Resonate AI

How EU-Based Consultancies Are Adopting AI:

Capabilities, Governance, and the New Competitive Edge

November 2025

Adam William Chalmers

TABLE OF CONTENTS

EXECUTIVE SUMMARY

This report presents the first systematic analysis of how **EU consultancies are adopting, governing, and communicating artificial intelligence (AI)**. Drawing on data from **645 firms** in the EU Transparency Register, the study combines large-scale multilingual web scraping, translation, and computational text analysis to map adoption levels, discourse framing, and regulatory readiness.

Key Findings

- **Limited visibility:** Only **22.8%** of consultancies explicitly mention AI in their public materials. This signals early adoption *and* selective disclosure.
- **Emergent maturity:** The average **AI maturity score is 2.32 / 5**, indicating partial and experimental integration rather than organization-wide transformation.
- **Operational uptake:** Roughly **62.6%** of firms report some internal use of AI, and **57.3%** offer AI-related services to clients, showing that capability and market offering evolve together.
- **Governance gap:** Mean **policy strength is 0.51 / 2** and **EU AI Act readiness 0.39 / 3**, revealing that regulatory preparedness lags behind technical experimentation.
- **Regional contrasts:** Austria (41%), Ireland (31%), and France (30%) lead in public AI visibility, while Germany (16%) and Belgium (19%) remain below the EU average.
- **Typology:** Three-quarters of firms are **Sceptics** (low talk, low action), 9% **Fast Followers**, 3% **Leaders**, and 0.5% **Talkers**, reflecting wide heterogeneity in AI engagement.
- **Technological focus:** Adoption concentrates on **predictive analytics (31%)**, **Data science (31%)**, and **Automation (16%)**, with **Generative AI (9%)** still at the pilot stage.

Overall, EU consultancies are in a **transitional phase**, balancing cautious exploration with growing regulatory awareness. A small group of **Leaders** already demonstrates mature, compliant, and transparent AI integration that may soon define market benchmarks.

INTRODUCTION

Artificial intelligence (AI) is reshaping the professional-services sector, yet a central question remains unanswered: **how effectively are consultancies, the firms advising others on digital transformation, adopting AI themselves**? The introduction of the EU Artificial Intelligence Act (AI Act) has intensified this tension. Consultancies must now demonstrate not only technical capability, but also governance, accountability, and readiness for a new regulatory environment.

This report presents the first large-scale, data-driven analysis of AI adoption among European consultancies. Unlike surveys or self-reported assessments (which often overstate readiness and rely on selective disclosure) this study uses a novel multilingual web-scraping and text-analysis pipeline to observe what firms publicly communicate and publish about their AI use, governance structures, and regulatory engagement. This method captures real behaviour rather than perceived or aspirational claims, offering a uniquely reliable view of how AI adoption is unfolding across the sector.

Consultancies occupy a pivotal position in Europe's AI ecosystem. They are the intermediaries who translate AI technologies into organisational practice, interpret the EU AI Act for clients, and influence policy through their role as registered lobbyists. Their internal adoption patterns therefore provide early signals of how AI capability, responsible governance, and regulatory compliance will diffuse across the wider economy. Put simply: consultancies shape how AI becomes "real" inside organisations.

Drawing on data from **645 consultancy firms** listed in the EU Transparency Register, this report maps three interconnected dimensions of AI engagement:

1. **Adoption and maturity**: the extent to which firms use AI in their own operations and in client-facing services. Here, maturity refers to the degree to which AI is intentional, embedded, governed, and connected to real workflows, rather than isolated tools or marketing claims.
2. **Governance and readiness**: the presence of oversight structures, ethical principles, and preparedness for obligations under the EU AI Act.
3. **Discourse and framing**: how firms communicate AI's value, risks, and regulatory implications, and how their rhetoric aligns (or does not align) with their capabilities.

Across these three dimensions, a clear picture emerges: only **one in four** consultancies publicly reference AI, and the average maturity score is 2.3 out of 5, signalling early-stage or experimental integration rather than organisation-wide adoption. Beneath this modest aggregate, however, lies a highly uneven landscape: from **AI-Sceptics** with virtually no adoption, to **Fast Followers** building structured capabilities, to a small group of **AI-Leaders** integrating AI with governance, ethics, and regulatory readiness.

The analysis moves beyond description to investigate what drives movement along this adoption ladder. Regression results show that capability and governance (not size, budget, or rhetoric) predict whether a consultancy becomes a Leader. Firms that use AI internally, apply it in client services, and formalise oversight mechanisms are those best positioned to meet both market demand and regulatory expectations.

This report captures a pivotal moment for the consulting industry. As AI regulation crystallises and client expectations accelerate, consultancies will play an outsized role in shaping how "trustworthy AI" is implemented in practice. Their evolving behaviour (adoption, governance, and communication) will ultimately influence not only their own competitiveness but also the real-world impact of Europe's AI regulatory framework.

WHY CONSULTANCIES MATTER FOR EUROPE'S EMERGING AI LANDSCAPE

Consultancies play a uniquely influential role in Europe's emerging AI landscape. They are, of course, not the primary developers of frontier AI models. They do, however, occupy a strategic position where technology, regulation, and organisational practice intersect. Understanding how consultancies adopt and govern AI therefore offers critical insight into how AI capability (and compliance with the EU AI Act) will diffuse across the broader economy.

First, consultancies act as **intermediaries of technological adoption**. Most organisations do not engage directly with complex AI systems; they rely on external advisers to evaluate tools, benchmark risks, redesign processes, and guide implementation. As a result, the consulting industry serves as a transmission mechanism for AI knowledge, influencing how technologies are selected, integrated, and scaled. If consultancies are slow to adopt AI internally, lack governance structures, or rely on rhetoric rather than substance, these weaknesses are likely to cascade across the clients they advise.

Second, consultancies are key players in the evolving regulatory ecosystem for AI. As the EU AI Act introduces new obligations around risk management, documentation, transparency, and oversight, organisations will increasingly turn to advisory firms for support. Many consultancies already **offer AI governance, risk assessment, compliance audits, and training programs**. Their own readiness, policy frameworks, ethics principles, governance structures, therefore provides a preview of how effectively the EU AI Act may be implemented in practice. High readiness among consultancies could accelerate compliant adoption across sectors; low readiness could become a bottleneck.

Third, consultancies are not only service providers but also **policy actors**. Every firm in this dataset is registered in the EU Transparency Register, indicating formal engagement in public policy or regulatory affairs. Through white papers, consultations, events, and lobbying, consultancies shape how AI is framed, as an opportunity, a risk, a compliance issue, or a transformation driver. Their discourse influences not only clients but also policymakers and public debates.

Finally, **consultancies operate in a sector where credibility is a core asset**. Expertise, trust, and reputation are the foundations of their business models. This makes them an ideal lens for examining the alignment (or misalignment) between **AI talk** and **AI action**. A consultancy that markets AI expertise without internal capability risks damaging the very trust that underpins its value proposition.

