Leaders use it First: How Leadership Role Modeling Shapes the Sources and Grounding of AI Self-efficacy in Knowledge-intensive Roles in Morocco

Amine Zenjari, Oumnia Sidqui and Amal Zaki
Groupe ISCAE, Casablanca, Morocco

azenjari@groupeiscae.ma
osidqui_doct27@groupeiscae.ma
azaki_doct27@groupeiscae.ma

Abstract: Artificial intelligence is reshaping knowledge-intensive work, yet organizations repeatedly confront a gap between technical access to AI tools and workers' confidence in using them. This exploratory qualitative study asks how leadership role modeling shapes the sources and the grounding of AI self-efficacy among knowledge workers, and how that self-efficacy relates to job performance. Drawing on eleven semi-directive interviews with knowledge workers in consulting, finance, technology, academia and medicine, coded in two verified cycles (446 references, 32 nodes), we bring the Technology Acceptance Model (Davis, 1989) into dialogue with Bandura's (1997) Social Cognitive Theory. Four patterns emerge. Formal training was rare and, where mandatory, experienced as box-ticking; self-efficacy was built through self-directed enactive mastery. Visible AI use by managers and senior professionals operated as a signal of social legitimacy, distinct from managerial sponsorship, which on its own left practice unchanged; where senior use was absent or discreet, use contracted to a narrow perimeter of low-stakes tasks. Peers, not leaders, were the main channel of vicarious know-how. Finally, the grounding of self-efficacy, in a competence the worker retains or in the tool's availability, tracked the framing managers gave to AI and governed the durability and calibration of confidence rather than the size of the reported performance gain. We revise the mechanism chain accordingly and offer it as propositions for quantitative testing; the contribution is to reposition leadership as a cognitive, not merely behavioral, enabler of AI adoption in an under-studied emerging-economy setting.

Keywords: Leadership role modeling, AI self-efficacy, Sources of self-efficacy, Social legitimacy of use, Mastery appraisal, Morocco

1. Introduction

Generative AI has compressed into a few years a transformation of knowledge work that earlier office technologies took decades to accomplish. Field studies now document sizeable gains in productivity and quality among professionals who delegate drafting, synthesis and exploratory reasoning to large language models (Brynjolfsson et al, 2025; Dell'Acqua et al, 2023; Noy and Zhang, 2023). One result is especially telling: Noy and Zhang (2023) found that exposure to generative AI raised not only output but participants' self-efficacy, a sign that what is shifting is capability belief and not merely performance. For organizations, accordingly, the strategic question has moved from whether to adopt AI to a more puzzling one: why adoption stalls even when the technology is free, capable and demonstrably useful.

The conventional response is to invest in access and training, yet provisioning licences and running sessions does not reliably yield confident, habitual use. The decisive construct is not whether a worker understands how an AI tool operates but whether they believe they can use it effectively for their own tasks: what Bandura (1997) called self-efficacy, and what information-systems research later operationalized as computer, and now AI, self-efficacy (Compeau and Higgins, 1995; Marakas et al, 1998). Even where AI integration is found to raise that belief, as in a recent survey of 441 employees in the United Arab Emirates (Shishakly, Khattak and Abukhait, 2026), the instruments capture confidence in working with AI rather than anything about the capability the worker retains, and leadership enters, if at all, as a contextual condition rather than a mechanism.

Two literatures speak to this gap but rarely converse. The Technology Acceptance Model (TAM; Davis, 1989) and the Unified Theory of Acceptance and Use of Technology (UTAUT; Venkatesh et al, 2003) explain adoption through cognitive perceptions, chiefly perceived usefulness and ease of use, yet treat capability beliefs thinly, as antecedents of intention. Social Cognitive Theory (Bandura, 1997) specifies how self-efficacy is built, through enactive mastery, vicarious experience, verbal persuasion and physiological states, but neither tradition has theorized, for the AI context, the social origin of capability belief inside organizations. Where leadership appears in adoption research it is usually framed as distal top-management support, while the proximal, daily mechanism, whether a worker sees their manager open an AI tool, reason with it and act on its output, has attracted little attention, despite Bandura's (1997) claim that similar, credible and proximate models are the most potent source of vicarious efficacy information.

We address this gap through an exploratory study of eleven knowledge workers across five domains, recruited in Morocco. Morocco is theoretically informative: as an emerging economy of relatively high power distance (Hofstede, 2001), leader behavior carries pronounced symbolic weight, which sharpens the visibility of the modeling mechanism. The paper repositions leadership from a distal behavioral lever to a proximal cognitive enabler of capability beliefs; it specifies the grounding of AI self-efficacy, in a retained competence or in the tool, as socially conditioned and as the property that governs the belief's durability; it delineates a boundary of TAM, perceived usefulness being distinct from and insufficient for self-efficacy; and it contributes evidence from an emerging economy largely absent from the literature. With eleven purposive interviews and retrospective self-reports, we make no causal or generalizable claims (Baron and Kenny, 1986); the output is a mechanism model expressed as testable propositions.

