

The age of the agents

The underlying impacts of agentic AI



Contents

3	Executive Summary
4	Introduction How agentic AI is rewriting the rules of data and business
5	Impact 1 Putting Agentic AI into practice: evolving digital twins
9	Impact 2 Data governance in the Agentic Age
13	Impact 3 Agent-ready data: unearthing the hidden cracks
17	Impact 4 A new type of ROI: how trust earnings pay dividends for ethical AI
21	Impact 5 Human data and AI talent in an agentic world
24	Concluding thoughts
25	Reflection in practice
26	References

Executive summary

The weight of this emerging technology is revealing cracks in the foundations on which agentic AI stands: our people, processes, and technology. The risks, rewards, and uncertainties of this technology are so profound that it has the world's attention, as we wait to see exactly how to harness it – and what its impact will be.

Success demands a holistic strategy that creates a secure, trusted environment to scale experimentation into innovation. Companies that move early will unlock measurable upside: faster product-to-shelf, stronger customer relationships, and a workforce ready for the next chapter of AI-driven growth.

This paper explores five of the critical impacts of agentic AI on our operations, our data, and society:



Agentic digital twins:

From static simulations to dynamic self-learning, these systems show the art-of-the-possible for agents to optimize supply chains, packaging, and logistics in real time.



Data governance in the agentic age:

Governance must evolve into intelligent, automated frameworks to keep pace with continuous agent-driven data flows.



Agent-ready data:

Trusted, permissioned, and certified data is the foundation for safe autonomy, reducing risk and enabling faster innovation.



Trust earnings as ROI:

In reaction to agents, ethical AI and transparency now drive measurable revenue and loyalty, making trust a strategic asset.



Human data and AI talent:

Agents reshape entry-level roles, creating long-term succession risks. Organizations must act now to secure future capability.

Introduction

How agentic AI is rewriting the rules of data and business

We've reached a global inflection point as technology and business converge. Data is more than a source of business insight; it is intrinsic to how business systems and products operate. The emergence of agentic AI is poised to merge these worlds beyond separation.

Unlike traditional models, agents don't just predict; they act. They follow triggers, make decisions, and orchestrate workflows, creating a new paradigm for how organizations engage with intelligence at scale. For the first time, employees can conceptualize and interact with AI in a way that feels tangible; AI models are immaterial, AI agents are our colleagues.

In commercial models, this shift is profound. From marketing personalization and supply chain agility to sustainability tracking and data monetization, agentic AI transforms every link in the value chain. For most organizations, experimentation is already afoot.

Inside technical teams, discussion too often focuses on the practicalities of AI deployment and adoption:

frameworks to advance POCs to MVP, perhaps considering challenges like AI literacy for adoption or broader questions of governance for 'human-in-the-loop' oversight and responsibilities. But the movement towards agent-centric AI raises deeper, yet unanswered questions about AI: What can these systems really do for us? How do we trust them? Who are we in an agentic enterprise?

This paper explores five underlying impacts that agents bring to operations, the data they rely on, and to us – the people behind it all. These five impacts are real-world challenges observed – and tackled – by Kubrick's technical leadership and SMEs working with clients and technology partners as they drive towards an agentic future.

We invite you to reflect on your own goals and maturity, as well as the risks and rewards at stake. With roadmaps, use cases, and measurements, this paper is designed to help benchmark current capability and refine priorities in order to build an ecosystem where human talent and AI agents collaborate to drive value.

Impact 1

Putting Agentic AI into practice with digital twins

Putting Agentic AI into practice with digital twins

Preliminary agentic workflows tend to focus on repetitive tasks with low-level decision-making, such as invoice processing or marketing automations. This is the breeding ground to experiment with agentic autonomy in multi-step processes, where outputs are easily validated by humans. However, the value is limited to a matter of personal productivity. The systems where agents will scale impact will be the processes that are already augmented by other forms of AI.

Take the digital twin: virtualizations of systems and products that model their lifecycle, combining data and AI. In global organizations, which span every part of the value chain from raw materials and manufacturing to last-mile delivery and in-home use, conventional digital twins still struggle to keep pace with this level of complexity. They can model static simulations and what-if dashboards that explain the past but not necessarily anticipate what's next.

Now consider the agentic digital twin: a self-learning environment where AI agents don't just simulate operations but act, adapt, and optimize continuously across production, logistics, and consumer touchpoints.

