Unlocking AI's Impact in Italy: moving from AI adoption to real transformation

Insights and lessons learned from an H-FARM Business School study of large Italian companies

Dicembre 2025

H-FARM
BUSINESS
SCHOOL

Susanna Sgarbossa
Lead researcher of the project
and student at UC Berkeley

Pierluigi Fasano
Professor and coordinator of the Competence Centers
at H-FARM Business School

Federico Donati
Program manager of the Competence Centers
at H-FARM Business School

H-FARM
BUSINESS
SCHOOL

In recent years, artificial intelligence has become a constant presence in business language, a symbol of modernity and progress that now seems necessary to mention in every strategic plan. Almost all organizations claim to use it or are ready to do so, but, as highlighted in the MIT *State of AI in Business 2025* the gap between statements and actual results remains significant with 95% of organizations reporting no measurable return despite massive investment.

It is from this observation that the purpose of this white paper takes shape, which aims to systematically analyze how Italian companies are introducing AI into their processes and what organizational conditions are decisive in consolidating its impact. To address this question, we designed and administered an in-depth survey to nine large Italian companies and conducted a series of qualitative interviews with their AI leaders and program owners. Altogether, these organizations employ approximately 28,000 people across the country and represent a diverse mix of structures, levels of complexity, and business sectors. This combination offered a close and multifaceted view of how AI is being interpreted, governed, and translated into organizational practice.

In addition to the Italian sample, the research also included an interview with a U.S.-based technology manager, formerly at Google and now at Meta, whose experience provides a useful point of comparison between the Italian context and more mature AI ecosystems. This perspective helped situate the findings within a broader international landscape and highlight differences in expectations, capabilities, and organizational pace.

9
large Italian companies

1
survey + qualitative interviews

95%
of organizations reporting no measurable return despite massive investment

28,000
people across the country

AI adoption is widespread, but the transformation it promises is still emerging

The combined analysis of survey data and interview insights paints a complex picture: on one hand, companies show a concrete ability to launch use cases in many functions, from marketing to operations, customer service to HR, and to roll out training programs and awareness initiatives that have reached both the corporate population and management levels. In most cases, people participate with interest, recognize the usefulness of the tools, and do not express significant fears about the technology, a factor that contradicts the narrative that cultural resistance is the main barrier to adoption. On the other hand, the transformation generated by AI is proceeding at a slower pace than the spread of initiatives would suggest. Many projects continue to be managed using traditional approaches: defined perimeters, high levels of control, and linear work cycles, which are difficult to adapt to a technology characterized by continuous evolution.

This asymmetry between widespread adoption and incomplete transformation is the hallmark of the current phase. AI is present and producing tangible results, yet it is evolving at a speed that far outpaces the ability of organizations to absorb and institutionalize it.

As a consequence, the technology is not yet fully incorporated into decision-making processes or into organizational learning models. Initiatives work, but often remain isolated; people are highly motivated, but the infrastructure needed to support continuous evolution is still emerging. The technology demonstrates its value, but the organizational structures that should enable it to scale and endure are not yet fully adapted, a gap that is understandable given how rapidly the AI landscape continues to accelerate.

In this transitional scenario, the contribution of this white paper is to highlight what is happening within Italian companies: an extensive experimentation phase that has generated significant results but now requires a leap towards more coherent, stable models capable of accompanying AI in its evolutionary pace. What will determine the direction of the coming years will not be the number of projects launched, but the ability of organizations to transform these initiatives into a continuous learning process, in which technology, governance, and culture can proceed in alignment.

The following sections reconstruct this path, analyze the evidence gathered, and identify the factors that enable the transition from adoption to effective transformation, offering useful insights for companies that, despite having already embarked on the path of AI, need to consolidate its impact and turn it into lasting value, as well as for those that are only now beginning their journey and seeking to build the foundations for effective adoption.