2. Theoretical Background and Propositions

Knowledge work centers on non-routine cognitive tasks, precisely the domain where generative AI is at once most disruptive and most ambiguous. Unlike deterministic software, an AI assistant requires the user to develop a tacit, probabilistic sense of where it can be trusted and how to steer it: a craft of prompting, verification and integration learned through use rather than instruction, in which errors can be confident and plausible (Dell'Acqua et al, 2023). Confidence in this craft, rather than familiarity with an interface, governs sustained adoption.

That reframing exposes a tension between two logics of adoption. TAM (Davis, 1989) and UTAUT (Venkatesh et al, 2003) predict intention well but say little about how the enactment that follows reshapes the user, or about capability belief as a distinct, cultivable construct. Social Cognitive Theory supplies the complement, locating the engine of behavior in self-efficacy and ranking its sources, with enactive mastery and vicarious experience the most potent (Bandura, 1997). A further tension separates perception from capability: perceived usefulness is a judgment about the tool's value, whereas self-efficacy is a judgment about the self's capability, increasingly treated as a distinct construct (Wang and Chuang, 2024). A worker may rate AI highly useful while feeling incapable of extracting that usefulness, and it is in that gap that adoption stalls. Beneath both lies the tension between the individual and the social: though the efficacy-to-performance relationship is robust across domains (Stajkovic and Luthans, 1998; Judge et al, 2007), the most potent non-enactive source of self-efficacy, vicarious experience, is irreducibly social, so the organizational distribution of observable models becomes a structural determinant of who develops capability belief. Taken together, these tensions yield our first proposition.

> ***Proposition 1:*** *Formal AI training experienced as compliance-driven, mandatory, generic and decoupled from observed use does not meaningfully increase workers' AI task self-efficacy, which is built primarily through self-directed enactive mastery.*

If training is not the lever, leaders move to the centre of the account. Much competence is acquired by observing models perceived as credible, similar and high in status (Bandura, 1997), and within organizations leaders are uniquely positioned as such models: visible, status-laden, and read as signaling what the organization values, as distinct from what its policies state. When a manager visibly uses AI, they transmit two things at once: procedural information about how the tool is used, and a symbolic signal that such use is legitimate and rewarded. This dual transmission is the heart of leadership role modeling. It is adjacent to, but distinct from, the leader-expectancy mechanism of the Pygmalion tradition (Eden, 1992): there, leaders raise performance by communicating expectations; here, they raise capability belief by enacting a behavior that followers can observe, imitate and feel sanctioned to attempt.

Three leadership streams are each individually insufficient here: role-modeling research rooted in social learning (Bandura, 1977) has concentrated on ethical conduct rather than on task-specific capability belief; sensegiving (Gioia and Chittipeddi, 1991) illuminates the legitimizing half of our mechanism but has rarely been connected to self-efficacy; and e-leadership scholarship (Avolio, Kahai and Dodge, 2000) has treated technology as the medium of leadership rather than its object. The leader is, in our account, at once a behavioral model, a sensegiver and a digital exemplar.

The mechanism is sharpened in high power-distance settings (Hofstede, 2001), where the symbolic content of what a leader does, and conspicuously does not do, carries amplified weight, so that the absence of modeling may read as tacit non-endorsement. This also reframes perceived usefulness: a favorable usefulness judgment motivates engagement, which produces mastery, which builds efficacy; but under uncertainty, a worker who believes AI useful may still not initiate sustained engagement absent a credible signal that doing so is

appropriate. The chain is, in other words, conditional on a legitimating signal, which leaders are best placed to supply.

> ***Proposition 2:*** *Visible AI use by leaders increases workers' AI task self-efficacy through two coupled pathways, symbolic legitimization and vicarious learning, which may be carried by different sources; managerial sponsorship without visible use does not, and the absence or discretion of leader use contracts engagement to a narrow perimeter of low-stakes tasks.*

> ***Proposition 3:*** *Perceived usefulness increases AI self-efficacy only indirectly, through engagement and mastery, and this indirect chain initiates only in the presence of a legitimating signal, of which leadership role modeling is the most efficient but not the only source.*

The efficacy-to-performance link is not mechanical, and this is where the grounding of self-efficacy enters. Enactive mastery is the most potent source of efficacy because success supplies information about one's own capability (Bandura, 1997); under assisted work, the question is to whom that information accrues. Two workers may produce identically strong AI-assisted outputs yet ground the confidence that follows differently: one in a competence she retains, the prompting, judgement and verification she exercised; the other in the tool's availability. Our review named these groundings durable and tool-contingent and argued that they should diverge when the tool is withdrawn (Sidqui and Zenjari, 2026), the divergence on which the ironies of automation turn (Bainbridge, 1983); that the two beliefs are separable is visible in instrument work, where the one item about performing without the tool loads on general rather than AI self-efficacy (Prather et al, 2026). The observable trace of grounding is the appraisal a worker gives to her assisted successes, what our coding manual calls mastery appraisal, and it is not fixed disposition: leaders who model AI as a craft may lead followers to ground their confidence in the craft, whereas leaders who present it as a convenience that delivers may lead them to ground it in the tool. The chain is finally bound by context: where the cost of AI error is high, as with clinical liability, threat perceptions dominate and even available models may be discounted.

