UNIVERSITY OF SOUTHERN DENMARK

Cand.merc. Strategy & Organization · Master's Thesis

Project Hephaistos – Research Summary

A qualitative study of AI-integrated software development teams

18 June 2026

AUTHOR
Rasmus Heydt Wraae

SUPERVISOR
Nicole Richter - Faculty of Business and Social Sciences.

Executive summary

Artificial intelligence (AI) is no longer something only a few teams have access to. As the tools become available to everyone, the advantage moves away from having the technology and toward how well a team actually integrates it into its daily work. This puts the human side of the team back at the center of performance. If almost everyone can use the same tools, then the difference between teams comes down to the people using them.

This report summarizes a qualitative study that asked a single question: which human characteristics, at the team level, do the academic literature and industry practitioners identify as drivers of performance and wellbeing in AI-integrated teams, and where do the two perspectives agree or disagree. To answer it, the study draws on eight interviews, seven with practitioners who already work AI-natively in software, data, and AI consulting, and one with an academic expert. Each interview was analyzed theme by theme and compared, point by point, against what the team-effectiveness literature already claims, using a shared framework so that the two sides could be held against each other directly.

The comparison produces a consistent pattern. The two perspectives converge on who makes an AI-integrated team work, and diverge on what happens inside it. They agree on the human characteristics to look for and on the conditions to provide, but AI reorders which ones matter most, pushing proactive agency, a product and value-creation mindset, AI literacy, and calibrated judgment to the front. Technical execution skill, the ability to write the work by hand, matters less than before, while experience remains an asset, since it is what lets a person judge whether the AI's output is any good. Where the two sides part is the team-level processes that the literature has studied for years and the practitioners barely named.

For the participating companies, the message is direct. As execution moves into the tooling, hiring shifts from raw technical skill toward disposition, toward how a person thinks, takes ownership, and works. It is important to mention here that the performance gains the practitioners describe are perceived rather than measured, so they are best read as a credible signal rather than proof. Overall, AI did not make people matter less in these teams. It reordered what matters, and raised the stakes on getting it right.

Background & Purpose

The advantage that comes from a new technology rarely lasts. In the late nineteenth century, firms could replace their steam engines with electric motors, yet many saw no immediate gain in productivity. What eventually drove performance was not electricity itself, but the way organizations redesigned their work around it, the workflow, the coordination, and the division of labor (Brynjolfsson & McAfee, 2014). The pattern recurs with general-purpose technologies. Early advantage comes from access to the technology, but as it matures, the advantage shifts from the technology toward how well an organization integrates it. The same shift is now unfolding with AI. As the tools become available to everyone, the difference between teams comes down less to whether they use AI and more to who is using it.

Team composition has been studied for decades and has a strong foundation in the literature. But that foundation was built for teams without AI, and AI changes both the environment and the way the work is done. This raises a set of open questions. Does the significance of the well-established factors change once AI enters the team? Are factors that used to matter less becoming more important? And what should an organization looking to implement AI look for in the people it hires? To address these questions, the study is built around a single research question:

> **RQ Which human characteristics at the team level are identified, by the academic literature and by industry practitioners, as drivers of team performance and wellbeing in AI-integrated teams, and where do these two perspectives converge or diverge?**

The question has two sides, an academic one and a practical one, and setting the two next to each other is what lets the study ask where they agree and where they diverge.

It is important to be clear about what the study does and does not cover. Empirically, it focuses on AI-integrated tech and knowledge teams, and specifically on early adopters, meaning people who already use AI workflows, agents, or automation for the majority of their daily work rather than experimenting with it. Performance is understood broadly, covering quality, speed, efficiency, and innovation, with wellbeing treated as a co-equal outcome rather than a secondary one. What the study does not do matters just as much. It does not measure the degree of AI integration, and it does not measure team performance

directly, since eight qualitative interviews cannot do that. It identifies what practitioners and the literature say drives performance, not the size of any effect. It also does not address the technical, ethical, or societal dimensions of AI.