Active Engagement, Balanced Teams, and Leadership Support Drive Adoption

Our analysis shows that the introduction of artificial intelligence has already produced measurable effects within organizations. The benefits observed are particularly evident in day-to-day activities, where the technology reduces time, increases output volume, and lightens the operational load. In several cases, the introduction of AI has generated tangible efficiency gains, reducing the time required to complete routine activities and streamlining internal processes. Other teams have reported the ability to handle higher workloads with the same resources, or an improvement in output quality, including a more consistent flow of commercial leads supported by analysis and content-generation tools.

These specific examples, although different in nature and scale, show a consistent pattern: in activities where AI is effectively integrated and used regularly, the impact is clear and immediate. These are not structural transformations, but concrete improvements that demonstrate how AI is already capable of producing operational value.

Crucially, leadership involvement has been a universal factor in AI adoption so far. In every organization, top executives have played a role in championing AI initiatives, often acting as visible sponsors at the highest levels, such as board-level advocates or dedicated C-suite champions. About half of the companies kicked off their AI programs through a clear top-down mandate, signaling management's commitment from the outset. The others describe a more hybrid path: certain AI use cases emerged organically from the ground up, but then scaled only once they received strong endorsement and resources from leadership. Tellingly, none of the cases were purely bottom-up or isolated initiatives; alignment with senior leadership proved essential to advance AI implementation beyond the pilot stage.

≈50%

signaled management's commitment from the outset

≈50%

use cases emerged from the ground up but scaled only with strong leadership endorsement

0%

none of the cases were purely bottom-up or isolated initiatives without alignment with senior leadership

This leadership support is reinforced by a notable and often underestimated dynamic: employees are not resisting AI, they are embracing it with enthusiasm.

Across almost all of the companies analyzed, teams reported curiosity, willingness to engage, and a pragmatic attitude toward experimenting with new tools. Cultural resistance, a commonly feared barrier, has been limited and isolated, far less significant than challenges related to skills, governance, or technology. In short, people are ready to adopt AI, and leadership is backing them, the question is whether the governance structures can keep pace.

Many firms have rolled out engagement initiatives, pilot projects, listening sessions, and feedback cycles to integrate AI tools into real processes and refine them through employee input. Several organizations have created internal communities for AI knowledge-sharing and introduced roles, such as AI Champions, to support teams throughout adoption. Training programs further reinforce this momentum: companies are investing in mandatory AI courses, hands-on workshops, and initiatives to bring even less tech-oriented staff closer to AI. Where the benefits are immediately visible, enthusiasm rises quickly; where the context is more complex, adoption progresses more slowly, requiring continuous guidance and support.

>20%

change-management personnel

≈50%

of the companies invest predominantly in technology and tools

In building AI capabilities, however, the composition of AI functions still leans heavily toward technical roles, with strategic and change-management expertise added only gradually. In most companies, the majority of people involved in AI teams belong to data and technology functions, while change-management personnel often constitute less than 20% of the team. Budget allocations mirror this imbalance: about half of the companies invest predominantly in technology and tools, with fewer dedicating equal resources to the strategic and organizational shifts required for AI to scale.

Measurable Gains Are Emerging, but Measurement Systems Are Not

If there is a brake on this transformation, it lies in the difficulty of measuring AI's impact beyond individual projects and translating results into an enterprise-wide perspective. Today, most companies monitor AI performance at the level of specific processes or use cases, relying on local indicators that capture improvements within a single workflow but do not connect to a broader organizational view. As a result, measurement remains fragmented.

Only a small number of companies have begun to develop holistic, program-level measurement systems capable of assessing the impact of AI adoption in a truly comprehensive, end-to-end way.

Without this overarching perspective, even successful initiatives risk remaining isolated, and organizations struggle to identify where AI could generate additional value or how local learnings can scale into broader transformation. Several respondents explicitly acknowledge that their performance frameworks are still in an early stage, creating ambiguity about progress and weakening the ability to steer AI as a coherent portfolio rather than a collection of experiments.