> ***Proposition 4:*** *The framing leaders give to AI use, as a craft to be verified or as a convenience that delivers, shifts the grounding of workers' AI self-efficacy toward a competence they retain (durable) or toward the tool's availability (tool-contingent).*

> ***Proposition 5:*** *The grounding of AI self-efficacy conditions its durability and calibration: durably grounded self-efficacy survives the tool's withdrawal and stays bounded by verifiable tasks, whereas tool-contingent self-efficacy anticipates degradation without the tool and extends beyond what the worker can verify.*

> ***Proposition 6:*** *Sectoral context and organizational support moderate the strength of the modeling chain: high task-level threat (liability, hallucination risk) and weak institutional signaling attenuate it, while supportive, low-threat contexts strengthen it.*

3. Methodology

We adopt an exploratory, interpretive qualitative design, on the reasoning that where the aim is mechanism discovery rather than effect estimation, qualitative inquiry is the appropriate instrument (Gioia et al, 2013). We are not testing the propositions above; we derived them through the dialogue between data and theory and present them as the study's output, distinguishing throughout what participants reported from the mechanism we infer.

The evidence comes from eleven semi-directive interviews with knowledge workers purposively sampled across consulting (four), finance (two), technology (two), academia (one) and medicine (two), identified through the authors' professional networks and by referral, and screened so that all held genuinely knowledge-intensive roles; they ranged from daily, unbounded users to a physician in specialty training with no professional use. Selection maximized variation along two dimensions fixed before fieldwork, the visibility of senior AI use and the task-level stakes of AI error: consulting and technology anchored the high-visibility, lower-stakes condition, medicine the low-visibility, high-stakes condition, finance and academia intermediate positions (Table 1). One finance participant, a Moroccan professional recruited in Morocco like the others, works for the Canadian entity of an international group; the physician in specialty training (P-052) supplies contextual material only and is reported as an exploratory interview.

Table 1: The eleven cases

Case	Sector	Lang.	Questions asked (of 37)	Senior AI use observed	Formal training received	Mastery appraisal (grounding)
P-011	Consulting	FR	30	Informal sharing; senior use not explicit	None; informal support	Durable
P-012	Consulting	FR	34	Not reported; manager neutral	Core modules; self-training	Tool-contingent
P-014	Consulting	EN	37	Discreet; seniors use it less than juniors	None	Durable
P-015	Consulting	FR	37	Visible; described as normalizing use	Resources and peer exchanges	Durable, explicit arbitration
P-021	Finance*	EN	29	Visible; manager and executives	None	Tool-contingent
P-022	Finance	EN	28	Manager supportive; use unspoken in team	Mandatory course	Durable claim, low engagement
P-031	Technology	FR	26	Visible senior use; also models for others	None; peer sharing	Mixed, time-dependent
P-032	Technology	FR	18	Not explicit; continuous informal exchange	None	Durable (thin data)
P-041	Academia	FR	28	Supervisor uses rarely; peers visible	None; requested	Balanced
P-051	Medicine	EN	30	None	None	Durable
P-052	Medicine	FR	14	None in practice	None	No data (exploratory)

* Moroccan professional recruited in Morocco, employed by the Canadian entity of an international group. Mastery appraisal read from nodes 2.1 (competence retained) and 2.2 (tool availability).

An eleven-participant sample is a deliberate choice: the logic of qualitative sampling is depth within a bounded phenomenon, not statistical representativeness (Lincoln and Guba, 1985). Interviews followed a thirty-seven-question semi-structured guide covering adoption history, episodes of use and failure, training and support received, the manager's and organization's position, what participants observed of senior and peer use, and how they explained their own successes. The guide was not administered uniformly: six questions were put to all eleven participants, and coverage ranged from fourteen to thirty-seven questions (Table 1). Interviews were conducted in French (six) or English (five), recorded with consent and transcribed verbatim; excerpts from French interviews were translated by the research team, cross-checked for fidelity of meaning and are marked "tr." below.