Agentic digital twin in action: Use case



Supply chain complexity & volatility

Digital twins already help with scenario planning for the many disruptions facing supply chains (supplier delays, regulation changes, port congestion, weather etc.), but they remain read-only. With predictive AI, digital twins become prescriptive systems that move from heuristic planning to dynamic optimization, but agentic AI can take autonomy further. When the twin detects risk, e.g. port delay or raw material shortage, the agent can propose or even trigger scenario responses, e.g., reroute shipments, adjust inventory buffers, renegotiate supplier orders, or shift production volumes - all within governed boundaries.




Synchronized "plant-to-shelf" orchestration

In retail, as an example, upstream (manufacturing, packaging) and downstream (distribution, freight, warehousing) are tightly coupled with volumes and freshness. A unified agentic twin can continuously align production schedules with transport capacity, SKU-level demand signals, and packaging constraints. This helps minimize waste, reduce stockouts or overproduction, and improve fleet utilization. At scale, these advanced twins help companies anticipate bottlenecks to optimize buffers and tune the supplier base.



Shipping, logistics and cold-chain optimization

In shipping and logistics, agentic twins can autonomously optimize routing, carrier selection, dynamic re-scheduling, and fuel/emissions trade-offs, as prototyped by IBM and Imperial College London [1]. In cold-chain logistics slight deviations in temperature or scheduling can degrade product. AI + digital twins are starting to drive smarter allocation of inventory by shelf life and routing efficiency.



Agentic brand/packaging twins

Agentic packaging twins change static design cycles into living, self-optimizing systems that simulate and respond to real-world feedback. To improve sales performance, an agentic twin could model how packaging performs in different retail environments by testing shelf placement, lighting, and shopper visibility virtually before launch. To improve design requirements, it could propose design variants optimized for local regulations, material availability, recyclability, and consumer preferences, all while keeping alignment with sustainability and brand integrity principles.

Key enablers and constraints

Dimension	Enablers	Constraints / Risk	Mitigation
Data & Connectivity	High-frequency telemetry (factories, logistics, warehouse), LLM-ready corpora (specs, contracts, logs), semantic integration (KGs)	Many suppliers and 3PLs may lack real-time data	Start with pilot tiers. Use hybrid mode (agent suggests, human approves) to manage risk
Semantic alignment and knowledge graph	Embedding meaning (SKUs, nodes, constraints) enables agents to “understand” domain concepts	Without semantics, agents hallucinate or misact	Invest early in ontology, taxonomy, and lineage
Surrogate or hybrid models	Use physics-informed models, surrogate ML models, or fast simulators inside twin to let agents experiment	High-risk areas (product safety, regulatory) are not suitable for experimentation without guardrails	Must ensure fallback to validated high-fidelity models in critical decisions
Governance and auditability	Agents must log actions, trace decisions, allow rollback, and enforce safety thresholds	Overreliance on autonomous agents without oversight is high risk	Design with human-in-the-loop checkpoints and escalation paths, with in-built audit trails.
Pilot-to-scale architecture	Design modular agent ecosystems and orchestrate toolchains (e.g., via Model Context Protocol) to enable tailored solutions that scale	Ensuring flexibility and scalability without sacrificing the ability to deliver custom high-value systems	Use open integration standards and frameworks such as OpenUSD (for 3D design in digital twins) and OPC UA / ISO 23247 (for industrial interoperability with sensors and devices).

Example roadmap for deploying agentic digital twins



1. Select high-leverage domains

Begin with logistics, shipping, and inventory buffers for a subset of SKUs or corridors. For instance, cold-chain distribution routes in a single region.



2. Build a sandbox twin & agent framework

Create a domain scoped twin: a shadow agent that proposes but doesn't enact, evaluated over 3 to 6 months and uses scenario stress tests, e.g. port closure.



3. Governed execution in low-risk lanes

Deploy agentic actions for less critical flows, e.g., replenishment, carrier switching, with a human-in-the-loop supervision and evaluation model.



4. Incremental expansion & agent cooperation

Link separate agents, logistics, production, packaging, via a meta-agent or orchestration, enabling cross-domain negotiation inside the larger twin.



5. Continuous learning and model refresh

Use feedback loops: monitor agent decisions vs outcomes, recalibrate models, update surrogate models, and enhance knowledge graph coverage.