DATA & METHOD

This study uses an innovative, large-scale, **multilingual web-scraping** and **text-analysis pipeline** to map how EU-based consultancies describe, adopt, and govern AI. Unlike surveys or interviews, which rely on self-reporting and often overstate readiness, this method captures what firms publicly disclose in practice, offering a more objective and behaviour-based view of AI engagement. The dataset comprises **645 consultancy firms** registered in the EU Transparency Register that had valid, publicly accessible websites as of October 2025. This represents *all* firms registered and identified as "Professional consultancies" on this date. These firms represent consultancies operating across all 27 EU Member States, with the largest concentrations in Belgium, Germany, France, Italy, Spain, and the Netherlands. All companies were screened for accessibility and inclusion based on three criteria: (1) Valid URL and domain availability; (2) Active HTTP/HTTPS pages retrievable via automated crawling; (3) Presence of substantive textual content (minimum 500 characters). The resulting corpus reflects a diverse and multilingual snapshot of the European consulting industry, spanning sectors from management and legal consulting to digital and technology advisory.

DATA COLLECTION AND VARIABLE CONSTRUCTION

To extract text from each firm's public website, a custom, multilingual scraping pipeline was developed in Python, using *requests*, *BeautifulSoup*, and *trafilatura*. For every consultancy, up to five priority pages were crawled within a depth of two to three link levels. To ensure ethical and responsible data collection, the crawler followed polite scraping (1–3 second randomized delays), robots.txt compliance, and custom user-agent identification. Extracted text was cleaned to remove menus, cookie banners, footers, disclaimers, and duplicate content. Because 65% of firms publish primarily in non-English languages, all relevant text was **translated** into English using GPT-4o-mini, enabling full cross-country comparability while preserving local nuance.

The final dataset contains **645 firm**s and **97 structured variables**, capturing adoption behaviours, governance practices, regulatory references, discourse frames, and technology focus. Specifically, each firm's data record captures multiple dimensions of AI engagement, including:

- **Adoption Indicators**: internal AI use, client-facing AI services, innovation claims, and AI maturity (0–5 scale).

- **Discourse Features**: tone (positive, neutral, critical), discourse frame (e.g., efficiency, compliance, transformation), and language markers of hype or uncertainty.
- **Governance & Ethics**: presence of AI ethics statements, codes of conduct, governance structures, and references to safety or oversight.
- **Regulatory Readiness**: engagement with the **EU AI Act** (score 0–3), including mentions of compliance, risk classification, and readiness programs.
- **Technological Focus**: mentions of AI subfields (predictive analytics, data science, automation, generative AI, machine learning, etc.).
- **Training & Capability**: estimated share of employees trained in AI-related skills, inferred from textual mentions.

Quantitative variables were normalized to a 0–1 range or standardized (z-scores) for cross-variable comparison.

AI-related content was identified using an expanded multilingual keyword lexicon covering: (1) core AI concepts (AI, machine learning, LLMs); (2) applications (predictive analytics, automation, NLP, computer vision); (3) responsible AI and ethics terminology; (4) regulatory references (EU AI Act, GDPR, CSRD). Each detected keyword triggered a **±150-word context window** to preserve semantic framing.

Rule-based and semi-automated classification mapped content into higher-order analytical categories: Tone & sentiment; Framing (efficiency, transformation, risk, compliance, opportunity); Ethical principles (fairness, transparency, accountability, privacy, bias); Regulatory engagement (EU AI Act, GDPR, CSRD references); and Technology mentions (data science, automation, generative AI, etc.)

Finally, composite indicators were developed for three key variables: (1) **AI Maturity (0–5):** governance (2 pts), training (1), use/application (1), ethics/compliance (1); (2) **Policy Strength (0–2):** codes, governance, and safety mechanisms; and (3) **EU AI Act Readiness (0–3):** extent of explicit regulatory alignment. These scores were used to construct indices for **Action**, **Governance**, and **Discourse**, underpinning the typology and correlation analyses reported in later sections.

Manual quality checks were performed for a random 10% of cases to validate: (a) Website reachability and text accuracy; (b) Correct language detection and translation fidelity; and (c) consistency in keyword extraction and classification. Automated diagnostics recorded **scrape confidence**, **keyword density**, and **text length** to flag potential anomalies.

While this approach is useful for overcoming challenges associated with surveys and self-reporting, it nevertheless has its own limitations. Notably, the analysis captures only publicly available discourse, not confidential project-level AI implementation. Additionally, “translation

artifacts"[1] may slightly alter tone or nuance in multilingual content while the "keyword-based" detection used in this analysis may undercount emerging AI terminology not covered by the lexicon. Finally, it is important to acknowledge the data only represent a temporal snapshot (October 2025), not longitudinal change.

A full account of crawling parameters, translation settings, keyword dictionaries, and scoring logic is provided a separately published **Appendix: Methodological Details**.

[1] Translation artifacts are small distortions introduced when text is translated from one language to another, often shifting tone, nuance, or emphasis.

This section analyses how EU consultancies are adopting artificial intelligence (AI) in their operations and service offerings. The analysis focuses on three complementary dimensions: whether firms publicly reference AI, their internal maturity in deploying AI tools, and the extent to which AI is integrated into client-facing work. Together, these measures provide a multi-layered picture of where the consulting sector stands in Europe's AI transition.

As shown in **Figure 1**, only 22.8 percent of consultancies in the sample (N = 645) explicitly mention AI in their online materials. **The remaining 77.2 percent make no reference to AI at all.** This is a strikingly low level of public AI engagement for a sector expected to advise clients on digital transformation. Despite the intense visibility of AI in public debate, corporate strategy, and regulatory planning, most consultancies appear cautious about linking themselves to AI in a public or strategic way.

Figure 1. Number of Consultancies Mentioning AI

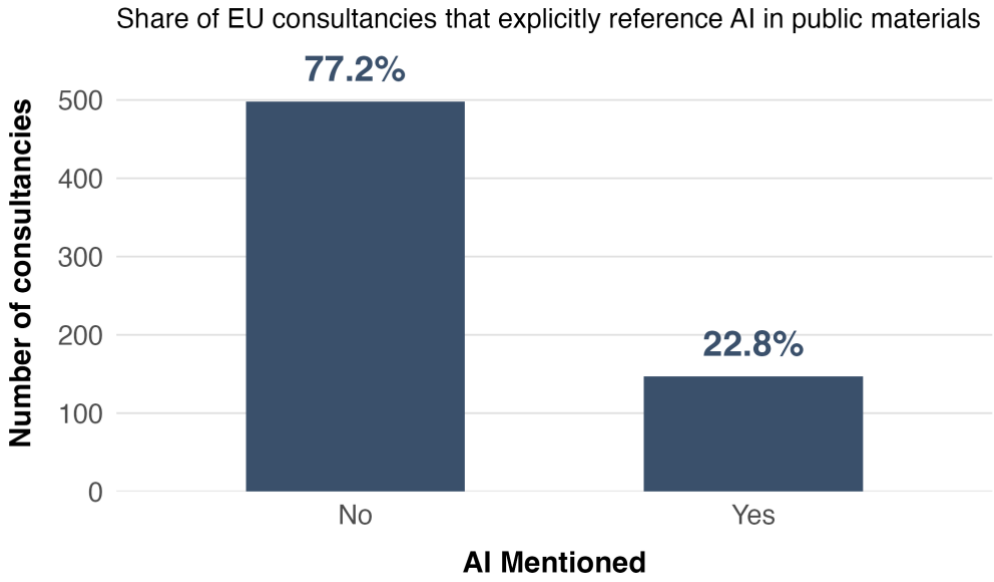

There are several plausible explanations for this limited visibility. Many consultancies are likely still in the early stages of capability-building and may not yet possess the internal skills or technical foundations required to confidently present themselves as AI-ready. Because credibility is the core currency of the consulting sector, firms may also be exercising reputational caution: it is safer to say nothing than to risk overstating expertise that clients could later challenge. Regulatory uncertainty adds another layer. With the EU AI Act approaching implementation, some firms appear to be waiting for clearer guidance before making public claims about AI services or compliance readiness. Finally, market positioning plays a role. Smaller or more traditional consultancies may not see AI as central to their offering,

either because their client base is not yet demanding it or because they are focusing on established service lines.

