



International Journal of Research Publication and Reviews

Journal Homepage: www.ijrpr.com ISSN 2582-7421

AI in Mergers and Acquisitions Due Diligence: A Study of Perceived Artificial Intelligence Capability and Outcomes Among Finance Professionals

Chalasani Rushi Ram Chowdary¹ and Dr. Ajay Kumar Varshney²

¹ Research Scholar, Koneru Lakshmaiah Education Foundation (KLEF), Vaddeswaram, Andhra Pradesh, India

² Professor, Department of Business Administration, Koneru Lakshmaiah Education Foundation (KLEF), Vaddeswaram, Andhra Pradesh, India

Email: rushiramchalasani@gmail.com

DOI : <https://doi.org/10.55248/gengpi.07.1026.3237>

ABSTRACT

Due diligence is among the most information-intensive stages of a merger or acquisition, and artificial intelligence is increasingly proposed as a means of accelerating it and improving its quality. Evidence on how practitioners connect AI capability with specific due-diligence outcomes remains limited. This study investigates how finance and M&A professionals relate perceived AI capability to perceived efficiency, accuracy and risk identification in due diligence. Drawing on task–technology fit theory, the study uses an exploratory, cross-sectional survey of investment banking and M&A professionals. A structured five-point Likert questionnaire was analysed with reliability analysis, Pearson and Spearman correlations and sensitivity analyses, and the results were contextualised against published industry evidence. The study distinguishes perceived effectiveness from objectively measured performance and examines how consistently perceived capability is associated with each outcome. Based on these insights and the wider literature, the paper proposes a five-stage framework for governed, human-verified AI-assisted due diligence, covering task-fit screening, data readiness, AI first-pass analysis, human verification, and governance and performance measurement. The paper contributes to research on AI in M&A and technology fit, and offers investment banks, advisers and technology providers a practical basis for responsible AI adoption and a foundation for future empirical testing.

Keywords: artificial intelligence; mergers and acquisitions; due diligence; task–technology fit; practitioner perceptions; human–AI collaboration; AI governance

1. Introduction

Mergers and acquisitions allow firms to enter markets, acquire capabilities and pursue synergies, but the value of a transaction depends heavily on how accurately the acquirer understands the target before committing capital. Classic work on acquisition management emphasises that value is created or lost through the quality of pre-deal understanding and post-deal integration (Haspeslagh & Jemison, 1991). Due diligence is the stage at which that understanding is built. Financial statements, contracts, tax records, regulatory filings, intellectual-property documents and operational data must be reviewed under tight deadlines by teams of bankers, accountants, lawyers and consultants. Because the acquirer typically knows less about the target than the seller does, the quality of the information obtained in this stage directly shapes pricing, deal terms and the decision to proceed, and the volume of material makes the process slow and creates a risk that relevant information is missed (Angwin, 2001).

Artificial intelligence, and in particular natural language processing (NLP), machine learning and generative models, is increasingly applied to document-heavy analysis of this kind. Such tools can in principle locate clauses, extract data points, compare documents across a data room and flag discrepancies for professional review. Consultancy surveys indicate that adoption is moving quickly. Deloitte's 2025 survey of 1,000 senior corporate and private-equity leaders in the United States, conducted in the first half of 2025, reported that 86% of responding organisations had integrated generative AI into M&A workflows, 65% of them within the previous year (Deloitte, 2025a, 2025b). Bain & Company's survey of 300 M&A executives reported that 45% used AI tools in M&A in 2025, more than double the prior year (Bain & Company, 2026). A McKinsey survey reported that respondents using generative AI in M&A cited average cost reductions of roughly 20% and, for 40% of respondents, deal cycles that were 30–50% faster, while only about 30% of respondents engaged with generative AI at moderate to high levels (McKinsey & Company, 2026). These figures are self-reported by advisory firms' survey respondents, use different definitions of AI and adoption, and do not measure verified performance.