6. Capability advancement

Finetune semantic modelling, delineate guardrails, and design agent orchestration patterns tailored to constraints, e.g. shelf life, freshness, regulatory zones.

Making agentic digital twins a reality

Implementing agentic digital twins will demand strong technical capability with a comprehensive understanding of the constraints/risks and how to mitigate them.

However, the most important component is a deep understanding of the industry's unique complexities. In retail for example, the realities of SKUs, packaging formats, cold-chain logistics, supplier networks, and the

constant trade-offs between cost, freshness, and waste are embedded in the data, but only interpretable by those who understand the pressures behind them.

Ultimately, agentic digital twins allow us to reimagine our systems as context-aware, trustworthy, and scalable for tomorrow. The impact is derived from alignment with the operational realities of today.

Impact 2

Data governance in the Agentic Age

Data governance in the Agentic Age

AI is only as good as the data it uses; concern for “garbage in, garbage out” is still true. Without clean, trustworthy data, automation amplifies errors instead of insights. For global enterprises, where data spans brands, products, markets, and regulatory environments, data governance is a strategic imperative.

Traditional methods of cataloguing, lineage tracking, and static quality checks cannot keep pace with dynamic, AI-driven environments. These approaches create snapshots of data that quickly become outdated.

However, highly manual, documentation-heavy activities can now be augmented with automation, pattern recognition, and natural language understanding. In turn, this AI-assisted data management will enable us to scale and adapt our governance systems to meet the complexities of agentic flows.

The four pillars of AI-enabled data management

1. Active metadata

Most organizations still treat metadata catalogues as lists rather than living systems, using tools that help automate the population of attributes and relationships. However, these catalogues become outdated snapshots.

Active metadata ensures governance reflects reality, not history, which is critical for environments where data changes continuously. Active metadata integrates with marketplaces, identity management, and data quality tools to surface insights instantly.

Key components to an active metadata catalog:

1

A meta-model model:

A blueprint defining relationships, attributes, and governance rules, providing the backbone for consistency as AI scales.

2

Interactivity:

Catalogs enable discovery, decision-making, & compliance. Interfacing catalogues through marketplaces helps business users engage.

3

Enablement:

Go beyond evangelism and training by embedding catalogues into workflows and decision-making to show their value.

Active metadata shifts the mindset from documentation to intelligence, turning metadata into a live layer that drives governance and business enablement.

2. AI-assisted metadata discovery and curation

Change in products, features, and regulations is continuous - thus defining business metadata is a never-ending task. LLMs help solve the “blank page problem” by generating conceptual models and draft definitions, reducing effort while improving consistency.

There are two approaches to using LLMs for metadata discovery. By creating generic models for definition and then refining to specific domains, governance teams can evangelize a broad foundation across the enterprise. Alternatively, creating a domain-focused model, aligned to a real, contextual question raised by the business, can help illuminate the concepts that need to be governed


while demonstrating value quickly. For example, using an LLM to answer questions such as: “What should my nutritional and allergen labelling requirements be based on in Singapore?”

As with any AI implementation, using LLMs for metadata discovery should not remove the human in the loop. The governance of metadata requires judgement, context, and accountability. In both approaches, AI provides a foundation from which they can scale, significantly reducing effort while improving the consistency and quality of business definitions.

3. Automated lineage and impact assessment

Traditional lineage shows “what connects to what.” AI adds interpretation - explaining why connections exist and what changes mean. LLMs can read SQL scripts and transformation logic to detect semantic shifts (e.g., revenue calculation changes) and flag them for review.

For example, if a transformation script changes from *revenue = gross sales* to *revenue = gross sales – discounts*, an LLM can recognise the semantic shift and flag it as a potential business definition change for review by stewards or data owners. Thus, AI enables lineage to become proactive rather than descriptive. It no longer just shows “what happened”; it helps anticipate “what is changing”, “what it means”, and “who needs to know”.



4. Data quality automation

Data quality underpins trustworthy AI, but static rules cannot keep up with the volume and complexity of data. AI uses metadata and lineage to infer quality checks, prioritizing issues by business impact.

Example: If “Customer ID” joins multiple systems, AI suggests uniqueness and completeness checks automatically. Alerts are sent to responsible stewards or product owners to confirm.