These factors help explain why AI visibility remains the exception rather than the norm among EU consultancies. The consulting sector appears to be at the very beginning of its public AI transition, with discourse diffusion lagging well behind the prominence of AI in media, policy debate, and regulatory planning.

Looking beneath the European average, clear geographical differences emerge. **Figure 2** shows the share of firms mentioning AI by country for all countries with at least 10 consultancies represented. [2] Austria stands out, with **41%** of its consultancies referring to AI, almost double the EU-wide average. Ireland and France also show high visibility, at **31%** and **30%** respectively, reflecting their strong digital economies and close ties to multinational technology clients. Spain and Italy, with mentions in the mid-20s, indicate an increasingly active AI discourse in Southern Europe. By contrast, Belgium and Germany display notably lower visibility at **19%** and **16%** respectively. This is striking given their large consulting markets and, in Belgium's case, proximity to EU institutions. These lower figures likely reflect different communication strategies rather than absence of activity. Firms in more established markets may prioritise risk management and restrained public messaging over promotional signalling, particularly before regulatory obligations are finalised.

Figure 2. Consultancy AI use by Country

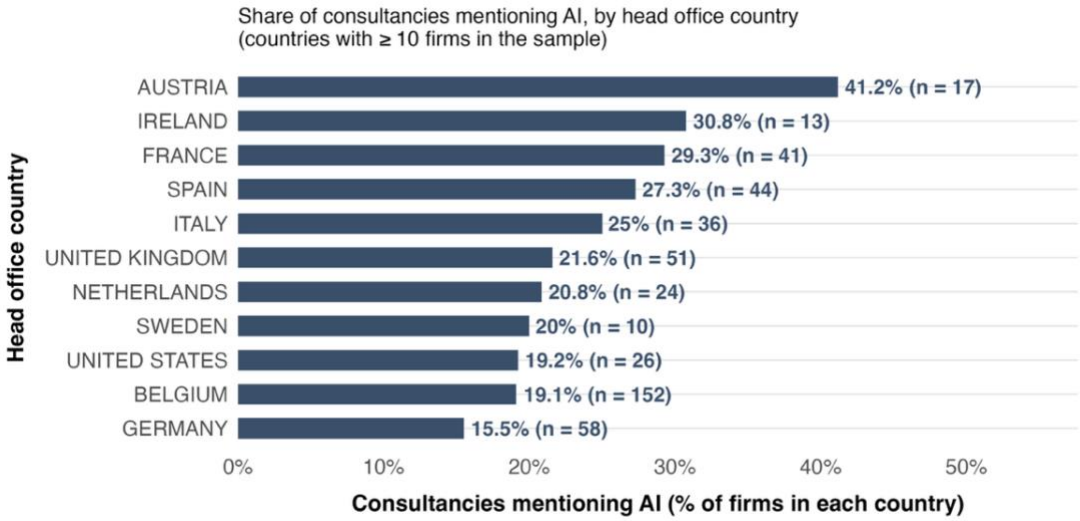

[2] Austrian consultancies appear as the top-ranked country because the metric reports the **share of firms** in each country that reference AI, not the absolute number.

Taken together, these patterns show that AI visibility across EU-based consultancies is uneven and shaped by national context. Smaller and digitally oriented economies such as Austria and Ireland are more vocal in presenting AI as part of their corporate identity, while larger incumbents in Germany, Belgium, and parts of the Benelux region adopt a more cautious approach. Importantly, high visibility does not necessarily signal deeper technical maturity. In many countries, public discourse appears to be advancing more quickly than operational integration.

MAPPING AI MATURITY ACROSS THE SECTOR

We next turn to the question of **AI maturity**, which in refers to the depth, structure, and intentionality of a firm's AI capability. In simple terms, AI maturity captures how far a consultancy has moved from experimenting with isolated tools toward using AI in a consistent, governed, and organisation-wide way. A mature firm does not just "use" AI; it has processes, oversight, training, and integration that make AI part of how the organisation works. Our maturity scale ranges from 0 to 5, where: 0 indicates no identifiable AI activity, 1–2 reflect early experimentation or ad-hoc tool use, 3 represents emerging, structured use with some governance in place, and 4–5 signal organisation-wide adoption supported by formal governance, training, and clear links between internal use and client-facing services.

Figure 3 shows the distribution of maturity across the sector. The average score is 2.3, placing most firms firmly in the early-to-mid adoption range. Only a small group of consultancies reach the highest levels (4 or 5), indicating genuinely strategic and integrated AI capability. These include **CIN Consult GmbH**, the only firm scoring a full 5, and a handful of Level-4 firms such as **DEXAI srl**, **ADS Insight**, **Considerati Holding B.V.**, **Pangea Public Affairs**, and **FortyEight**. These organisations show not only evidence of AI use but also structured governance, client-facing services, and examples of AI being embedded into advisory work.

Figure 3. Consultancy Firms' AI Maturity

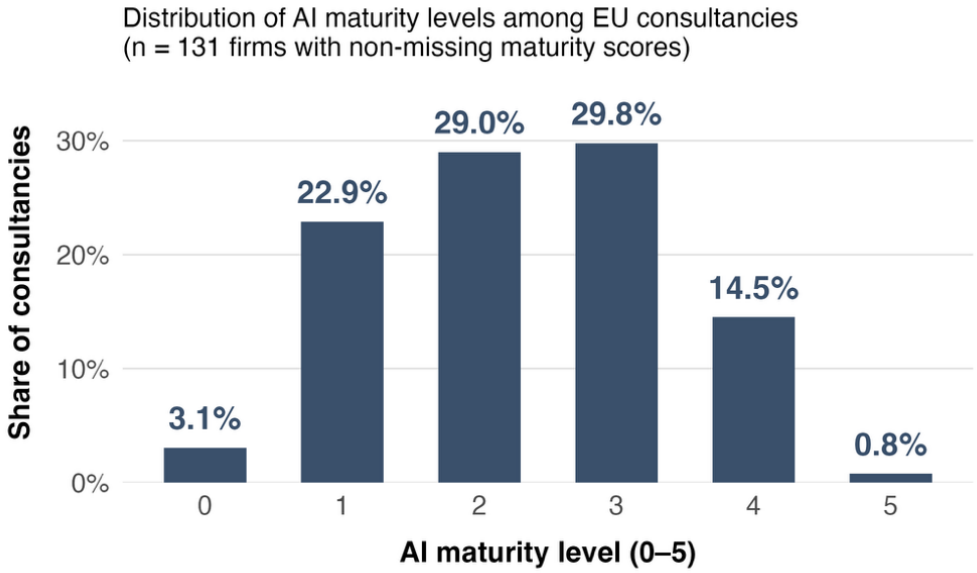

At the other end of the spectrum, a significant proportion of firms score 0 or 1, showing little or no visible AI activity. Examples include **Lighthouse Europe**, **Influera Sverige AB**, **WeeFin**, and **I-DS Koucing & Konsalting**. These firms either make no reference to AI at all or do so only in extremely general terms that lack concrete operational substance. For most of these organisations, AI appears tangential rather than transformational.