A related hurdle is the classic question of ROI. One of the most frequent obstacles cited is the difficulty of justifying AI investments in concrete financial terms. Many AI projects do not have immediate or easily quantifiable payoffs, leading to uncertainty or skepticism from stakeholders who must allocate budgets. This emphasis on short-term justification can become a bottleneck: it slows decision-making and can cause promising experiments to pause if their value cannot be expressed in the metrics the organization is accustomed to. Notably, the challenge of proving ROI emerged far more often than concerns about employee resistance or lack of sponsorship. The combination of fragmented measurement and ROI uncertainty remains a central issue to resolve before AI can progress from opportunistic use cases to a sustained, organization-wide transformation.

>

most companies monitor AI performance

<

only a small number of companies measure holistically

ROI »

employee resistance or lack of sponsorship

Governance and Organizational Complexity Slow Down Scaling

100%

organizations that have introduced some form of preliminary governance

<

organizations that have elevated AI oversight to a dedicated C-level role or specialized function

Another critical factor tempering the pace of AI transformation is the way companies manage and govern their AI efforts. Across all organizations, some form of governance has already been introduced: dedicated committees, cross-functional task forces, and formal roles such as AI Champions have become common tools for coordinating initiatives and supporting deployment. Some companies have embedded AI governance within existing structures most often under IT or innovation. While this approach integrates AI into familiar organizational frameworks, it also means that initiatives may receive competing signals or priorities depending on the hosting function. In these cases, coordination can become slower or more fragmented, with delays from lengthy reviews or ambiguity about who truly owns AI decision-making.

Advancing AI requires alignment across functions with different priorities, processes, and skill sets. As a result, governance is precisely where digital transformation meets the limits of existing structures: companies must navigate silos, refine roles, and often rethink decision-making models.

While all firms have introduced preliminary governance mechanisms, only a few have taken the next step of elevating AI oversight to a dedicated C-level role or specialized function. Notably, several organizations plan to make this shift as their AI programs mature, recognising that current arrangements, based on committees and dispersed responsibilities, are not sufficient to support enterprise-wide adoption in the long term.

Before an AI solution can even be deployed, companies have learned that they must define clear internal guardrails. Almost all organizations have drafted internal guidelines that specify how AI can be used, which data can be processed, which public tools are permitted, and what risks must be mitigated.

This regulatory framework is essential for accountability, confidentiality, and security, but its very strength introduces friction: approvals become longer, experimentation slower, and each new use case must undergo a sequence of reviews from legal, security, and IT teams. In practice, AI adoption is unfolding within narrow, carefully controlled boundaries that protect the organization but also slow the kind of rapid, iterative experimentation that AI would require.

In parallel, companies continue to struggle with foundational issues around data and skills. Data governance remains uneven: ensuring data quality, accessibility, and enterprise-wide management is still a major challenge. Skills gaps, particularly in advanced data and AI engineering, further slow progress, forcing companies to rely on external expertise or to build internal capabilities through training programs.

Finally, the way many companies design AI projects reflects the logic traditionally used for technology initiatives: phased planning, rigid scopes, and responsibilities concentrated within technical teams. While this approach allows use cases to be deployed in an orderly manner, it limits the continuity of change. AI is not just a set of tools but a technology that reshapes how people work, coordinate, and make decisions. It requires frequent adjustments, iterative cycles, and rapid learning, features that traditional project models are not designed to support. Until governance structures, data foundations, and organizational processes evolve to match this pace, AI initiatives will continue to advance, but not yet transform.

»

organizations that have drafted internal guidelines/guardrails specifying AI use, data, and risks

>

many companies design AI projects using traditional logic

"AI Enablers & Blockers"

What accelerates, and what slows down, AI transformation in Italian companies.