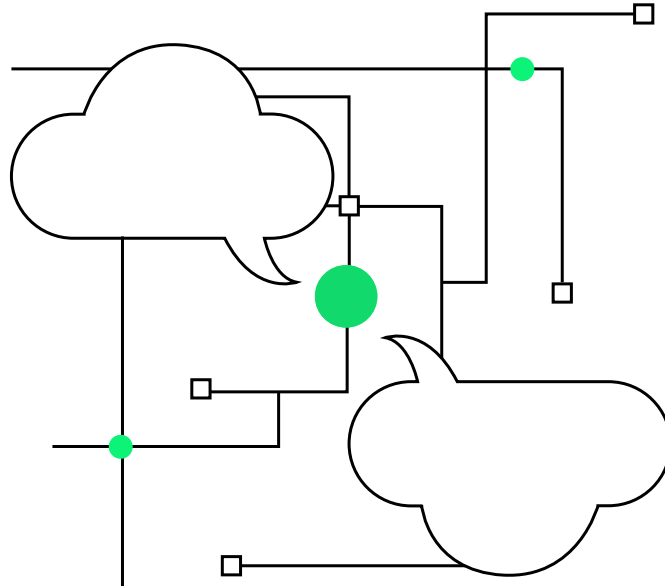
Once established, these rules can be operationalised through monitoring pipelines and observability tooling.

Over time, feedback loops refine the rule library, learning which anomalies were material and which were noise.

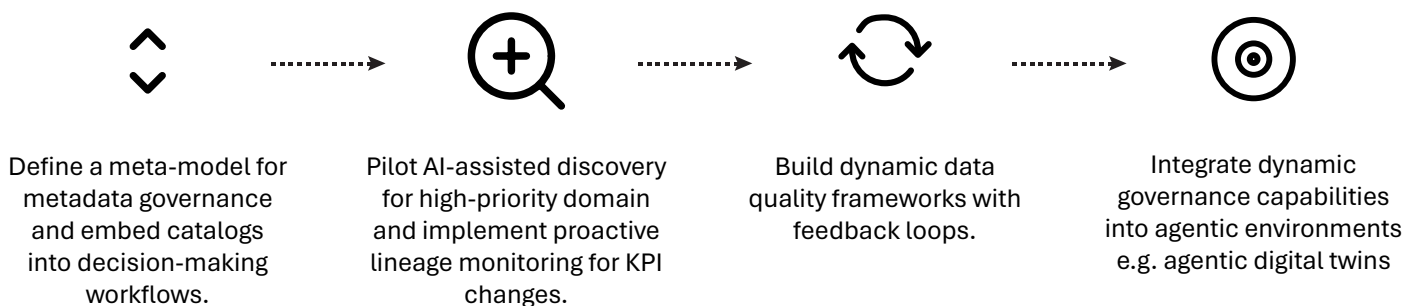
The result is a dynamic and context-aware data quality framework. Instead of maintaining thousands of static checks, organizations can move towards rule discovery, validation, and alerting that adapts as data and definitions evolve. It turns data quality from a defensive exercise into a proactive, intelligence-led capability — one that continuously reinforces trust in the organization’s most important data assets.

AI-automated, human-led

Ungoverned automation of data management components risks spreading errors quickly -thus none of these AI systems should be left ungoverned. AI will not replace the governance professional, but it will redefine their craft. The role of data stewards and owners will shift from documenting what already exists to curating and validating what AI discovers. In this way, we can build data management processes that keep up with the complexities of product and value chains, rather than timestamping information that has already changed.



Example roadmap for building AI-assisted data governance




We can combine human judgement with machine-scale insight, creating an intelligent and trustworthy governance model to reduce effort and achieve the promise of data governance: confidence, context, and control at scale.

Impact 3

Exposing hidden cracks with agent- ready data

The unseen risk



Centralized data teams carefully engineer, manage, and control data systems to feed downstream analytics and custom AI products. Now, plug-and-play agents can avoid the carefully curated foundations and operate in the messy top layer: work suites, collaboration tools, and unstructured content.

Security gaps that were tolerable when humans were the primary consumers of data become critical when agents take over. Research from Varonis [2] shows most organizations are experiencing 4 key security gaps:



Stale users: Nearly 90% of organizations have accounts unused for 90+ days, exposing sensitive data to dormant identities.



Over-permissioned cloud roles: Thousands of overly permissive policies in environments make enforcing least privilege nearly impossible.