This distribution reveals a sector that is still in the process of building foundational competence. Many consultancies appear to be experimenting with discrete tools (predictive-analytics modules, workflow automation, or isolated generative-AI pilots) rather than adopting AI as a coherent, cross-cutting capability. Mature integration remains rare. Few firms demonstrate the governance structures, training pathways, or systematic use of AI required for organisation-wide transformation.

These findings suggest that while AI has clearly captured the industry's attention, it has not yet reshaped day-to-day consulting practice. The consulting sector is still moving from curiosity toward consistency. AI is present, but not yet pervasive.

To understand how consultancies are using AI, we distinguish between two separate variables: **Internal AI Use,** capturing whether a firm mentions using AI in its own operations: for example, using AI for research, analytics, automation, workflow optimisation, drafting, or internal decision support. This variable only records concrete references to INTERNAL PRACTICE, not marketing claims. **Client-Facing AI Services** captures whether a firm explicitly offers AI-related services to clients, such as AI strategy, model development, automation consulting, analytics solutions, governance support, or AI Act readiness assessments These variables allow us to observe whether consultancies "practice what they preach," and whether commercial AI services are rooted in genuine internal capability.

Results are visualised in **Figure 4**. Across the sample, **62.6%** of firms report some form of internal AI use, while **57.3%** present AI-related services to clients. The strong overlap between the two measures indicates that consultancies typically commercialise the capabilities they first adopt internally. Firms that use AI in their own workflows are also those confident enough to advise clients on AI adoption, governance, or compliance.

Figure 4. Internal v External AI Use

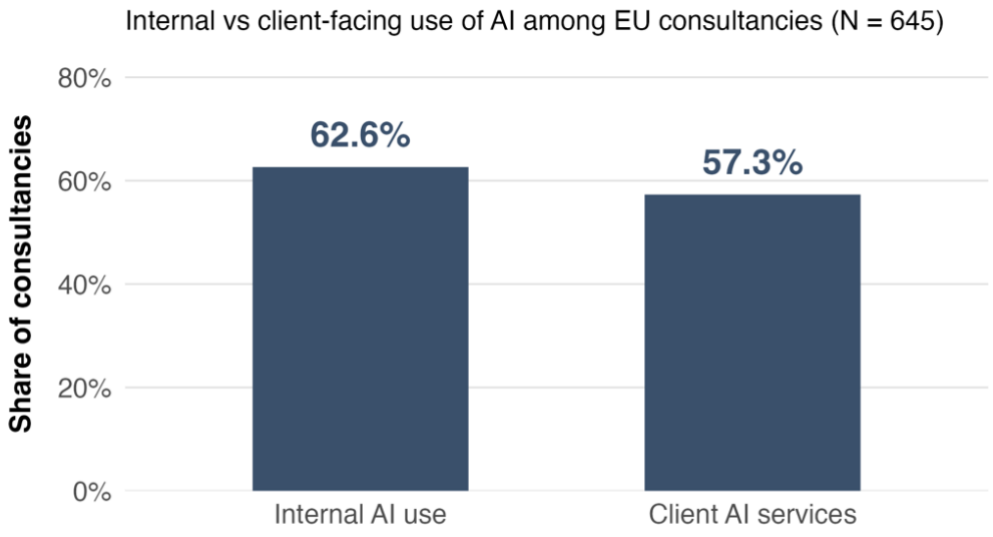

The reverse is equally telling. Very few firms market AI services without mentioning some form of internal AI use. This pattern suggests that consultancies understand the reputational risk of appearing to sell expertise they have not themselves implemented -- particularly in a sector where credibility and trust are central to client relationships.

In short, internal adoption and external service offerings move together. Consultancies that treat AI as a strategic enabler inside the organisation are the same firms positioning themselves as leaders in AI-enabled advisory work. Internal

capability appears to be both a foundation for credibility and a catalyst for commercialisation in the AI consulting market.

Conversely, consultancies that have not implemented AI internally rarely provide AI advisory or technical services externally, underscoring the tight coupling between **internal competence and market offering**. This alignment suggests that firms treating AI as a strategic enabler within their own operations are the ones positioned to lead in the emerging AI-consulting market.

DISCOURSE & FRAMING

This section examines how consultancies publicly talk about AI, what they emphasise, what they downplay, and how their language reflects broader regulatory and strategic pressures. Two types of variables are analysed. The first captures **discourse frames**, meaning the dominant themes or lenses through which AI is described (e.g., as an opportunity, a risk, a compliance issue). The second captures **ethical principles**, meaning the specific moral or governance concepts invoked when firms discuss responsible or trustworthy AI. Together, these measures reveal how consultancies position AI in their professional identity and client-facing narratives.

AI DISCOURSE FRAMES

The discourse-frame variable records any passage where a firm describes AI using thematic language that signals its purpose, value, or implications, for example, describing AI as a tool for efficiency, as a source of risk, or as something requiring compliance. These frames are coded from the text that surrounds AI-related keywords, capturing the firm's **overall storyline** rather than isolated phrases.

As shown in **Figure 5**, a small number of frames dominate the sector's public communications. The most frequent frame is **Compliance** (32.2%), followed by **Opportunity** (22.5%), **Transformation** (17.8%), **Risk** (14.3%), and **Efficiency** (13.2%). This ranking is highly revealing. The prominence of compliance indicates that many consultancies already conceptualise AI primarily through the lens of rules, standards, and regulatory alignment, an unsurprising pattern given the impending implementation of the EU AI Act. Rather than presenting AI as an unbounded innovation frontier, firms often describe it as something that must be governed, controlled, and integrated responsibly into existing structures.

Figure 5. Dominant AI Discourse Frames

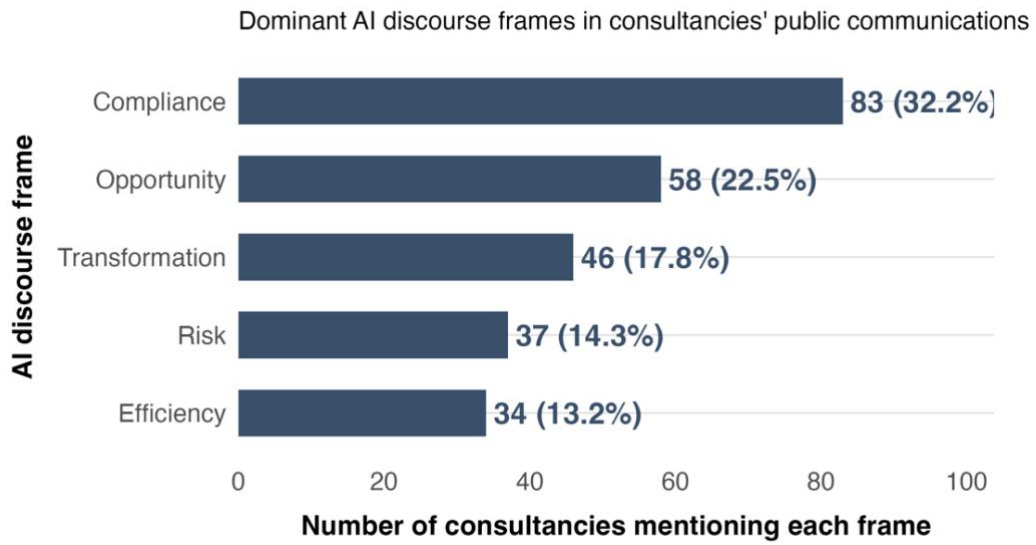

Opportunity and **Transformation** frames highlight AI's potential to open markets, create value, and modernise organisational processes. For example, one Brussels-based consultancy promises to "provide clarity and a clear roadmap for your organisation, turning regulatory challenges into a competitive advantage," explicitly positioning AI and regulation as a source of strategic opportunity rather than constraint. **Risk** frames, by contrast, stress the need for safeguards: another firm emphasises that it "focus[es] on helping clients manage the risks associated with bias, fairness, and data integrity in their AI deployments," underscoring concerns around reliability and responsibility. Efficiency-oriented language adds a further layer, particularly around internal use of AI: one consultancy notes that it "utilize[s] automation tools to streamline tedious research and reporting tasks, freeing up [its] consultants to focus on high-value client engagement and strategic advice," framing AI as a way to optimise workflows rather than reinvent the business model.