The same sources, and the wider literature, point to risks. Data security, data quality and model reliability are among the most frequently cited concerns (Deloitte, 2025a; PE Professional, 2025); legal analysis argues that reliance on flawed AI-generated information can undermine the validity of transaction agreements and engage the liability of tool providers and directors (Herbosch & Mertens, 2025); and research on automation shows that human reviewers may over-rely on automated output unless the system and the review process are designed to calibrate trust (Lee & See, 2004;

Parasuraman & Manzey, 2010). An experiment with consultants likewise found that generative AI raised performance on some tasks and lowered it on a task outside the model's capability frontier (Dell'Acqua et al., 2023). Whether AI helps in a given due-diligence task therefore depends on task fit, data readiness and verification, not only on tool capability.

This matters from a finance and investment-banking perspective for three reasons. Diligence errors can transmit into valuation and deal terms; time pressure makes the efficiency claims commercially attractive; and the personnel who will rely on or supervise AI output are the practitioners whose beliefs shape adoption, training and governance. Those beliefs are the focus of this paper. A central distinction runs through it: perceived effectiveness (what practitioners believe AI can do) is not measured effectiveness (what AI demonstrably does against a verified baseline). A survey of opinions cannot establish that AI improves transaction outcomes.

Research gap. Academic research on AI in M&A is developing but uneven (Section 2). Firm-level studies examine AI capability and acquisition strategy; an experimental study examines AI agents' evaluation of deals; legal scholarship examines liability; and conceptual papers describe potential applications. Industry surveys report adoption and expected benefits. Within the sources reviewed for this paper, we did not find a study that measures, with separate multi-item scales, how practitioners connect perceived AI capability to the specific outcomes of due diligence (efficiency, accuracy and risk identification) and reports the reliability and stability of those associations. This is a statement about the targeted search described in Section 2.8, not a claim that no such study exists. The gap addressed is therefore primarily one of practitioner evidence and measurement; the paper does not claim to fill a gap in objectively measured AI performance, which remains open.

Research questions. RQ1: How favourably do respondents rate AI capability and its due-diligence outcomes? RQ2: Is perceived AI capability associated with perceived efficiency, accuracy and risk identification, and how stable are these associations across samples and respondent groups? RQ3: Which barriers do respondents emphasise, how do they view human verification, and are these views consistent in direction with published industry evidence?

Contributions. First, the study measures perceived capability and three perceived outcomes with separate multi-item scales and reports their reliability. Second, it tests the capability–outcome associations and examines how stable they are when respondents with identical answers to all items are set aside, showing that two of three associations are fragile and one is likely inflated by measurement overlap. Third, it reports exploratory descriptive comparisons by AI use, age and AI familiarity. Fourth, it proposes a five-stage framework for governed, human-verified AI-assisted due diligence together with measurable indicators; the framework is derived from the findings and the literature and has not been empirically validated. The paper proceeds as follows: Section 2 reviews the literature; Section 3 presents the theoretical framework and hypotheses; Section 4 describes the methodology; Section 5 reports the results; Section 6 discusses them; Section 7 presents the proposed framework and implications; Section 8 covers limitations and future research; and Section 9 concludes.

2. Literature Review

2.1 Due diligence and information asymmetry

Acquirers face information asymmetry about targets, and due diligence is the principal mechanism for reducing it. Angwin (2001) shows that the conduct of pre-acquisition due diligence and the use of professional advisers vary with national and organisational context, so diligence is shaped by more than technical analysis. Two features matter for AI. Diligence is document-intensive and time-boxed, so analytical capacity is a binding constraint; and it is judgement-intensive, so errors of interpretation can be as costly as omissions (Haspeslagh & Jemison, 1991). The first feature suggests where AI may help; the second suggests why human verification remains central.