High-risk OAuth apps: One in four OAuth apps are unverified and high-risk, yet over half of employees use them - often without IT oversight.



Public links: 92% of organizations allow public sharing links, risking exposing confidential information.

For agents, these issues create gateways to bad outputs, breaches, and distrust. Agents searching through outdated, inconsistent files can amplify errors at scale. For federated enterprises, complexity multiplies: different business units operate distinct storage models and access controls. Shadow economies of Copilot and agent tools are emerging as employees adopt GenAI solutions in their personal work.

Enter agent-ready data

As organizations deploy agents, even in controlled environments, poor results can result in a loss of trust. Agent-ready data is permissioned, retrievable, and certified for agent consumption, provided the basis for workflows to achieve a minimal viable standard of trust from which to improve outputs.

- **Governed semantic access:** A unified access layer serving certified views of core entities (product, ingredient, claim, price, customer, supplier) while preserving business unit ownership.
- **Policy-aware retrieval:** Classify at ingest, enforce row/column filters at retrieval, and capture lineage by default for traceability.
- **Non-human identity hygiene:** Short-lived credentials, least privilege by default, and routine removal of stale users or tokens.
- **Trust signals at runtime:** Tag freshness, source, owner, and approval status. Block and notify when data fails trust checks.
- **Agentic data quality:** Use agents to detect downstream issues and automate fixes at scale.

Agent-ready data means less rework on outputs from agents with legal, regulatory, and brand teams and provide a traceable output to support approval processes.

Agent-ready data in action

Use Case	Opportunity	Risks without agent-ready data	Benefits with agent-ready data
Regulatory compliance in product claims	Agents automatically validate marketing claims against certified reference data.	Agents validate claims against outdated or inconsistent ingredient data from unstructured sources → non-compliant marketing, regulatory, reputational damage.	Traceable compliance, reduction in non-compliant content, reduced legal review cycles, and approvals.
Dynamic pricing and promotions	AI agents pull certified product and pricing data to adjust promotions in real time.	Agents rely on stale spreadsheets or uncontrolled sources. Pricing errors, margin erosion, inconsistent promotions across markets, customer distrust.	Governed semantic access and policy-aware retrieval. Real-time certified pricing, synchronized global promotions, improved trust.
Supplier risk monitoring	Agents scan supplier data for compliance signals (e.g., certifications, sustainability scores) and flag anomalies.	Agents scan fragmented supplier records. Miss expired certifications or ESG gaps. Procurement from high-risk suppliers, regulatory exposure.	Lineage-aware retrieval and trust signals. Proactive anomaly detection, compliance with ESG and procurement standards.


Example roadmap map for implementing agent-ready data

Building agent-ready data requires a taskforce approach that combines governance, automation, and trust signals:



1. Set a single enterprise data classification policy

Enforce classification at ingest and block unclassified content from retrieval. This ensures agents only access certified, structured data.



2. Create a non-human identity standard

Implement short-lived credentials and least privilege by default. Automate hygiene with scripts or AI to audit stale identities and revoke unused tokens quarterly.



3. Mandate attribute-based access control (ABAC) for AI retrieval

Attribute-based access control ensures granular permissions for agents, reducing risk from over-permissioned roles.



4. Publish certified reference entities

Establish governed semantic access for core entities: e.g., product, ingredient, price, supplier, and customer to create a single source of truth for agents.



5. Standardize DPIAs for agentic use cases

Auto-populate processing records from lineage and integrate runtime trust checks so outputs are blocked if data fails freshness or approval tests.



6. Measure trust and quality metrics

Track certified dataset coverage, stale identity count, AI output acceptance, and audit findings closed. Deploy agents to detect inconsistencies in unstructured data.



7. Establish monitoring and anomaly detection

Detect policy bypass, over-retrieval, and unusual data movement. Run quarterly checks on agent misuse to maintain compliance and brand safety.

Making agent-ready data a priority

Agents are already solving complex challenges, from real-time supply chain decisions to dynamic pricing. Agent-ready data means revenue protection, brand safety, and speed-to-shelf. Frame the business case in board language and measure ROI through the metrics such as: compliance exposure, operating cost, and AI output acceptance.

A cross-function task force, who understand the risks and benefits of agent-ready data can help position it as a strategic enabler for the next wave of AI-driven growth.