Analysis followed the principles of Gioia et al (2013) and was conducted in NVivo 12 in two verified cycles. First-order coding applied a thirty-two-node manual to 749 candidate extracts, each carrying the guide question it answered. A consistency check between the analysis workbook and the NVivo project then compared every planned reference with its coded counterpart: 73 per cent matched after the first cycle, and the divergences resolved into six conditions of application left implicit in the manual, fixed as dated arbitration rules and applied uniformly in a second cycle that returned 446 references at 77 per cent conformity. Two rules matter here: observation that legitimizes use and exchange that transmits know-how were separated by mechanism rather than by rank, and a negative answer to a question that was asked, such as "no training", was coded as a documented absence rather than left as silence. Second-order analysis for this paper read the nodes bearing on the leadership question across cases by constant comparison, testing whether each candidate mechanism held where visible senior use was present and dissolved where it was absent. The five themes below draw on the nodes for training and enactive mastery (each carried by all eleven cases), visible senior use and organizational position (nine and eleven), vicarious experience and verbal persuasion (ten and eight), the grounding of confidence and its anticipated fate without the tool (nine), and fluent failure, delegation boundary and proxy agency (nine each). Case coverage matters more than counts: a node carried by nine cases supports a corpus-level claim, a node carried by one supports only an illustration.

Trustworthiness is addressed through the criteria of Lincoln and Guba (1985): credibility through maximum-variation sampling, the retention of a non-adopting case, a negative-case node (seven references, five cases)

and absences coded as results; transferability through the case attributes in Table 1, though we make no transfer claim; dependability and confirmability through the two dated coding cycles and the traceability of every reference from theme to node to source extract. We do not overstate these safeguards: the consistency check is a within-team audit, not independent inter-coder agreement; the manual was built for a wider programme on the durability of AI self-efficacy, so the leadership analysis is a secondary reading of the corpus; and proximity to the Moroccan context risked reading leadership into accounts, a risk we mitigated by disciplined first-order coding and by foregrounding the cases where the mechanism failed.

4. Findings

Formal training was scarce: eight of eleven participants reported none when asked ("Nothing", "No trainings", "None"). Provision, where it existed, took the form of core modules (P-012), learning resources and exchanges of good practice (P-015), and a single mandatory programme, the only one in the corpus delivered as formal instruction:

> *"Yes, there was a training, I attended because it was mandatory, it was provided by an external consultant to help you prompt and use the tools. It was just checking a box for me." (P-022, finance)*

What built confidence instead was doing. Assisted enactive mastery is the only source present in all eleven interviews, and participants described it as self-directed: "I have always been self-taught for subjects like that, I learn by doing" (P-031, technology, tr.); "I learned above all through personal experimentation, testing the tools, comparing the results and progressively refining my practice" (P-011, consulting, tr.). The gap was experiential, not informational, consistent with Proposition 1. Two qualifications follow. Four participants asked for more training, but of a different kind, "for anything beyond asking it and it helps you, over prompts and tricks" (P-022); and the one favourable account of support, P-015's, combined resources with visible senior use and peer exchange, foreshadowing that what mattered was the presence of models, not the format of provision.

Where managers and senior professionals used AI visibly, participants described a change in permission rather than in know-how:

> *"The use of AI by managers and senior profiles is visible. That contributes to normalizing its use within the team." (P-015, consulting, tr.)*

> *"She uses AI a lot for emails and summaries and schedules for her upcoming weeks. […] Knowing that she uses it too means that I am on the right path." (P-021, finance)*

P-015's manager "encourages experimentation while reminding us of the importance of verifying the results" (tr.), a framing that models AI as a craft; in technology, a senior colleague's use worked as motivation, "since he uses it too, it pushes me to improve the way I use it" (P-031, tr.). This is the dual transmission anticipated by Proposition 2, with the symbolic half doing most of the work: what participants took from senior use was that use was sanctioned. The contrast cases separate modeling from sponsorship. P-022 works under a manager he calls "very supportive of AI", in an organization that circulated an AI strategy and made training mandatory, yet this support "has not changed my way of using it"; in his team "everybody uses it but we don't talk about it much", and his own use is confined to proofreading and quick calculations. P-014 reports that "the seniors use it a bit less I would say, juniors it's very obvious. I'd say it's more discreet", and restricts AI to proofreading and paraphrasing. In medicine the absence is complete: "They don't use it, they barely understand it" (P-051). In none of these cases did absent or discreet senior use produce low declared confidence, P-014 is "very confident in my ability given enough time", but it coincided with a narrow perimeter of low-stakes tasks. In a high power-distance context (Hofstede, 2001), senior silence appeared to be read as tacit non-endorsement: use continued but stayed discreet, the pattern Beane (2019) calls shadow learning, and did not expand. What absent modeling withholds is not confidence but the licence to extend it.