Enablers

- Strong and visible leadership sponsorship
- High employee curiosity and willingness to experiment
- Structured focus on change management & training
- Use cases already delivering measurable operational benefits
- Early governance structures in place
- Integration of AI-supporting roles (e.g., AI Champions)
- Increasing attention to data quality and accessibility

Blockers

- Pressure for immediate and quantifiable ROI
- Heavy reliance on external vendors
- Fragmented measurement systems
- Overly restrictive and bureaucratic policy frameworks
- Data governance not fully ready for AI at scale
- Organizational structures not yet adapted to AI's pace
- Traditional, linear project management approaches

A Look at U.S. AI Workflows: A Useful Contrast

A contrasting perspective emerges from the interview with a U.S.-based technology leader, formerly at Google and now at Meta, which offers a concrete view of what happens inside more mature ecosystems. The most striking difference is not technological sophistication but the way AI is positioned in day-to-day work. At Meta, AI is embedded from day one: Metamate, trained on the company's internal documentation, and DevMate, capable of assisting across the entire code-to-production workflow, are introduced during onboarding as standard tools rather than experimental features. The expectation is immediate: employees are encouraged to use these systems to speed up their work, reduce friction, and focus their time on higher-value activities. At Google, this shift is even more explicit: AI use has been formally integrated into the yearly performance review, signalling that adopting and mastering AI is no longer optional but part of what defines individual contribution.

What stands out in the interview is the simplicity and directness of the message: while in many European companies the first conversations about AI tend to revolve around ethical risks, compliance constraints, or the need to draw careful boundaries before experimenting, the approach described here is almost the opposite. Employees are exposed to the tools immediately, they see concretely what they can do, and they are encouraged to test them in the flow of their daily work rather than in isolated training environments.

When someone uncovers a more efficient workflow, a useful prompting technique, or a way to automate recurring steps, this insight is not treated as personal advantage but as something to circulate to colleagues, since professional progression is tied to contributing to the productivity of the entire team.

In this environment, learning does not happen in predefined sessions but through a continuous cycle in which people try, adjust, teach one another, and incorporate what works into their routines without excessive procedural friction.

A second relevant element concerns the freedom employees have in choosing their tools. At Meta, individuals can work with all major external models, ChatGPT, Claude, Gemini, alongside internal systems, and are explicitly encouraged to use whichever solution produces the best outcome for the task at hand. This openness reinforces the same pragmatic orientation: what matters is the quality and speed of the result, not the vendor behind the model. It also creates a context in which teams evolve at the same pace as the technology, comparing outputs, adopting improvements quickly, and refining their practices as models advance. Rather than anchoring decisions to a single standard or a predefined technological perimeter, the focus is on enabling people to discover what truly works in practice and to share those insights widely, allowing the organization to learn faster than any formal process could manage.

What this perspective highlights is not a model to replicate, but a possibility: when organizations treat productivity gains as an opportunity to broaden what teams can build rather than something to control, experimentation becomes a driver of progress. Accessible tools, clear expectations, and shared incentives enable teams to test ideas quickly and turn small discoveries into broader practices. The interview shows how this mindset allows structures to evolve with the technology instead of reacting to it, suggesting that organizations capable of aligning governance and learning rhythms with the pace of AI will be best positioned to transform early pilots into lasting, scalable impact.

From Adoption to Maturity: The 3P Path to Sustainable AI

The research shows that Italian companies are now moving beyond the initial adoption phase and entering a more delicate stage: one in which technology is no longer the key variable.

Use cases have been launched across functions, employees have begun to integrate AI into their daily routines, and the first benefits are already visible in efficiency, workload reduction, and output quality. Adoption is widespread, leadership sponsorship is strong, and the organizational energy around AI is real.

Yet this momentum has not yet become structural. Early progress remains local, often limited to specific teams or processes, while the organizational conditions needed to scale it across the enterprise are still emerging. The asymmetry highlighted throughout this research is clear: AI is advancing faster than the structures meant to support it. Departments experiment; people learn; technology evolves; but the governance, processes, and data foundations still follow logics designed for stable, linear projects rather than for a continuously changing technology.

This is where continuity becomes the decisive variable. Transformation accelerates when AI is supported by adaptable processes, clear rules that evolve with use, sustained investment in skills, and the ability to learn from daily interactions with the tools. Conversely, it slows down when initiatives remain isolated, when training is episodic, when governance becomes more rigid than the technology it intends to regulate, or when policies, although necessary, are designed to minimize risk rather than maximize learning.