Impact 4

The ROI of ethical AI: measuring trust earnings

The ROI of ethical AI

Predictive, generative, and now agentic AI are not only shaping our products and processes, but customer perception. Increasingly, customers reward companies not simply for deploying sophisticated AI but for doing so responsibly, transparently and with clear respect for their data. Trust, long treated as a soft sentiment, has become a measurable driver of growth.

Customers who view their providers as innovative and responsible spend significantly more than those who do not share that view. The uplift, which is through trust, is substantial enough to be considered a new form of ROI.

Consumers may pay up to 25% more for AI-enabled tools that demonstrate responsible data use and transparent decision-making [3]. When customers place a premium on ethical implementation, ethical implementation becomes a revenue strategy.

Trust and ethics are rarely quantified to date with the same level of precision as financial metrics. Leadership teams acknowledge their importance but struggle to demonstrate how responsible data use translates into revenue and loyalty. The challenge is to embed it into customer and data strategy.

From principles to performance

While building customer-facing agentic AI, organizations must define what ‘trust ROI’ and ‘ethical ROI’ represent.



Trust ROI refers to the value created when customers believe in how an organization uses their data and deploys AI on their behalf.

Ethical ROI refers to the investment in AI governance to reduce bias, increase transparency and security; and enable sustainable customer engagement.

These definitions are only meaningful once linked to measurement. The organizations that recognize which measures are most relevant to them begin will gain a structural advantage.

Measure 1:

Spending behaviour

By embedding trust measures into existing feedback, such as renewal conversations or onboarding surveys, organizations can correlate high trust scores with commercial performance. The ‘trust premium’ becomes visible in the behaviors of customers with the highest trust rating: basket size, repeat purchases, acceptance of cross-sell attempts and adoption of new AI-enabled features. The uplift can then be modeled like other revenue driver.

Measure 2:

Customer churn

Over half of consumers have abandoned a provider following a data-related loss of trust [3], making trust erosion a quantifiable liability. Forward-thinking organizations conduct rapid ‘trust pulse checks’ after incident and track customer retention over time to identify the cost of trust loss and drive investments in governance, communication, or model oversight.

Measure 3:

Advocacy

When customers feel confident in a company’s AI practices, they are more likely to mention those characteristics in their reviews. Combining review analysis, referral rates, and sentiment analysis from social listening can help establish the link between transparency and growth.

Measure 4:

Monetization of ethical AI

Some companies are beginning to test ethical AI as a value proposition, offering versions of products which transparency, explainability, or user control. A/B tests can reveal whether customers convert at higher rates when responsible AI features are visible, thereby turning ethical ROI into a direct revenue driver.

Measure 5:

Risk reduction

Fairness audits, explainability tooling and privacy-by-design frameworks all carry cost, but they offset larger regulatory and reputational liabilities. Organizations that take a structured approach can model the value of avoiding these events, using industry breach data or previous incidents to estimate the financial impact and ROI of prevention.

Measure 6:

Brand integrity

Many organizations now incorporate data and AI ethics into their brand tracking studies, monitoring whether customers view their AI systems as fair, secure or transparent. Over time, these measures correlate strongly with revenue growth and retention.

Moving from insight to action

Translating these measures into strategy requires the operational conditions for trust to accumulate. This typically begins with governance: cross-functional ethics boards that can evaluate the design and deployment of AI systems, challenging decisions through the lens of customer value, societal and human impact; and regulatory compliance.

Equally important is transparency. Companies that explain how their AI works, using accessible language and visible

guardrails, earn customer confidence. When transparency is embedded in the customer experience, trust becomes an everyday interaction.

Customer control plays a critical role. Providing options for how data is used and the ability to update permissions signals maturity. The introduction of controls often leads to deeper data sharing. Trust enables data depth, and data depth enables better AI.

Example roadmap for enabling trust earnings measurement



1. Define trust and ethical ROI metrics

Establish clear definitions for trust ROI & ethical ROI linked to measurable outcomes (spend uplift, churn reduction, advocacy rates). Integrate trust questions into existing customer feedback (onboarding, renewals).



2. Embed trust measurement into CX and data strategy

Pilot trust pulse checks after key events (feature launches, data policy updates). Correlate trust scores with revenue, retention, and advocacy to quantify the trust premium.