These patterns suggest that the consulting industry's AI narrative is pragmatic and institutionally aligned rather than speculative or disruptive. Firms acknowledge AI's potential to create value and transform processes, but they do so in a measured way, tightly coupled to regulatory compliance, risk management, and incremental efficiency gains rather than to radical technological evangelism.

ETHICAL PRINCIPLES REFERENCED

The ethical-principle variable records explicit references to concepts such as transparency, fairness, bias, accountability, and privacy, terms that signal a firm's approach to responsible or trustworthy AI. Only passages explicitly connecting ethics to AI activity are included, ensuring the measure captures meaningful

engagement rather than generic corporate language. The results are presented in **Figure 6**.

Figure 6. Dominant Ethical Principles

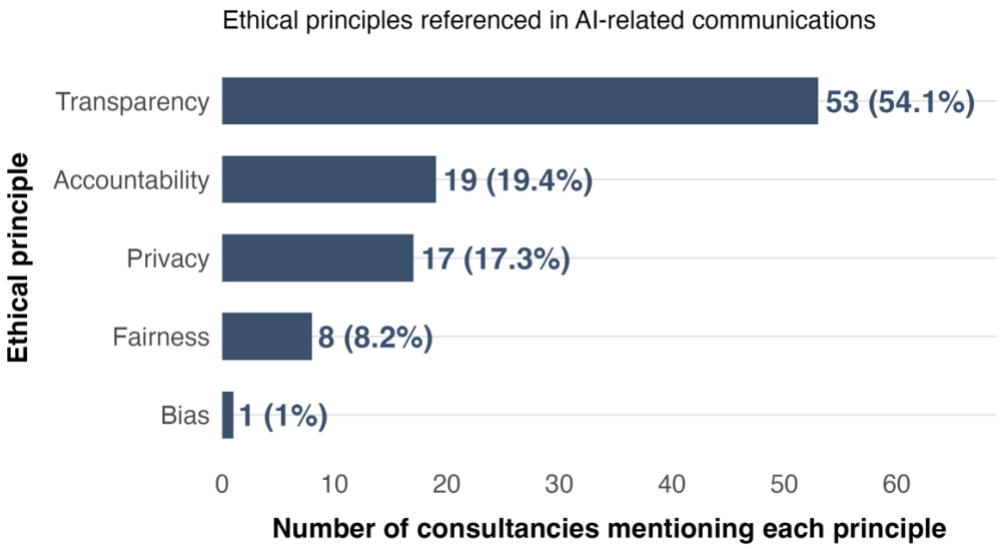

The results reveal a relatively narrow ethical vocabulary. **Transparency** (54.1%) is by far the most frequently cited principle, appearing more than two times as often as **Accountability** (19.4%) and more than three times as often as **Privacy** (17.3%). **Fairness** (8.2%) and **Bias** (1%) appear only sporadically. This concentration around transparency reflects how consultancies align themselves with regulatory and client expectations for explainable, auditable AI systems. Transparency is a safe and institutionally salient principle, one strongly associated with compliance, documentation, and governance. Mentions of Accountability and Privacy reinforce this orientation, drawing on familiar norms from the EU's broader data-protection framework. By contrast, deeper socio-technical concerns such as Fairness and Bias receive minimal attention, suggesting that equity-oriented AI ethics remain peripheral in the sector's public positioning.

Across both discourse and ethics, the overarching pattern is consistent. Consultancies frame AI in ways that are cautious, institutional, and regulation aware. Compliance, transparency, and risk management dominate the narrative, while more expansive or transformative framings play a secondary role. This signals that the industry is positioning itself not as a source of radical innovation but as a trusted intermediary, an actor that helps clients navigate AI safely within well-defined legal and organisational boundaries.

In practical terms, this language serves a dual purpose: it reinforces credibility in a sector where trust is central, and it aligns the industry with the governance principles embedded in the EU AI Act. The result is a discourse that emphasises safe,

explainable, and compliant AI, grounded in the values and expectations of Europe's regulatory environment rather than in technological evangelism or disruption.

AI TECHNOLOGY FOCUS

To understand *how* consultancies are using AI (not just whether they use it) we analyse the specific technologies and subfields they reference in their public materials. The **technology-focus variable** captures any explicit mention of an AI subdomain such as predictive analytics, automation, generative AI, or machine learning. Only concrete references are counted; vague language about "innovation" or "digital solutions" is excluded. This approach shows where firms are directing their technical attention and which kinds of AI are becoming part of their service offerings or internal workflows.

Figure 7 summarises the seven most frequently cited AI technologies. The results show a strong concentration in a narrow set of mature, data-centric tools. **Predictive analytics** (30.9%) and **Data science** (30.9%) account for nearly two-thirds of all mentions, indicating that most consultancies are using AI in ways that build on existing business-analytics capabilities. These technologies are familiar, low-risk, and align closely with traditional consulting work such as performance dashboards, forecasting, and data-driven decision support. They also integrate easily with standard enterprise software and require less specialist talent than more advanced AI systems.

Figure 7. AI Technologies used by Consultancy Firms

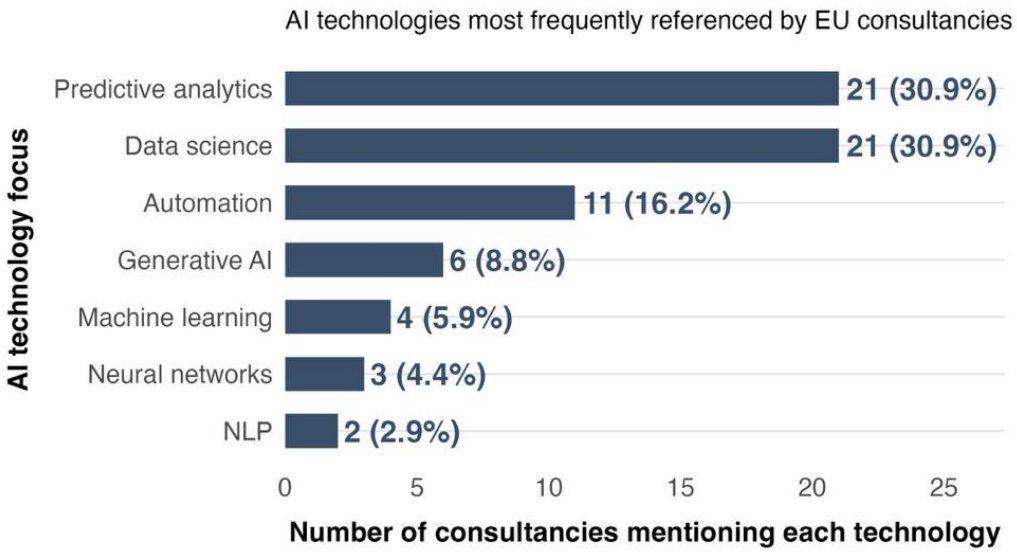

Automation and process AI (16.2%) appears next, reflecting growing interest in workflow optimisation, efficiency gains, and Robotic Process Automation (RPA).