A third pattern qualifies the leadership account: the procedural half of vicarious learning, how to prompt, iterate and verify, was transmitted mostly laterally. P-011 learned from observing others "uses I would not have thought of alone, in particular how to formulate requests, iterate, or use AI on very concrete tasks" (tr.); P-041 first tried the tool after a colleague "fascinated by GPT presented it to us" (tr.); P-031 was himself the model, "as they saw me work, they asked how I did it and they started" (tr.). Verbal persuasion was mostly absent, four participants reporting that their capability with AI had never been remarked on. Legitimacy descends from seniors; know-how circulates among peers.

Usefulness on its own did not produce capability belief. The physician who judged the technology capable, "AI is able to diagnose rare disease", added in the same breath, "I don't use it" (P-051); P-022 rated the tool's

answers "7 out of 10" yet said there had "never been a success story with AI". A favourable judgment of the tool coexisted with flat engagement, the seam between TAM and Social Cognitive Theory that Proposition 3 anticipates: usefulness motivates but does not constitute capability. Where the chain did run it ran on trial, "the more I use AI successfully, the more comfortable I am exploring new uses" (P-015, tr.). One case bounds the leadership version of the claim: P-041, with no institutional support and a supervisor who "uses it herself but not often" (tr.), shows the densest engagement in the corpus, initiated on peer models. The legitimating signal appears necessary; leaders are its most efficient but not its only source.

Among participants with comparable daily use, the appraisal given to success, and the grounding of the confidence that followed, diverged. Some grounded it in a competence they retained: "I am definitely the expert, the prompts are mine and the instructions are mine, it is there to assist only" (P-022, finance); "with equal access to the tool, the difference is not about access to AI, but about the capacity to use it well" (P-011, tr.). Others grounded it in the tool:

> *"If I am being honest, it felt like I was wasting my time doing it without AI; with AI it felt like having a magic wand that allowed to do things instantly." (P-021, finance)*

> *"AI saved me that day: I was able to meet the delivery deadline while producing quality work, which I could not have done without Copilot." (P-012, consulting, tr.)*

The first is the durable grounding, the second the tool-contingent one (Sidqui and Zenjari, 2026); in our data structure, mastery appraisal is the aggregate dimension that carries both and is treated as a candidate indicator of durability. Where the difference comes from is readable, tentatively and on three cases, in line with Proposition 4. P-015's manager frames AI as experimentation plus verification, and she alone in the corpus arbitrates explicitly between the two groundings ("the success comes both from the tool and from my expertise", tr.) and states her self-efficacy by task rather than globally, "8 out of 10" for writing and synthesis, "around 6 out of 10" for complex analyses (tr.), the generality dimension Bandura (1997) distinguishes from level. P-021's executives and manager present Copilot as a productivity tool to adopt, and his vocabulary is that of magic wands and "a superpower". P-012's manager is "neutral" and her clients advise using AI "on low value-added tasks" (tr.); her confidence rests predominantly on the tool.

What grounding governed was not the size of the performance gain, which nine of the ten participants who answered reported whichever grounding they held, but the durability of the confidence built on it. Asked what would happen without the tool, tool-contingent participants anticipated collapse, "I would probably panic. I no longer know how to function without it" (P-012, tr.), "I would go back to doing the tasks as I used to, which takes more time and less confidence" (P-021); durably grounded participants claimed continuity, "I'd do it on my own" (P-014), "I could still carry out my tasks thanks to my knowledge, but it would take me more time" (P-015, tr.). Tool-contingent confidence also reached further than the worker could verify: P-012 could "calmly answer my client's various requests, including on subjects I do not master, because I know AI will be there" (tr.), the one clear instance of uncalibrated trust in the corpus (Lee and See, 2004). This is Proposition 5: grounding conditions the durability and calibration of self-efficacy rather than its conversion into performance. Durable grounding did not, however, protect capability itself: P-014 claims an intact expertise while noting that "the more I rely on it the less my own personal touch is felt in my work".

Where legitimacy was absent and the cost of error high, the chain did not run at all, and medicine displayed this failure mode most starkly, illustrating Proposition 6:

> *"Using AI in this field is very reckless, since it tends to hallucinate things, the treatment can be wrong, so who will bear the responsibility of this mistake? […] Who are you going to sue if a mistake happens?" (P-051, medicine)*

P-051 has "never used it in something that matters" and remains "still not confident, I am skeptical about using AI"; the physician in specialty training uses no AI professionally and reports declining trust "because we do not know where the data is stored" (P-052, tr.). Academia, contrary to our expectation, did not echo this pattern: P-041 met the same fluent failures, "ChatGPT gave me fictitious sources, articles that did not exist" (tr.), but in a domain where the error is checkable by the worker and borne by her, responded with a delegation boundary, "I delegate no control to AI" (tr.), and kept engaging. The chain is not self-starting: it requires a legitimating signal and a threat environment in which the worker can detect and own the error. Remove both, as medicine removes both, and the mechanism does not run, however capable the technology is judged to be.