Study Design & Method

The study is qualitative and exploratory, which fits a question about AI-integrated teams where the established literature is still thin. The overall design follows the research onion of Saunders et al. (2023). The philosophical position, established with the reflexive HARP tool, is critical realism, and the approach is abductive, moving back and forth between the interview data and the existing theory rather than testing fixed hypotheses. Since the aim is to understand what practitioners experience and then to compare it with the literature, this positioning fits the work better than a purely deductive design would.

The evidence base is eight semi-structured interviews, seven with industry practitioners and one with an academic expert. The seven practitioners all qualify as early adopters, meaning they use AI workflows, agents, or automation for the majority of their daily work. Six were recruited through the network of the German company TalentFormation, and the remaining two, one practitioner and the expert, were arranged by the author. The expert is an associate professor at the Technical University of Denmark who researches AI in organizations and software development. The interviews were conducted in German, Danish, and English, mostly online, and followed a guide divided into seven content blocks that keep the conversation on track while leaving room to follow what each participant raises. It is important to mention here that the sample is not meant to be statistically representative. The design rests on Yin's (2018) replication logic rather than on sampling logic, so the right question is not whether eight interviews are many, but whether they are enough for a multiple-case study of this kind. By going for depth over breadth, they are.

ID	Role/Context	Sector	Sex	Original Language	Length	Format	Conducted By
P1	Senior consultant + senior practitioner at IT consultancy	Software / IT consulting	M	German	79 min	Paired one-on-two	TalentFormation
P2	Independent advisor on AI / agentic strategy to VC, PE and startups	AI strategy advisory	M	German	38 min	One-on-one	TalentFormation
P3	Independent software / AI consultant	Software / AI consulting	M	German	59 min	One-on-one	TalentFormation
P4	CTO at small software company	Software development	M	German	59 min	One-on-one	TalentFormation
P5	Software engineer / consultant at a software firm	Software development	M	German	51 min	One-on-one	TalentFormation
P6	Co-founder and developer at small dev studio	Software development (Next.js)	M	German	47 min	One-on-one	TalentFormation
P7	AI / workflow lead at a digital media network	Digital media / AI workflows	M	Danish	50 min	One-on-one	Author
E1	Associate Professor, Strategy and Leadership Development, researching AI in organizations and software development	Academic at DTU	F	English	41 min	One-on-one	Author

Table 1: Information on interview participants

The analysis is organized around the Input-Mediator-Output-Input (IMOI) framework (Ilgen et al., 2005), shown in Figure 1. The framework treats a team as a sequence, where inputs such as the characteristics of the people feed into mediators such as coordination and trust, which in turn produce outputs such as performance and wellbeing, with a feedback loop back to the inputs over time. The IMOI framework is used twice in the study, once to structure the academic literature and once to structure the interview findings. This is what makes the later comparison possible, since both sides are mapped onto the same structure and can be held against each other cell by cell.