The potential of AI will be unlocked in organizations capable of simplifying decision-making processes, improving coordination between functions, strengthening data infrastructure, and redesigning governance to match the speed of technological change.

To interpret this turning point, we propose the **3P framework: People, Problems, Platforms**, a synthesis derived from the insights gathered through the survey, the interviews with AI leaders, and the contribution of external experts involved in the research. Rather than being a theoretical model, the 3P framework reflects the practical lessons shared by practitioners and captures the core conditions that companies must strengthen to progress from initial adoption to true AI maturity.

Each dimension represents one of the structural shifts required to move from experimentation to scale. Together, they offer a guiding path for organizations: how to introduce AI effectively, how to consolidate early successes, and how to transform momentum into a stable, enterprise-wide capability.

Below, the 3P framework and its corresponding best practices outline the organizational foundations that can help convert Italy's fragmented progress into a more continuous, systemic form of transformation.

A practical framework to help organizations with the accelerating pace of AI.

People

Turning Enthusiasm into Organizational Capability

The main challenge for AI adoption is the transition from leadership enthusiasm to structural maturity, requiring robust governance and stable internal skills.

Problems

From Local Proofs of Concept to Scalable Learning

AI maturity requires shifting from isolated pilots to strategic, scalable programs, driven by holistic measurement tied to business goals.

Platforms

Building the Foundations for AI to Evolve

The structural foundations are the main bottleneck, creating a speed gap that demands a shift to enabling governance and a dedicated organizational backbone for scaling AI.

People:

Turning Enthusiasm into Organizational Capability

If there is one insight that clearly emerges from the survey and the interviews, it is that *people are not the obstacle to AI adoption in Italy*. Employees approach AI with curiosity and a pragmatic willingness to experiment, especially when benefits are visible in their daily work: shorter execution times, fewer repetitive steps, and better-quality outputs. The research shows a workforce that is ready to engage and learn, as long as the organization provides clarity, guidance, and accessible tools.

Leadership has been fundamental in creating the conditions for this openness. In almost all companies analyzed, AI initiatives were legitimized at the highest levels; yet sponsorship alone is not sufficient to sustain momentum over time. This sponsorship has been essential to move early experiments forward, but it also highlights the need to carefully align expectations. The issue is not that AI is unable to deliver on its promise; rather, expectations can sometimes be shaped around forms of value that are not immediately measurable or that rely on direct financial returns which certain initiatives are not designed to generate. Some benefits of AI adoption are indirect, cumulative, or linked to improvements in quality, speed, or decision-making—areas where calculating a direct ROI is inherently difficult.

The transition to structural maturity requires organizations to move from *leadership enthusiasm* to *leadership architecture*: clear ownership, shared expectations, dedicated roles, and governance mechanisms that ensure continuity beyond individual initiatives.

Another emerging theme is the need to internalize skills. While early experiments often depend on external vendors, long-term success requires a shift towards internal capability-building: training, cross-functional roles, and the ability to translate operational problems into AI opportunities. Relying too heavily on outsourcing may offer short-term acceleration, but it limits long-term autonomy, slows organizational learning, and weakens ownership of the technology.

Moreover, both the need to calibrate leadership expectations and the opportunity to harness employees' enthusiasm point in the same direction: sustained investment in skills, internal talent development, and cultural evolution is essential. Only by strengthening these internal foundations can organizations convert early interest and top-down commitment into a mature, resilient capacity to adopt and scale AI over time.

Crucially, capability-building is also what enables organizations to manage expectations coherently. By raising AI literacy across teams, companies can create a shared language and a shared understanding of value, reducing the risk of misaligned ambitions and enabling more grounded, sustainable decision-making.

Best Practices for People

↳ Ensure strong leadership commitment and aligned expectations

AI adoption requires visible and sustained commitment from the top. Establishing a clear and shared vision, articulated and sponsored by a dedicated C-level leader, ensures that objectives and expectations are aligned across the organization. This shared understanding helps prevent overestimations and conflicting priorities, while creating the coherence needed to translate early experimentation into long-term capability. A strong C-level sponsor anchors the vision, reinforces alignment, and keeps the organization focused on realistic, strategically meaningful outcomes.