These applications appeal to clients because the benefits (cost reduction, time savings, error reduction) are concrete and measurable. They also carry limited reputational and regulatory risk, making them attractive for consultancies still building foundational experience with AI.

Generative AI (8.8%), despite its visibility in public debate, appears in only 9 percent of firm profiles. This suggests that adoption remains exploratory: consultancies may be experimenting with text generation, summarisation, or ideation tools internally, but few yet present generative AI as a formal service line. Mentions of more advanced techniques, like **Machine learning** (5.9%), **Neural networks** (4.4%) or **Natural language processing (NLP)** (2.9%) are even rarer. These domains typically demand specialised expertise, robust data pipelines, and higher levels of risk and governance, which many consultancies have not yet developed.

What does this mean? The distribution of technological focus highlights a consulting sector **adopting AI pragmatically rather than disruptively**. Firms gravitate toward tools that extend their existing analytics capabilities, fit easily into established processes, and deliver reliable client value without requiring high levels of technical investment or regulatory risk. This aligns with earlier findings: only a small group of **Leaders** experiment with frontier AI models, while most consultancies frame AI as a logical extension of their data-analytics work.

The modest uptake of generative AI and other advanced techniques indicates that the industry remains in a cautious, proof-of-concept phase. Consultants may use generative tools internally for productivity gains, but they are not yet widely commercialising these capabilities. In essence, the technological profile of AI adoption shows an industry investing in incremental improvement rather than transformative innovation, more “enhanced analytics” than “AI revolution.”

GOVERNANCE & EU AI ACT READINESS

This section assesses how well EU consultancies are preparing for the governance and regulatory obligations that accompany AI adoption. Two core variables are examined. **Policy strength** measures the extent to which a firm has established internal structures for managing AI responsibly, like formal principles, governance committees, risk procedures, or documented oversight mechanisms. It is scored on a scale of 0 to 2, where 0 indicates no visible governance activity and 2 reflects clear, structured, and explicit oversight. **EU AI Act readiness** captures whether a firm demonstrates awareness of, or preparation for, upcoming regulatory requirements. This includes references to the EU AI Act, risk classification, conformity assessment, or readiness programmes, scored from 0 to 3. Together, these measures show not only whether consultancies use AI but whether they are building the governance systems required to manage it in line with European expectations.

Across the sector, both measures remain at an early stage. The average policy strength score is 0.51 (out of 2), while the average readiness score is 0.39 (out of 3). These values indicate that most consultancies either lack formal governance structures or have not yet translated general principles into actionable policies. In real-world terms, this means that although many firms may experiment with AI tools, few have established the documentation or oversight practices necessary to support clients in compliance-heavy environments or to satisfy emerging regulatory expectations.

The relationship between the two variables reveals an important pattern. The scatter plot shown in **Figure 8** suggests only a loose connection between policy strength and regulatory readiness. Some firms describe internal ethics or governance principles but make no explicit reference to the EU AI Act. Others mention the Act or compliance work but lack visible internal policies. These two aspects of governance (internal structure and regulatory alignment) appear to be developing along different trajectories. This separation likely reflects the sector's transitionary state: firms may understand that regulation is coming but have not yet built the internal processes required to operationalise it. A small group of consultancies, likely large multinationals or firms with strong legal and risk-advisory capabilities, score highly on both dimensions, indicating more mature and integrated governance systems.

Figure 8. AI Policy Strength and EU AI Act Readiness

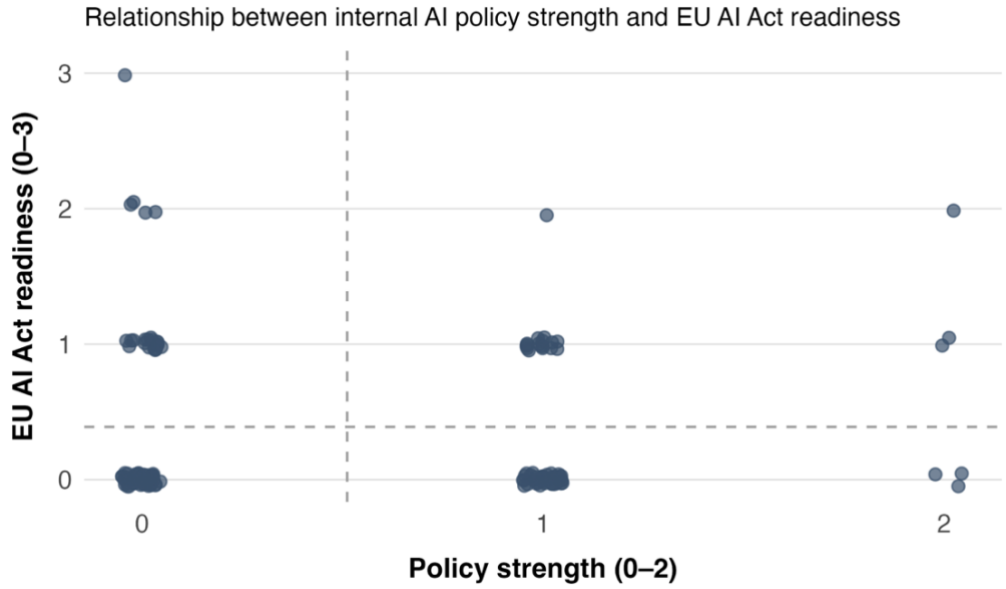

A clearer pattern emerges when governance is analysed alongside AI maturity. **Figure 9** shows that as firms advance in technical adoption, their governance structures strengthen in parallel. Consultancies with no AI activity tend to have no policies at all. By maturity level three, internal governance becomes more visible, with many firms beginning to codify their principles or describe oversight processes.

At the highest maturity levels, firms show both well-developed policy frameworks and meaningful engagement with the EU AI Act. This gradual rise indicates that governance and regulatory preparedness evolve in step with deeper technical adoption. As firms integrate AI into real workflows, they begin investing in the structures needed to manage risk, ensure accountability, and align with regulatory expectations.