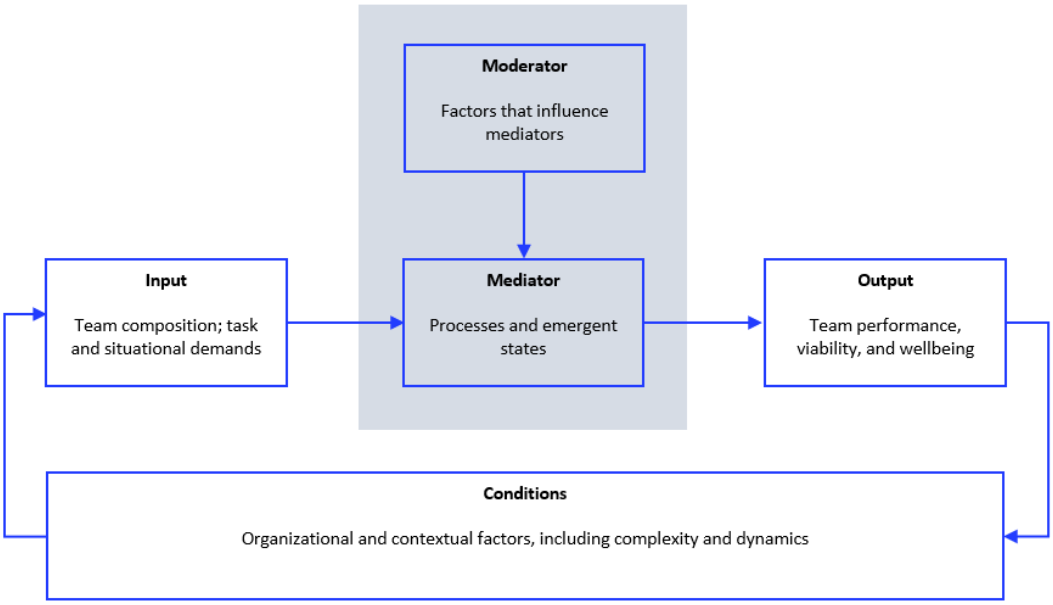


Figure 1: The IMOI framework, adapted from Ilgen et al. (2005) and Kozlowski and Ilgen (2006). The grey shading groups the Mediator stage, which consists of team processes and emergent states together with the factors that shape them.

The interviews were analyzed with Template Analysis (King & Brooks, 2017), using a hybrid coding strategy that combines a priori themes from the literature review with themes that emerged from the data, managed in NVivo. Three further steps make the evidence as credible as a qualitative design allows. First, the German interviews were coded in parallel by the supervisor as a second coder, which gives a cross-language check on the coding. Second, the seven practitioner cases were analyzed as the primary evidence, serving as literal replications of one another. Third, the expert interview was added as a theoretical replication, following Yin's two-tail logic, to test the practitioner findings against an independent academic view. The result of mapping the literature and the findings onto the same framework is the cross-walk presented in the later comparison.

Findings

The findings are organized on the IMOI framework, moving from the people on the team (inputs), to how they work together (mediators), to what the team produces (outputs), and finally the conditions and moderators that surround them. Most of the weight sits on the inputs, since this is where the practitioners had the most to say and where the study makes its main contribution.

Across the seven practitioners, a consistent set of human characteristics came up as the ones that separate strong performers in AI-native work. The classic picture of a good team member did not disappear. What changed is the order of importance, since AI pushes a particular set of characteristics to the front. Figure 2 maps the full input side as it emerged from the coding, and the seven themes are described below.

The single strongest finding, raised by all seven, is agency. The practitioners described it as taking ownership and acting on a problem instead of waiting to be told what to do, together with the willingness to make decisions under uncertainty and to push an idea through. As the tools take over more of the execution, the person who decides what is worth doing, and then takes responsibility for it, becomes the one who makes the difference.

Almost as prominent is what the practitioners called a product and value-creation mindset. They described a shift away from simply executing a task toward caring about the value behind it and the user it is meant for. When the AI can produce the work, the harder and more valuable skill becomes knowing which work is worth producing. This was one of the clearest patterns in the data, and as the comparison will show, it is also one the literature has barely touched.

All seven also pointed to AI literacy, the ability to actually get good results out of the tools, paired with enough technical grounding to judge whether those results are any good. This is where the practitioners drew an important distinction. The value of writing the work by hand has fallen, since the AI now does much of it. But experience has not lost its value. It is exactly what lets a person look at the AI's output and tell whether it can be trusted. What loses value is hand-execution, not knowing things.

Closely related is a cognitive style built on systems thinking and critical evaluation. The practitioners valued people who can hold the whole system in view, and who do not take the AI's output at face value but check it. As the work shifts away from producing the answer, judging the answer becomes the harder part of the job. Since the tools and ways of working also change constantly, adaptability and a willingness to keep learning came through as characteristics in their own right, since a person who was effective with last year's tools is not automatically effective with this year's. Clear communication and social sensitivity rounded out the picture, since even as teams get smaller, the ability to express an idea precisely and work well with others still matters.



Figure 2: The Input coding structure (I1–I7): threads, sub-themes, and themes. Own creation in NVivo.

Motivation and drive appeared too, though less universally than the others. What makes it interesting is the twist the practitioners gave it. AI raises output for everyone, but a motivated person turns that into better outcomes, while an unmotivated one simply produces more of the wrong thing faster.