5. Discussion

Prevailing treatments emphasize leadership's behavioral and structural influence: leaders sponsor initiatives, allocate resources and communicate vision, thereby raising usage. Our data separate that sponsorship from something more proximal. A supportive manager, a circulated strategy and a mandatory course left one participant's practice unchanged; a manager seen using the tool for her own emails told another that he was "on the right path". What changed, when workers observed senior use, was internal: a revision of what they believed they were permitted to do. Modeling functioned as a legitimacy signal that lowered the perceived risk of extending use, unlocking the engagement that mastery requires, which reframes leadership from behavioral influence toward cognitive enabler.

The same evidence extends self-efficacy theory while bounding TAM, in two ways the original design did not anticipate. First, the two components of vicarious experience were carried by different people: legitimacy descended from seniors, while procedural know-how circulated laterally among colleagues, and in one case flowed upward from a junior who became the team's model. Bandura (1997) treats observation as a transfer of capability information, "someone like me can do this"; what senior use transmitted here was closer to "engaging with this tool is permissible for someone like me", a signal about legitimacy rather than capability. Whether that is a variant of vicarious experience, a form of persuasion enacted rather than spoken, or a source the four-source taxonomy does not contain is a question we carry forward rather than settle; in hierarchical settings, at least, the transfer of permission and the transfer of technique are separable. Second, observation appeared to shape the mastery appraisal downstream: where a manager framed AI as experimentation to be verified, the follower grounded her confidence in her own developing skill and stated it by task; where executives framed it as a productivity tool to adopt, the follower grounded it in the tool. Grounding did not govern the size of the reported performance gain, which was nearly universal; it governed whether self-efficacy was durable or contingent, anticipating collapse without the tool (Bainbridge, 1983; Sidqui and Zenjari, 2026). This matters for measurement: a survey that finds AI integration raising self-efficacy (Shishakly et al, 2026) cannot distinguish the two groundings when its items ask about working with the tool, and the one instrument that carried a tool-free item found it loading elsewhere (Prather et al, 2026). The same data bound TAM: perceived usefulness was associated with favourable judgments of the tool but neither with engagement nor with capability belief. TAM and UTAUT are not wrong but incomplete for AI, modeling the antecedents of intention while under-specifying the social and experiential process by which capability belief is built.

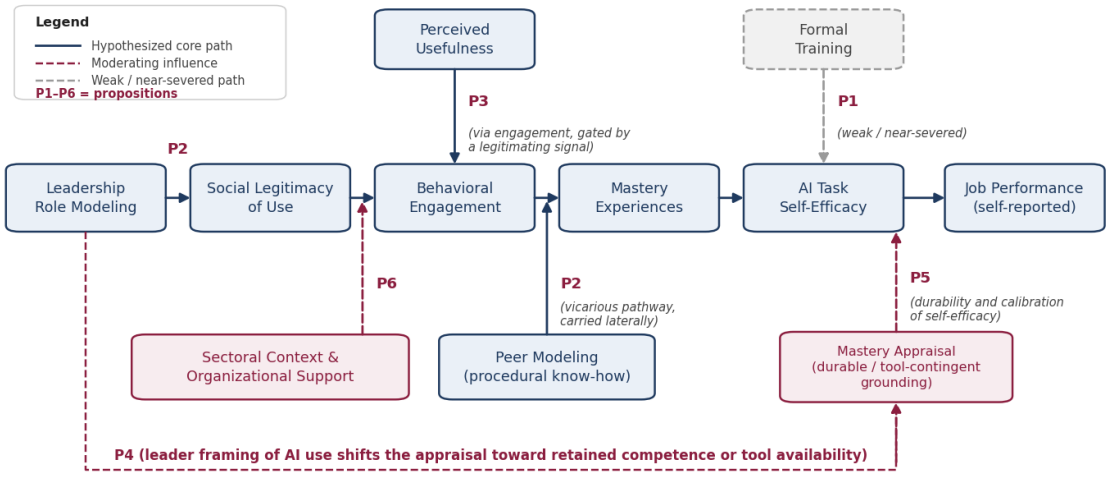


Figure 1: A mechanism model of leadership role modeling, social legitimacy of use and the grounding of AI self-efficacy in knowledge-intensive roles, as revised after analysis

Synthesizing the propositions yields a revised mechanism model (Figure 1). The core path runs from leadership role modeling through social legitimacy of use, behavioral engagement, mastery experiences and AI task self-efficacy to job performance; peer modeling feeds engagement and mastery with procedural know-how, in parallel with the legitimacy leaders supply. Sectoral context and organizational support moderate the upstream link from legitimacy to engagement, and the grounding of self-efficacy, durable or tool-contingent and itself

shifted by how leaders frame AI use, conditions the durability and calibration of the belief that mastery produces. Formal training appears as a near-severed path (Proposition 1).