Figure 9. Governance by AI Maturity

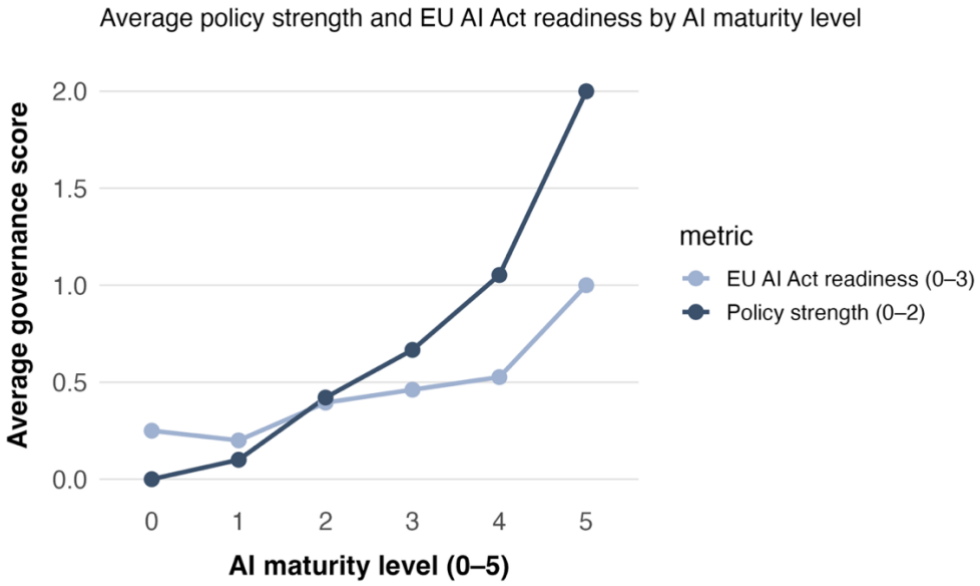

Taken together, these findings suggest that governance in Europe's consulting sector is emerging but far from complete. Most firms are still in the initial phases of building internal oversight, while only a small minority have reached levels that resemble comprehensive, compliance-ready systems. As the EU AI Act moves toward full implementation, the ability to demonstrate credible AI governance will become increasingly central to the industry's market positioning. Firms that lag behind may face challenges in winning client trust, particularly in risk-sensitive or regulated sectors, whereas those that mature early will be better placed to lead in the coming AI compliance landscape.

AI ADOPTION ARCHETYPES: SCEPTICS, FAST-FOLLOWERS, LEADERS, TALKERS

Firms can be grouped into four archetypes (**Figure 10**). The largest group, **Sceptics**, representing around 73 percent of the sample, make little reference to AI and show minimal use in their operations or services. Their communications typically focus only on traditional consulting offerings, with no attempt to signal digital or AI capability. The next group, **Fast Followers** (18%), consists of firms that have begun adopting AI tools and building governance structures but have not yet emphasized AI in their public identity. Their language is pragmatic and tool-focused rather than strategic or rhetorical. For example, one firm highlights that "Merakzy is the first web-based software for complete management and automation of work processes …," positioning AI as a functional enabler rather than a grand narrative. These firms often embed AI internally and use it to enhance operational workflows, even if their public positioning remains conservative.

Figure 10. Four AI Archetypes

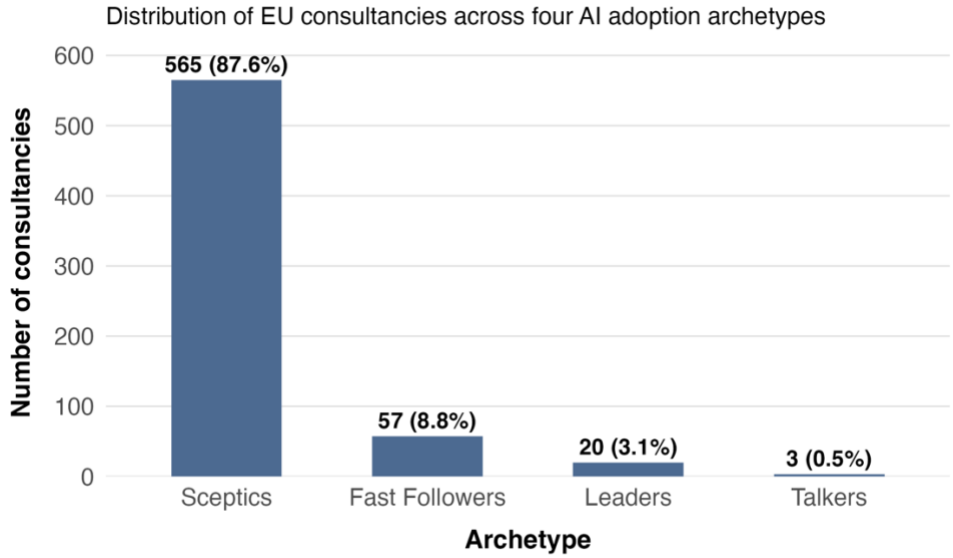

A much smaller set of **Leaders** (around 3%) combine mature internal capability, structured governance, and clear external communication. Their discourse pairs concrete evidence of AI deployment with robust oversight structures. One leading consultancy explicitly states: "We help you establish robust internal governance frameworks, including risk assessment, documentation, and monitoring, ensuring your AI systems are trustworthy and accountable." This integration of capability, governance, and external communication demonstrates credible alignment between what these firms say and what they can deliver.

Finally, the **Talkers** (0.5%) use ambitious or visionary language but provide little evidence of operational integration. Their messaging often includes broad claims about transformation or sustainability without tying these narratives to actual AI use, governance, or compliance. For example, one firm describes "an industrial system that is both restorative and regenerative by design," invoking future-proof

sustainability rhetoric while making no concrete reference to AI deployment. This pattern suggests a reliance on high-level marketing language rather than substantive capability.

The typology presented here illustrates the unevenness of AI integration across the European consulting sector. While a small vanguard is moving decisively toward mature, governed adoption, most firms remain in early or exploratory stages, and a small minority rely on rhetoric in the absence of capability. This variation helps explain why AI visibility alone is an unreliable indicator of true readiness, and highlights the importance of evaluating adoption and discourse jointly rather than in isolation.

WHAT DRIVES AI ADOPTION?

To understand what differentiates the four AI adoption archetypes (Sceptics, Fast Followers, Leaders, and Talkers) we estimated a multinomial logistic regression model using **Sceptics as the reference group**. The model helps identify which organisational characteristics predict a firm's likelihood of moving from early-stage experimentation toward more advanced and structured forms of AI adoption. The predictors reflect both **resources** and **capabilities**. Data are taken from the information provided in the EU Transparency Register. **AI maturity** reflects how deeply AI is embedded in the organisation's own processes, as described earlier. **Policy strength** and **EU AI Act readiness** measure the robustness of a firm's governance framework and its preparation for forthcoming regulatory obligations. **Training coverage** captures whether firms refer to staff being trained or upskilled in AI. Finally, indicators for **internal AI use** and **client-facing AI services** identify whether firms use AI themselves or actively commercialise it.

The results, presented as Relative Risk Ratios in **Figure 11**, show a consistent pattern. **AI maturity is the single strongest predictor** of belonging to any category beyond Sceptics. Firms that have embedded AI into internal workflows or service offerings are far more likely to be Fast Followers or Leaders. Policy strength also plays a powerful role, particularly in distinguishing Leaders from all other groups. Firms with formalised governance structures (codes of conduct, oversight mechanisms, risk processes) are significantly more likely to appear in the most advanced adoption category.