Where the practitioners had a great deal to say about the people, they had much less to say about the team-level processes that sit between the people and the results. Coordination came up, mostly as an overhead to manage as the work speeds up, and so did the need for a team to move at a shared pace. But the deeper team dynamics that researchers spend a lot of time on, such as whether the AI is treated as a tool or as a teammate, or the sense of shared meaning in the work, were barely mentioned. Most practitioners spoke about the AI in tool and orchestration terms rather than as a teammate, and the question of meaning hardly surfaced at all. This near-silence is itself a finding, and the comparison returns to what it might indicate.

On results, the practitioners are united on one point and quiet on the rest. They are certain that AI makes work faster, and speed was by far the most common claim. It is important to mention here that these gains are perceived rather than measured, since the study does not measure performance, so they are best read as a strong and consistent impression rather than proof. Quality and efficiency were touched on but not developed in the same way. Alongside performance, wellbeing came through as a real concern, mostly around the cognitive load of working at AI speed and the risk of a new kind of burnout, which suggests that faster is not automatically better for the people doing the work.

The practitioners were also clear that these characteristics do not work in a vacuum. Two conditions shape whether they translate into performance. The first is the organization itself, its structure, its handovers, and how far along it is in its own AI transformation. The second is the enabling infrastructure, meaning access to the right tools and a workable approach to governance and data protection. On top of these, three factors moderate how strongly these characteristics translate into performance. Team size matters, and the practitioners leaned clearly toward smaller, more generalist teams, since AI lets fewer people cover more ground while every extra person adds coordination cost. Task type and risk matter too, since the practitioners themselves drew a line between low-risk software work and high-stakes domains such as medicine, where the same hands-off use of AI would not be acceptable. And the maturity of AI adoption matters, since a team only meets these dynamics once it is genuinely working AI-natively.

Taken together, the practitioner findings describe who makes an AI-integrated team work: a reordered set of human characteristics, a handful of conditions that let them translate, and an honest gap where the team-level processes would be.

The comparison: Literature vs Practice

This section sets the practitioner findings and the academic literature side by side, cell by cell, on the same IMOI framework. Read as a whole, one pattern stands out. The two perspectives converge on the people and on the conditions around them, they diverge on the team-level processes in between, and on the outputs they land somewhere between the two, since the expert and the literature qualify the practitioners' speed claim rather than deny it. Table 2 presents the full cross-walk, with the relation between the two sides marked for each construct, and the walkthrough that follows takes it one cell at a time. In the table, each row is labeled by where the construct sits in the framework, with **I** for inputs, **M** for mediators, **O** for outputs, **C** for the surrounding conditions, and **X** for the moderators, the levers that shape the team's mediating processes.

Construct	Literature	Practitioners	Relation	Cross-walk note
I1 Agency & ownership	Kovbasiuk (proxy); Frese & Fay (Thin)	7/7; strongest input finding	Converge (under-anchored)	Strong in practice, weak in literature; AI elevates it.
I2 Product & value-creation mindset	—	6–7/7	New	No literature analog; candidate original contribution.
I3 Calibrated cognition & critical evaluation	Bansal; Lee & See; Zhang & Chen	7/7	Converge	E1 confirms and adds failure modes (time pressure, illusion of mastery).
I4 Adaptability & learning	Bell; LePine; Edmondson (2001)	7/7	Converge	Classical openness and learning-orientation result holds.
I5 Technical foundation & AI literacy	Brynjolfsson; Dell'Acqua	7/7	Converge (quantified)	Experience tension; literature and E1 tilt toward 'asset'.
I6 Interpersonal & social sensitivity	Woolley; Bell	7/7	Converge	E1 confirms; practitioners add empathy toward the AI.
I7 Motivation, drive & resilience	Gagné & Deci; Luthans	4–5/7	Converge (partial)	Output-vs-outcome distinction is novel; anchors not AI-specific.