↳ Invest in internal capabilities rather than relying solely on external support

External providers help accelerate early steps, but long-term maturity requires internal ownership. Building cross-functional expertise, people who understand both processes and technology, allows organizations to scale learning instead of repeating the same foundational steps with each new project.

↳ Link training with daily practice

Training generates impact only when it is connected to real tasks. Short, applied learning formats such as hands-on sessions, practical exercises, internal guides, and team-level micro-training help individuals incorporate AI naturally into their work. This approach avoids the gap between course knowledge and real competence and creates a rhythm of continuous learning instead of episodic education.

↳ Make internal communication part of the project

People's confidence in AI depends on what they know, what they expect, and how clearly the organization communicates benefits and limitations. Regular, transparent communication reduces uncertainty and fosters alignment, especially when adoption involves changes to workflows or responsibilities.

↳ Continuously observe the market and invest in research capacity

The speed of AI evolution makes it necessary to monitor technological trends systematically. Companies that dedicate time and resources to technology scouting, or establish small internal research teams, stay aligned with global developments.

Problem:

From Local Proofs of Concept to Scalable Learning

Italian companies often began their AI journey with small, well-contained use cases. These early pilots were essential to reduce risk, learn quickly, and demonstrate initial value. Yet, what emerges from this research is that remaining anchored to "small problems" becomes a bottleneck. The organizations that move fastest are those that treat small experiments as catalysts: they validate what works rapidly, extract the insights, and scale decisively. This shift, from isolated pilots to structured learning, is the real marker of maturity.

Scaling requires more than replication; it requires the ability to learn from practice. AI behaves differently once embedded in real processes, and many insights emerge only after deployment: data inconsistencies, user adaptations, unexpected opportunities, or inefficiencies in workflow design. Organizations that institutionalize short feedback loops are able to convert these insights into improvements before issues crystallize.

Likewise, measurement is not a late-stage check but a necessary ingredient from the start. Companies often assess performance at the level of individual projects or processes, but this is not sufficient to guide scaling. Effective problem selection requires a clear link to strategic objectives: AI initiatives should be framed and evaluated not only as technical experiments, but as contributions to the company's broader priorities. Establishing measurement criteria tied to these strategic goals from the outset enables organizations to learn, compare, and ultimately scale what truly matters. What is needed is a holistic performance-management framework that is defined from the earliest stages and grows with the program. Such a framework enables organizations to build confidence in AI's value, both for business stakeholders and for the leadership team whose expectations must be continually calibrated.

Best Practices for Problem

↳ Start small, scale fast

Starting small reduces initial uncertainty, but staying small limits impact. Organizations that progress transform early wins into broader improvements by expanding validated use cases across teams and processes. The shift from pilot to process must be intentional and timely to avoid fragmentation and to convert localized gains into structural change.

↳ Give people the opportunity to experiment directly on real tasks

AI reveals its true value only when tested in real operations, not in controlled simulations. Allowing teams to experiment with genuine workflows, while providing support to ensure safe use, accelerates learning and uncovers opportunities that would remain invisible in artificial environments. This practical exposure also helps identify non-intuitive configurations and workflow adjustments.

↳ Encourage short cycles of listening and adjustment

Critical information surfaces within the first weeks of adoption. Organizations that collect feedback early and adjust rapidly avoid the delays associated with long, linear project structures. These short learning loops help stabilize use cases and create a shared understanding of how AI should evolve within the process.

↳ Make measurement part of the design

Reliable measurement frameworks allow companies to compare outcomes across pilots, evaluate ROI, and justify scaling. But more importantly, a robust, program-level measurement model must make it possible to understand how AI is contributing to the organization's strategic objectives. When measurement is embedded into the design from the start, it strengthens strategic alignment, helps calibrate leadership expectations, provides transparency on progress, and builds confidence in AI's long-term business value.