Figure 11. Factors explaining AI Adoption Archetypes

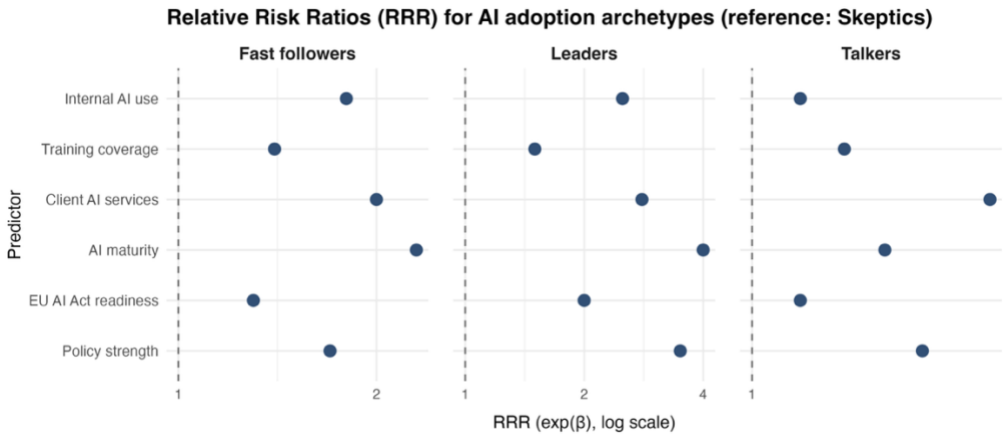

Marginal effects results translate these patterns into intuitive probability shifts. Results are presented in **Figure 12**. We can see that a one-unit increase in AI maturity raises the probability of being a Fast Follower by roughly three to four percentage points and simultaneously lowers the probability of being a Sceptic. Policy strength and EU AI Act readiness also increase the likelihood of being a Leader, though their effects are more modest in magnitude. These findings show that the movement from early-stage experimentation to structured adoption is gradual and cumulative: firms build competence first, then governance, and then external credibility.

Figure 12. Marginal Effects

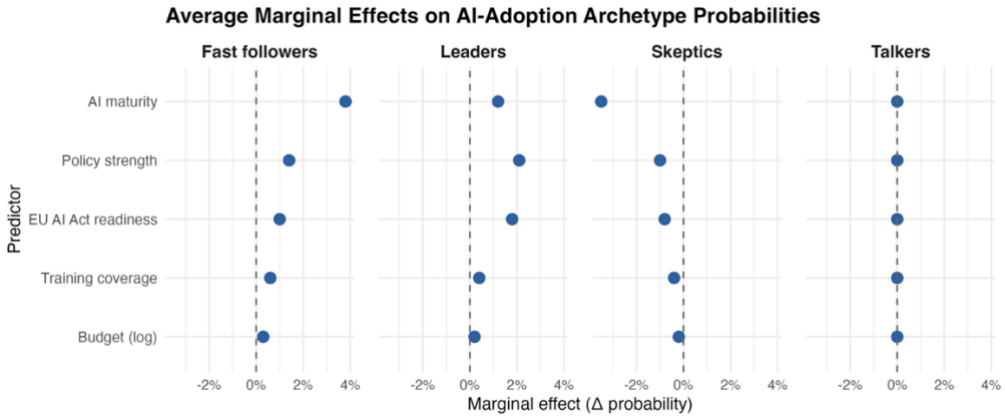

The Talkers (i.e., firms that speak frequently about AI but show little evidence of operational integration) stand out as an exception. Their category shows almost no association with any capability variable. They are neither better governed, nor more mature in their use of AI. This indicates that rhetorical alignment with AI is not a product of institutional strength but of communication choice. In the consulting context, where reputational risks are significant, this group remains extremely small.

Overall, the regression findings reinforce a clear conclusion: **consultancies advance along the AI adoption ladder through internal capability and governance, not through budget size or marketing emphasis**. Leaders are those who combine mature internal systems with credible governance and regulatory readiness. Fast Followers are firms that have begun embedding AI in practical ways but are still building formal oversight. Sceptics remain unconvinced or unprepared. Talkers, meanwhile, highlight the limits of rhetoric: signalling alone rarely substitutes for capability in a credibility-based industry.

These results suggest that competitive advantage in the European AI-consulting market will increasingly depend on a firm's ability to match its public narrative with genuine operational and governance depth. Capability and governance (not hype) are what differentiate the sector's emerging leaders.

CONCLUSION

The findings of this report reveal a consulting sector in the early stages of a major technological transition. AI is clearly on the industry's radar, yet public engagement, operational maturity, and governance readiness remain uneven and often modest. While a small group of consultancies have begun to integrate AI into their internal processes, service lines, and governance structures, the majority are still experimenting cautiously or avoiding explicit association with AI altogether. This dual-speed pattern reflects both the sector's dependence on credibility and the regulatory uncertainty created by the forthcoming EU AI Act.

The evidence shows that **capability (not size, rhetoric, or market presence) is the strongest predictor of advancement**. Firms that actively use AI internally, formalise governance processes, train their workforce, and begin aligning with the EU AI Act are those most likely to lead the next phase of AI-enabled consulting. By contrast, firms that merely reference AI in promotional terms, or rely on general statements of innovation, rarely exhibit substantive adoption beneath the surface.

Europe's consulting sector therefore stands at a crossroads. With regulation tightening and client expectations rising, consultancies will soon be required to demonstrate not just interest in AI but operational competence, documented governance, and verifiable readiness. Those that invest early in structured adoption will gain a competitive edge, while those that remain hesitant risk falling behind a rapidly evolving market. As the EU AI Act enters implementation and organisations across Europe seek guidance, consultancies that can combine technical capability with credible governance will define the standards for trustworthy, compliant, and impactful AI use.

KEY INSIGHTS & RECOMMENDATIONS

Across the analysis, several themes emerge that carry clear strategic implications for consultancies, clients, and policymakers.

AI adoption is visible but shallow.
Fewer than one in four consultancies publicly associate themselves with AI. This reflects reputational caution and suggests that public discourse is advancing more slowly than regulatory and market signals.

Maturity remains modest.
An average maturity score of 2.32 out of 5 indicates that most firms are still in an exploratory or pilot phase. Structured roadmaps connecting experimentation to measurable business value are needed to progress beyond this stage.

Internal capability underpins credibility.
The close alignment between internal AI use and client-facing services shows that consultancies must "practice what they preach." Firms that embed AI into their own operations are the ones most likely to credibly advise clients.

Governance lags behind adoption.
Low policy strength (0.51/2) and AI Act readiness (0.39/3) highlight a growing compliance gap. Consultancies should prioritise establishing AI principles, oversight committees, risk registers, and documentation processes ahead of EU AI Act enforcement.

Geographical variation shapes competitive positioning.
Countries such as Austria, Ireland, and France illustrate how proactive communication on AI can strengthen national competitiveness and attract clients seeking compliance-ready innovation. Core markets like Germany and Belgium adopt a more cautious public stance.

Transparency will become a differentiator.
The prominence of compliance and transparency frames in public discourse shows that clients value clarity over hype. Consultancies that openly disclose their AI governance frameworks will gain trust and commercial advantage.

A dual-speed sector is emerging.
With 73 percent of firms classified as Sceptics and only 7 percent as Leaders, the consulting field risks developing an AI readiness divide. Policymakers, industry bodies, and professional associations may need to support capacity-building initiatives to ensure that smaller firms are not left behind.

About Resonate AI

Resonate AI is an intelligence platform that helps organisations analyse sustainability and non-financial disclosures, benchmark reporting quality, and produce compliance-ready reports. Using retrieval-augmented generation and machine learning, Resonate AI delivers fast, accurate, and transparent insights aligned with leading regulatory and reporting frameworks.

How to Cite

Chalmers, Adam William. 2025. *How EU Consultancies Are Adopting AI: Capabilities, Governance, and the New Competitive Edge*. Edinburgh: Resonate AI.