M1 Coordination & communication	Marks; Okhuysen & Bechky; Schmutz	Rich (coordination overhead)	Converge	Practitioners confirm HAT coordination cost; E1 adds asymmetric workload.
M2 Shared cognition / synchronized working	Cooke; Lewis; Zhang & Chen	M2a/M2b populated; M2c empty	Converge (partial)	Pace-sync matches pseudo-SMM; transactive memory never surfaced.
M3 Trust & psychological safety	Edmondson; Lee & See; Georganta	M3a/M3b populated; M3c empty	Converge + New	M3a/M3b converge; M3c (AI-mediated interpersonal trust) raised only by E1 - novel, measurement gap.
M4 AI role: tool vs teammate	O'Neill; Lyons	Thin (only P6)	Diverge	Literature: teammate framing matters; practitioners use tool / orchestration framing.
M5 Meaning & contribution	Project Aristotle	Empty	Practitioner gap	Literature predicts it matters; practitioner silence is the finding.
O1 Team performance	Brynjolfsson; Li; jagged frontier	O1a speed only; O1b/O1c empty	Qualify	E1: gains perceived, not measured; effect smaller at team level.
O2 Viability & wellbeing	Högemann; Liu & Li	O2a ceiling (3 cases); O2b strain	Converge	Core converges; E1 raises junior-cohort and relational concerns.
C1 Leadership & structure	Raisch & Krakowski; Kolbjørnsrud; Edmondson (2001)	Autonomy paradox; transformation radicality	Converge	E1 extends (policy gap) and raises AI-as-manager role inversion.
C2 Enabling infrastructure & governance	Kirchner; Papagiannidis	4/7; cleanest condition	Converge	Cleanest convergence in the comparison (GDPR / tool access).
X1 Team size & role structure	Li; Moe; Dybå & Dingsøyr	X1a 5/7; X1b 4/7	Converge	Strongest moderator; RQ-central; ties to the weakest-link logic.
X2 Task complexity & criticality	LePine; Dell'Acqua	Greenfield/legacy; criticality	Converge	E1 confirms (medicine vs IT consulting).
X3 AI integration maturity	Moore; Von Hippel	Spectrum: experimenting → AI-native	Converge	E1 confirms; bounds how far the findings generalize.

Table 2: The cross-walk: The literature review against the practitioner findings, by IMOI construct.

Starting with the inputs, the comparison delivers the substantive core of the answer. The literature and the practitioners converge on a recognizable set of human characteristics, and the expert supports the same cluster, but AI reorders the profile. Agency, product orientation, AI literacy, and calibrated cognition move to the front, while the classical personality core stays in place beneath them. Most of these converge quietly. Adaptability sits on the established openness and learning-orientation findings, interpersonal sensitivity on the social-sensitivity literature, and motivation on self-determination theory, though the practitioners add a distinction the general theories do not, that motivation now drives the outcome rather than the output, since AI lowers the cost of producing work but not the cost of caring whether it is any good.

Two characteristics stand out, and they are the most important result of the comparison. Agency (I1) is the strongest finding in the interviews, raised by all seven practitioners, yet the literature barely anchors it, with only thin and indirect support. The product and value-creation mindset (I2) goes further, since the literature review found no source at all that matches it. It is not a refinement of an existing construct, but a characteristic the practitioners surfaced that the literature has not yet named. I would argue that these two are the clearest candidates for an original contribution, since practice has run ahead of theory on both. For the company reader, the practical reading is that the personality core and AI literacy rest on solid research, while agency and the product mindset rest for now on a strong practitioner consensus rather than on tested evidence.

The mediators are where the two sides part. Here the practitioners had little to say, while the literature, speaking through the expert, has a great deal. Coordination and a shared working pace did surface, and they converge with the literature on coordination cost. But the deeper team-level states diverge. The tool-versus-teammate question, which the human-autonomy literature treats as central, was raised by only one practitioner, since most spoke about the AI as a tool to orchestrate rather than a teammate. Shared meaning did not surface at all, even though the literature predicts it matters. The most defensible reading is not that these processes have stopped mattering, but that the literature is ahead of practitioner articulation on this layer. The practitioners are early adopters absorbed in what to bring to the work and how to direct the tool, and they have not yet developed the vocabulary for the team-level states the literature has spent decades naming. This is the single most productive divergence in the comparison, since it marks where AI-integrated teamwork is least understood.