Two paradoxes follow. The first is AI as empowerment versus threat: the same technology that a finance analyst experienced as "a superpower" and a product lead used to "replace a Product Manager, an architect" (P-031, tr.), a physician experienced as a liability to be managed by avoidance. The valence appeared to be a property not of the tool but of the configuration around it, whether legitimacy was supplied and whether the error could be detected and owned by the worker. The second is leaders as enablers versus gatekeepers: a leader who models AI enables a self-reinforcing cycle of engagement, mastery and confidence, while a silent leader, especially in a high power-distance setting, withholds the signal that would unlock it, often unintentionally. Leadership silence is therefore not the absence of an intervention but an intervention in itself.

Practically, these findings invert the usual investment logic. Formal provision was rare and, where mandatory, confidence-neutral; the participants who asked for training asked for something task-anchored and advanced, not another prompting session. The most effective intervention in our data was a respected senior using AI in plain sight and, in the one case where self-efficacy was both high and well calibrated, narrating the verification. Managers might treat their own observable behavior as the adoption strategy, and the framing they attach to it as a lever on whether followers build a competence or a dependency, the more so in Morocco, where high power distance (Hofstede, 2001) may amplify both the value of leader engagement and the cost of leader silence. For researchers, the task-referenced self-efficacy participants volunteered, distinct levels for writing and for analysis, argues for instruments that ask about tasks rather than about the tool, and that include what the worker can do without it.

Five limitations bound these claims. Eleven purposive interviews support mechanism discovery, not prevalence: the mastery-appraisal reading rests on three cases and the sectoral contrasts on one or two participants per sector. The guide was not administered uniformly, and the withdrawal question on which the durability reading depends was put directly to six participants, with answers for three others recovered from elsewhere in their interviews. Self-reports are vulnerable to recall and rationalization, performance is self-reported throughout, and cross-sectional data cannot establish the temporal ordering the model implies (Baron and Kenny, 1986). The findings are situated largely in one country, whose power-distance profile both makes it informative and limits transfer. And a single team coded and interpreted the data. The propositions accordingly invite a quantitative phase: a multi-sector survey operationalizing leadership role modeling, peer modeling, social legitimacy of use and AI self-efficacy with its grounding, and a design in which performance is observed with the tool present and withdrawn, so that durable and tool-contingent self-efficacy can be told apart by behavior rather than by report. Intervention studies contrasting leader demonstration with training would sharpen these conjectures further.

6. Conclusion

Organizations have largely treated AI adoption as a problem of access and instruction. Our exploratory evidence points elsewhere: the binding constraint appears to be capability belief, and the most potent organizational lever on it may be not training, nor sponsorship, but the visible behavior of leaders. When leaders use AI first, and use it as an observable, learnable craft, they appear to supply the legitimacy that unlocks engagement, leave to peers the know-how that builds mastery, and set the framing that lets workers ground their confidence in a skill they own rather than in a tool they cannot do without. When leaders are silent, especially when hierarchy is salient, use goes discreet and the chain does not extend. We offer these mechanisms as propositions, not conclusions; eleven interviews cannot settle the question, but they can sharpen it, relocating the study of AI adoption from the cognition of the tool to the cognition of the self, and naming leadership, modeled rather than mandated, as the social act that bridges the two.

Ethics Declaration: Participation in this study was entirely voluntary and based on informed consent. Participants were fully briefed on the purpose of the research, the nature of their involvement, and their right to withdraw at any time without consequence. In accordance with Moroccan Law 09-08 on the protection of personal data (in alignment with GDPR rules), no sensitive data were collected, and all information was anonymized prior to analysis. Audio recordings were destroyed following transcription and anonymization, and anonymized data are securely stored for research purposes only. At the time of the study, formal ethical clearance was not required under the institutional policies of Groupe ISCAE for non-invasive research involving voluntary adult participants and anonymized data.

AI Declaration: A generative AI assistant (a large language model) was used to support the structuring, drafting and language editing of this manuscript. The study design, data collection, coding, thematic analysis,

interpretation and conclusions are the work of the authors. All AI-assisted text was reviewed, verified and edited by the authors, who take full responsibility for the content of the paper, including the accuracy of all reported data and quotations.

References

Avolio, B.J., Kahai, S. and Dodge, G.E. (2000) "E-Leadership: Implications for Theory, Research, and Practice", *The Leadership Quarterly*, Vol 11, No. 4, pp 615–668.

Bainbridge, L. (1983) "Ironies of Automation", *Automatica*, Vol 19, No. 6, pp 775–779.

Bandura, A. (1977) *Social Learning Theory*, Prentice-Hall, Englewood Cliffs.

