotechnical Commission (ISO/IEC) (2023), [ISO/IEC 42001:2023 Information technology—Artificial intelligence—Management system](#).

Mirabile, M., Corazza, G. E. and Alonso-Moral, J. M. (2026), “[Trust in human–AI collaboration in finance: A bibliometric–systematic literature review](#),” *AI & Society*, vol. 41, pp. 7625–7654.

National Institute of Standards and Technology (NIST) (2024), [Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile](#), NIST AI 600-1.

National Institute of Standards and Technology (NIST) (2023), [Artificial Intelligence Risk Management Framework \(AI RMF 1.0\)](#), NIST AI 100-1.

McKinsey & Company (2025), [How Finance Teams Are Putting AI to Work Today](#).

Organisation for Economic Co-operation and Development (OECD) (2019; updated 2024), [OECD AI Principles](#).