The outputs are where the two perspectives sit furthest apart, though the gap is one of qualification rather than contradiction. The practitioners are most confident about one thing, that AI makes work faster, and this is exactly the claim the literature and the expert are most cautious about. The practitioners report dramatic individual gains, building alone in months what would once have taken a team a year. But these are first-person impressions of individual output rather than measured team performance. The expert does not deny that gains exist, but she questions whether they are real or only perceived, since no one measures them, and the time spent writing and rewriting prompts may eat into the time saved. The literature points the same way, finding that AI raises individual performance far more than team performance, and that on the wrong task it can make performance worse. On wellbeing the two sides converge again, agreeing that the strain and cognitive load of working at AI speed are real, with the expert adding concerns about junior staff and the relational side of work. For a company, the honest reading is that the speed gains are credible but unproven, and should not be booked as guaranteed return.

The conditions and moderators bring the comparison back to broad agreement. Governance and infrastructure produce the cleanest convergence in the whole cross-walk, since both sides agree that tool access and a workable approach to data protection are basic enablers. Leadership and structure converge as well, with the expert extending the point toward the policy gaps that open up when organizations transform. Among the moderators, team size is both the strongest finding and the one the research question turns on. The practitioners argued that AI shifts the optimal team size downward, often sharply, and that roles blur as developers move toward product ownership. This converges with the literature, which finds that individual-AI pairs can match whole human teams and that adding more brings diminishing returns. It also ties back to the input side, since as a team shrinks, a single weak member counts for more, so the case for a small, generalist team is also a case for getting its composition right. Task criticality and adoption maturity converge more quietly, with both sides agreeing that high-risk work still needs a human owner, and that a team only meets these dynamics once it is genuinely working AI-natively.

One thread runs through the whole comparison and helps explain it. One practitioner framed it as amplification, that AI magnifies the differences that are already there, good and bad. It is why an unmotivated person now produces more output but no better outcome, why a single weak member counts for more as teams shrink, and why a team that is not moving at one pace falls apart faster than it used to.

If AI amplifies what is already there, then the human characteristics on the input side matter more rather than less, which is the opposite of the worry that AI would make people interchangeable.

What it means in practice

The clearest message of the study is about hiring. As the execution side of the work collapses into AI tooling, the balance of hiring criteria shifts away from raw technical skill and toward disposition, meaning how a person thinks, takes ownership, and works. One practitioner put a number on it, weighting personality at seventy to eighty percent and technical skill at only twenty to thirty. The exact split is his alone, but the wider data supports the direction, if not the figure. For a company hiring into an AI-integrated team, this is the single most useful takeaway, that the question has moved from what a candidate can build by hand toward how they think and what they bring.

What to screen for follows from the findings. The characteristics the comparison placed at the front are agency and ownership, a product and value-creation mindset, AI literacy, calibrated judgment, and adaptability. It is important to mention here that these do not all carry the same weight of evidence. AI literacy, judgment, adaptability, and the classical personality core have validated research behind them, which makes them defensible criteria to hire on. Agency and the product mindset do not, since as the comparison showed, the literature has barely caught up with them. The practitioners stressed both more than anything else, but for now the case for hiring on them rests on that consensus rather than on tested evidence, so they are best treated as promising criteria still to be confirmed.

A word of caution goes the other way. The shift toward disposition is not a reason to stop valuing experience. Experience is what lets a person judge whether the AI's output is any good. A company that read this as a reason to hire only juniors and let AI fill the gap would be reading it wrong.

The data also names clear deal-breakers. The practitioners ruled out the toxic, dark-triad traits, and treated a flat refusal to work with AI at all as a reason not to hire. Since AI amplifies what is already there, a capable but toxic hire does more damage in an AI-native team than in a traditional one, which makes the negative screen matter as much as the positive one.