Platform:

Building the Foundations for AI to Evolve

While enthusiasm and adoption are strong, the structural foundations of AI: data, governance, architecture, and integration, remain the most significant bottlenecks. Many companies operate with legacy systems, inaccessible data, slow authorization processes, and technical environments that were not designed for tools that evolve weekly. As a result, initiatives advance, but they do not yet consolidate into a capability that grows autonomously. The gap between the speed of AI and the speed of organizational infrastructure is now one of the defining challenges of this phase.

Governance reflects a similar dynamic. Policies, guidelines, and approvals are essential for protection, but when they become more rigid than the technology they regulate, they slow down learning. Mature ecosystems, such as the one described by the U.S. interviewee, show how organizations maintain safety without sacrificing evolution, allowing teams to experiment with real tools while keeping sensitive data protected. For Italian companies, the shift is to move from protective governance to enabling governance: still rigorous, but designed to support experimentation, scaling, and continuous adaptation.

Organizational structures represent the third foundational element of the platform. To scale AI effectively, companies need to build a true organizational backbone, one that goes beyond distributed responsibilities and coordination mechanisms. This backbone should include a dedicated AI function with clear C-level ownership, repeatable processes, and long-term capability-building. Without these structural foundations, AI risks remaining a collection of isolated initiatives rather than evolving into a scalable enterprise capability.

Best Practices for Platform

↳ Invest in data quality and accessibility

High-quality, well-governed, accessible data is the foundation of scalable AI. Improving data structures, clarifying ownership, and enabling controlled access allow organizations to multiply the impact of AI without requiring more complex models or additional tools.

↳ Use policies as a framework, not a brake

Policies should define clear boundaries while allowing teams to experiment safely and iterate quickly. Effective governance specifies what cannot be done, but leaves room for trial cycles, controlled testing, and the evolution of solutions through usage.

↳ Avoid treating AI as a traditional IT project

AI initiatives lose momentum when managed through rigid, sequential phases. Iterative releases, cross-functional teams, and multidisciplinary collaboration enable solutions to adapt and grow with user insights and evolving business needs. AI must be treated as a capability in motion, not a fixed implementation.

↳ Evolve the organization at the same pace as AI

As AI scales, the organization must scale with it. Early phases can be supported by committees, shared responsibilities, or coordination across existing functions, but these mechanisms quickly become insufficient as the volume and strategic relevance of AI initiatives grow. To match the speed at which AI evolves, companies need to build organizational structures that are equally agile.

Agility as a Strategic Capability: Organizations Must Be Prepared to Run at the Speed of AI

AI will continue to evolve faster than any organizational structure can naturally adapt. This tension is not a limitation, it is the defining condition of the current phase. The companies that will succeed are not those that try to match the speed of the technology, but those that **design themselves to move faster**: faster in learning, faster in adjusting, faster in scaling what works and abandoning what does not.

Preparing the organization to run at the speed of AI requires a combination of agile **skills**, **systems**, **and structural agility**. It means investing in people and internal capabilities so that teams can understand, evaluate, and apply new tools as they emerge. It requires continuously observing the market, translating external signals into informed decisions, and experimenting quickly with new solutions. It also demands the ability to scale promising initiatives with discipline, supported by strong governance mechanisms and holistic measurement frameworks that clarify how AI contributes to strategic objectives.

Most importantly, organizational agility depends on the willingness to revise and adapt strategies, structures, and roles continuously. As AI accelerates, organizations must evolve in parallel refining their operating models, reinforcing their capabilities, and realigning expectations across leadership and teams. In this sense, agility is not an operational feature but a strategic capability, one that brings together the three pillars highlighted throughout this report: People, Problem, and Platform.

Unlocking AI's Impact in Italy: moving from AI adoption to real transformation.
Insights and lessons learned from an H-FARM Business School study of large Italian companies

Susanna Sgarbossa, Pierluigi Fasano, Federico Donati

www.bs.h-farm.com
business.school@h-farm.com

H-FARM
BUSINESS
SCHOOL