Bandura, A. (1997) *Self-Efficacy: The Exercise of Control*, W.H. Freeman, New York.

Baron, R.M. and Kenny, D.A. (1986) "The Moderator-Mediator Variable Distinction in Social Psychological Research: Conceptual, Strategic, and Statistical Considerations", *Journal of Personality and Social Psychology*, Vol 51, No. 6, pp 1173–1182.

Beane, M. (2019) "Shadow Learning: Building Robotic Surgical Skill When Approved Means Fail", *Administrative Science Quarterly*, Vol 64, No. 1, pp 87–123.

Brynjolfsson, E., Li, D. and Raymond, L. (2025) "Generative AI at Work", *The Quarterly Journal of Economics*, Vol 140, No. 2, pp 889–942.

Compeau, D.R. and Higgins, C.A. (1995) "Computer Self-Efficacy: Development of a Measure and Initial Test", *MIS Quarterly*, Vol 19, No. 2, pp 189–211.

Davis, F.D. (1989) "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology", *MIS Quarterly*, Vol 13, No. 3, pp 319–340.

Dell'Acqua, F., McFowland, E., Mollick, E.R., Lifshitz-Assaf, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F. and Lakhani, K.R. (2023) *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality*, Working Paper 24-013, Harvard Business School, Boston.

Eden, D. (1992) "Leadership and Expectations: Pygmalion Effects and Other Self-Fulfilling Prophecies in Organizations", *The Leadership Quarterly*, Vol 3, No. 4, pp 271–305.

Gioia, D.A. and Chittipeddi, K. (1991) "Sensemaking and Sensegiving in Strategic Change Initiation", *Strategic Management Journal*, Vol 12, No. 6, pp 433–448.

Gioia, D.A., Corley, K.G. and Hamilton, A.L. (2013) "Seeking Qualitative Rigor in Inductive Research: Notes on the Gioia Methodology", *Organizational Research Methods*, Vol 16, No. 1, pp 15–31.

Hofstede, G. (2001) *Culture's Consequences: Comparing Values, Behaviors, Institutions, and Organizations Across Nations*, 2nd ed, Sage, Thousand Oaks.

Judge, T.A., Jackson, C.L., Shaw, J.C., Scott, B.A. and Rich, B.L. (2007) "Self-Efficacy and Work-Related Performance: The Integral Role of Individual Differences", *Journal of Applied Psychology*, Vol 92, No. 1, pp 107–127.

Lee, J.D. and See, K.A. (2004) "Trust in Automation: Designing for Appropriate Reliance", *Human Factors*, Vol 46, No. 1, pp 50–80.

Lincoln, Y.S. and Guba, E.G. (1985) *Naturalistic Inquiry*, Sage, Beverly Hills.

Marakas, G.M., Yi, M.Y. and Johnson, R.D. (1998) "The Multilevel and Multifaceted Character of Computer Self-Efficacy: Toward Clarification of the Construct and an Integrative Framework for Research", *Information Systems Research*, Vol 9, No. 2, pp 126–163.

Noy, S. and Zhang, W. (2023) "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence", *Science*, Vol 381, No. 6654, pp 187–192.

Prather, J., Margulieux, L., Kharitonova, Y., Yang, Y., Reeves, B.N., Denny, P., Gorson Benario, J., Holmes, E.D.V., Spaulding, E., Barbre, G., Blake, M. and Leinonen, J. (2026) "A Validated Scale Measuring Student Self-Efficacy for Programming with Generative AI", *Proceedings of the ACM Conference on* International Computing Education Research (ICER 2026), Vol 1, ACM, New York, pp 59–72.

Sidqui, O. and Zenjari, A. (2026) "Training Wheels or Real Balance? AI Self-Efficacy and Performance in Knowledge-Intensive Work: A Systematic Review", *Human Resources Management and Services*, Vol 8, No. 4, 6167.

Shishakly, R., Khattak, M.N. and Abukhait, R. (2026) "Artificial Intelligence in the Workplace: Examining the Role of Perceived Service Quality and Employee Crafting in Shaping Resilience, Self-Efficacy and Innovation", Current Psychology, Vol 45, Article 1349.

Stajkovic, A.D. and Luthans, F. (1998) "Self-Efficacy and Work-Related Performance: A Meta-Analysis", Psychological Bulletin, Vol 124, No. 2, pp 240–261.

Venkatesh, V., Morris, M.G., Davis, G.B. and Davis, F.D. (2003) "User Acceptance of Information Technology: Toward a Unified View", *MIS Quarterly*, Vol 27, No. 3, pp 425–478.

Wang, Y. and Chuang, Y. (2024) "Artificial Intelligence Self-Efficacy: Scale Development and Validation", *Education and Information Technologies*, Vol 29, pp 4785–4808.