Hiring is not the whole story. The same findings argue for smaller, more generalist teams, for keeping a human owner on high-stakes work, and for treating tool access and workable governance as basic enablers. And since working at AI speed raises the cognitive load, wellbeing needs managing rather than assuming away.

Overall, these are signals to build into how a team is hired and structured, not a finished scorecard. The study describes what to look for and how much confidence to place in each criterion, but it stops short of a validated hiring tool, since designing and testing one is the work a follow-up would take on.

Limitations & how to read the findings

A few limitations shape how the findings should be read. The study is qualitative and exploratory, and it rests on Yin's (2018) replication logic rather than on sampling logic, so the findings are best understood as a grounded picture and a set of candidate characteristics, not as tested predictors. They describe what practitioners and the literature say drives performance, not the size of any effect, and they do not support statistical generalization.

Two caveats matter most. First, the performance gains are perceived rather than measured, since the study does not measure team output, so the speed claim in particular should be read as a strong and consistent impression rather than proof. Second, the academic side of the comparison leans on a single expert, which means one part of the cross-walk rests on one voice rather than many.

The sample carries its own limits. Eight interviews is small in absolute terms, the seven practitioners were selected on purpose as early adopters rather than drawn at random, all seven are male, most were interviewed in German, and six were conducted through the partner company rather than by the author, which introduces a possible interviewer effect. The interviews were also analyzed across a German-to-English translation. None of this undermines the patterns the study found, but it does bound the range of contexts the findings can speak to.

This leads to the most important point about how to read the report. The practitioners studied here are early adopters working AI-natively in tech and knowledge work. Their experience is best read as a leading indicator of where mainstream teams are heading, not as a description of all teams today. That

forecast is bounded in two ways the data itself names, by how far a team has matured in its AI adoption, and by the risk of the work, since the same hands-off use of AI that suits low-risk software would not suit a high-stakes domain. Whether the findings carry over to less tech-heavy functions is a theoretical question, since the data does not include those teams. Read with these bounds in mind, the study offers a credible forecast and a set of characteristics worth testing, rather than a finished or universal rule.

Conclusion & what comes next

This study set out to answer a single question: which human characteristics at the team level the academic literature and industry practitioners identify as drivers of performance and wellbeing in AI-integrated teams, and where the two perspectives converge or diverge. The short answer is that they converge on the people and diverge on the team. The literature and the practitioners largely agree on which human characteristics matter, and on the conditions and moderators around them. They part on the team-level processes that sit in between. AI does not change which characteristics matter so much as it changes their order, pushing agency, a product and value-creation mindset, AI literacy, and calibrated judgment to the front, while the classical personality core stays in place beneath them.

The study makes three contributions. The first is empirical, a grounded picture of what practitioners working AI-natively identify as the characteristics that make a team perform, drawn from eight interviews and organized on the IMOI framework. The second is the comparison itself, and it is the headline. By mapping the literature and the practitioner findings onto the same framework, the study offers a principled account of where the two agree, where one flags something the other misses, and where they genuinely diverge, and it is this comparison that surfaces the two characteristics the literature has yet to develop. The third is a foundation for future research, since the study identifies and organizes the characteristics a follow-up would need to test.

That follow-up is the natural next step. The findings here are best understood as candidate characteristics rather than tested predictors, and the obvious way forward is a quantitative study that tests which of them actually predict performance, and that develops measures for the two the literature cannot yet anchor, agency and the product mindset. The groundwork for that study is already in place. For the participating companies, though, the practical value does not wait on it, since the report already offers a clear picture

of the human qualities that make AI-integrated teams work, and a sharper set of questions to hire and build against.

Overall, a common worry about AI at work is that it makes people interchangeable, that as the tools do more, the differences between individuals matter less. The findings point the other way. AI did not flatten the importance of human characteristics in these teams. It reordered them, and in doing so it raised the stakes on getting them right, since the same technology that lets a motivated person achieve more also lets an unmotivated one produce more of the wrong thing. If that holds as AI spreads beyond the early adopters studied here, then the question of which human characteristics a team is built from will become more important, not less.