3. Operationalize transparency and control

Introduce guardrails in AI-driven experiences (e.g., “Why this recommendation” features, data usage dashboards). Provide customers with clear options for data permissions and revocation.



4. Establish governance for ethical AI

Form a cross-functional ethics board to review AI design and deployment. Implement fairness audits, explainability tooling, and privacy-by-design frameworks.



5. Monetize ethical AI as a differentiator

Test A/B scenarios where transparency and explainability highlighted as product features. Track conversion and willingness-to-pay for ethical AI options.



6. Monitor and report trust ROI

Build dashboards that track trust metrics alongside financial KPIs. Include trust earnings in quarterly performance reviews to elevate accountability.

Embracing trust earnings as a critical measure defines the new standard for customer relationships. The companies that can demonstrate trustworthy data practices, ethical AI, and transparent governance protect themselves from downside exposure and unlock a upside: higher revenue, greater advocacy, stronger loyalty, and a more resilient brand.

Impact 5

Human data and AI talent in an agentic world

Human data and AI talent in an agentic world

Across the UK and US, graduate and entry-level job postings have fallen by a third since late 2022[4]. This isn't just another cyclical dip; it's a structural shift with deep implications for business succession and future capability – and it is accelerated by AI.



The agent effect on junior talent

Junior roles, which often involve well-defined, repetitive tasks, have been the first affected. According to Stanford University [5], employment for junior software developers fell by nearly 20% since late 2022 in the US, while in the UK, postings for AI-exposed roles have dropped 38% - almost twice the decline of other sectors. Freelancers and gig workers in creative and technical fields have also experienced significant reductions in work volumes and pay rates.

As entry-level IT vacancies decline and graduate schemes are reduced, planned hiring has reached its lowest point in over a decade. For most, the result is an ageing, costlier workforce with fewer succession options, especially in sectors where leaders are grown from within. Without a

steady intake of adaptable junior talent, organizations risk stagnation, wage inflation, and a future leadership shortage.

Beyond the immediate job losses, inflated AI expectations are reshaping the workforce, creating a top-heavy “diamond” structure instead of a healthy pyramid. Although AI investment is generating new opportunities, these are not emerging at the scale or speed required to replace lost entry-level positions. According to Stanford’s study [5], while entry-level employment in AI-exposed fields (such as software engineering and customer service) fell ~20% between 2022 and mid-2025, employment for older workers in those same roles actually rose.

The myth of AI as a talent substitute

At the root of this trend is a board-level belief that artificial intelligence will offset lost workforce capacity. But the data tells a different story; the impact and resulting productivity gains of enterprise AI and automation are lagging behind expectations. A recent study from MIT states that “Despite US\$30-40 billion in enterprise investment into GenAI, [...] 95 % of organizations are getting zero return”[6]. Moreover, EY’s 2025 survey found that, despite the introduction of AI, over two thirds of employees reported an increase in their workload, without tangible productivity gains from AI or the talent to help manage it [7].

Missing generation of talent today will create a leadership vacuum tomorrow. Succession cracks are already visible: more than half of company leaders have no identified successor, and only 50% of HR leaders feel prepared to

lead their organizations into the future [8]. Poorly managed succession can have severe consequences, including stalled projects, loss of institutional knowledge, and even destruction of market value; studies estimate that nearly \$1 trillion is lost annually in the US due to poorly managed C-suite transitions [9].

The innovation that junior talent brings, with fresh perspectives, digital fluency, and adaptability, is also at risk. Cutting roles and graduate schemes undermines inclusion and long-term competitiveness. Research consistently demonstrates that companies with diverse leadership teams outperform their peers on profitability and innovation. Scaling back such programs damages not only morale and reputation but also the creative edge needed to compete in a fast-changing AI landscape.

Taking action

This isn’t just an HR issue - it’s a strategic, boardroom-level risk. Organizations that will thrive in the AI era are those secure and develop their next generation of data and AI professionals. That means reinvesting in early-career talent, protecting graduate pipelines even when budgets tighten, and pairing this with continuous learning and AI literacy across the workforce. It also requires a proactive approach to succession planning, tracking the age and skills profile of critical teams, and ensuring that entry-level routes remain open to a broad range of backgrounds and perspectives.

The future of data and AI in business will be shaped not just by technology, but by people. The real question is whether organizations are narrowing their future talent runway to the point of no return. The time to act is now, before the next generation of leaders disappears for good.



Concluding thoughts

Having agency over our agentic future

Agentic AI is already reshaping how businesses think about data, automation, and decision-making. The implications go far beyond technology adoption; this is about reimagining operating models, governance frameworks, and talent strategies.

Agentic digital twins show what is possible when intelligence moves from simulation to action, unlocking agility across supply chains. Yet these capabilities depend on the integrity of the underlying data - data that is permissioned, retrievable, and certified for agent consumption. Without this foundation, autonomy is a risk rather than a reward.

Agentic AI demands trust from all parties: in the data being used by agents, trust in the outputs from business users, and trust in end products from consumers. Ethical AI becomes a competitive growth strategy that influences revenue, loyalty, and advocacy.

Agentic AI creates operational efficiencies, but it can't replace future leaders. Organizations need to prioritize their talent pipeline to create a workforce capable of working with AI.

Ultimately, we must reevaluate how we understand and interact with technology, treating data not as a static asset but as a living system, and AI not as a bolt-on tool but as an embedded capability.

Kubrick's approach

Founded with the mission to transform lives through data and AI, Kubrick focuses on the most pressing challenges facing organizations and their customers. In the age of Agentic AI, this is an existential question about people, principles, and purpose.

Over the last decade, Kubrick has hand-picked and trained over 4,000 data and AI professionals, creating a workforce of brilliant minds who bring a diverse perspective to technology and how it is applied to domains. Working with clients across the AI capability spectrum, Kubrick is helping to define roadmaps, create safe environments for innovation, helping accelerate deployment for first-mover advantage, and tackling each of the five underlying impacts explored in this paper.

The insights in this report are thanks to the first-hand experience and expertise of Kubrick's technical leadership and SMEs: Marija Milojevic (Agentic digital twins), Harry Hazelwood (AI-assisted data management), Jessica Macleod (Agent-ready data), Kornelia Lencz (ROI of ethical AI), Ralph Robinson (Future of talent), and Camilla Dickson (Research and editorial).

Reflection in practice

Reflect on the ideas explored throughout this paper and apply them to your own context to help sharpen your goals, challenge points, and areas for improvement in your AI strategy.

Agentic AI in practice: digital twins

- Where in your operations could agentic digital twins have the biggest impact?
- What would you need to trust a digital twin making autonomous decisions? What guardrails would you want in place?
- Where are you on the roadmap to deploying agent-driven simulations at scale?

Data governance in the Agentic Age

- What is your biggest governance gap to making agentic AI risky or difficult to scale?
- How mature are your AI-enabled governance practices?
- Who should own governance for autonomous systems?

Agent-ready data

- What are the biggest obstacles to creating high-quality, agent-ready data?
- What manual process would you hand over to an agent tomorrow? What process would you never hand over — and why?

Trust earnings as ROI

- How are your customers' trust expectations shifting as AI becomes more autonomous?
- Do you believe trust can be measured as ROI? If so, who is accountable?
- How transparent should autonomous systems be?

Human data & AI talent

- How is agentic AI changing your view of workforce planning in entry-level roles?
- What new capabilities will you need in three years that you don't have today?
- How are you thinking about reskilling or redeploying roles displaced by agentic workflows?
- What cultural or organizational barriers could slow adoption of autonomous systems?

References

[1] A Timms, A Langbridge, A Antonopoulos, A Mygiakis, F O'Donncha (2025) Agentic AI for Digital Twin. Available via IBM

[2] Varonis (2025) State of Data Security Report. Available via Varonis

[3] S Fineberg, S Hupfer, J Loucks, M Steinhart, S Mazumder (2025) Connected Consumer study. Available via Deloitte

[4] T Allas, A Goodman (2025) Not yet productive, already disruptive: AI's uneven effects on UK jobs and talent. Available via McKinsey

[5] E Brynjolfsson, B Chandar, R Chen (2025) Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence. Available via Stanford University

[6] A Challapally, C Pease, R Raskar, P Chari (2025) The GenAI Divide: State of AI in Business 2025. Available via MIT

[7] EY (2025) Workforce Reimagined Survey. Available via EY

[8] Robert Half (2025) Succession and Pipeline Planning. Available via Robert Half

[9] C Fernandez-Araoz, G Nagel, C Green (2025) The High Cost of Poor Succession Planning. Available via Harvard Business Review