2.2 AI across the M&A process: empirical and conceptual evidence

Firm-level evidence suggests AI affects which deals firms pursue as well as how they analyse them. Fang et al. (2026), a working paper that measures AI capability through employees' job skills, report that AI-intensive acquirers are more likely to acquire data-rich targets and hire data-analytics specialists, and show stronger deal performance without paying higher premiums. Liu et al. (2026), drawing on a resource-based view and published in *International Review of Financial Analysis*, report that AI adoption is associated with a higher likelihood of undertaking M&A, with information transparency and resource-reallocation efficiency as proposed mechanisms. Both studies are archival and concern firms' strategic choices and outcomes, not the conduct of due diligence tasks or practitioners' beliefs. Conceptual contributions, such as Rashid et al. (2025), describe how machine learning and NLP could support target identification and document analysis, but rest on secondary sources and report no primary data; they are best read as statements of expectation.

2.3 AI in due diligence and deal evaluation

Evidence on AI's performance in evaluation and diligence tasks is thin. Mirzayev et al. (2025), an SSRN working paper, assessed anonymised announcements of 109 deals between US public firms and report that multi-agent large-language-model systems showed potential for distinguishing value-creating from value-destroying deals; because the data are announcement-stage and the study is not peer reviewed in the version consulted, it indicates capability in a stylised evaluation setting rather than performance in live diligence. Thelisson (2024) examines, through a single in-depth case, how AI tools can support the screening and evaluation phases of a takeover, and Bedekar et al. (2024), a conference paper, discuss the potential of

machine learning, NLP and predictive analytics to improve decision accuracy and reduce processing time in due diligence. Disu et al. (2022) report efficiency and accuracy gains from AI in investment banking, but some headline figures in that paper differ between its abstract and tables, and it appeared in a venue outside the core finance literature, so it is treated here as indicative only. Taken together, these studies show plausible capability and early evidence, not verified field performance.

2.4 Practitioner and industry evidence

Large-sample industry surveys add adoption and expected-benefit data (Table 1). They indicate rapid uptake, concentration of use in early deal stages, and recurring concerns about data security, data quality and reliability. They are produced by advisory firms with commercial interests, are not peer reviewed, differ in population, wording and definitions, and generally report marginal percentages rather than relationships between beliefs. They therefore provide context and a benchmark for direction, not effect sizes. Trade reporting adds that investment banks are building proprietary deal datasets for AI tools while continuing to hire analysts, and that junior roles may shift towards analysis (Gillespie, 2025; Stewart, 2024); these are journalistic accounts and are used here only as background.

Table 1. Secondary evidence on AI adoption and benefits in M&A (figures as published; not pooled or directly comparable)

Source (type)	Population and timing	Findings relevant to due diligence
Deloitte (2025a, 2025b); consultancy survey	1,000 senior corporate and PE leaders, U.S.; first half of 2025	86% had integrated generative AI into M&A workflows, 65% of them within the past year; data security (67%) and data quality and availability (65%) were leading concerns (concerns as reported by PE Professional, 2025)
Bain & Company (2026); press release	300 M&A executives; survey reported January 2026	45% used AI tools in M&A in 2025, more than double the prior year; about one in three deploying AI systematically or redesigning processes
Bain & Company (2024); consultancy report	306 M&A practitioners; early 2024	16% had relied on generative AI in past deals; 80% expected to deploy within three years; data inaccuracy, privacy and cybersecurity were frequently cited risks; one user noted checking output can take as long as the time saved
McKinsey & Company (2026); consultancy article	Survey conducted in 2025; size not stated in the page consulted	Respondents using generative AI reported average cost reduction of roughly 20%; 40% reported 30–50% faster deal cycles; about 30% engaged at moderate to high levels; lack of expertise was the main adoption challenge
Alexis (2026a), reporting an Accenture survey	650 senior dealmakers, 12 industries, 24 countries; January 2026	Share investing in generative AI for pre-deal M&A rose to 46% in 2026 from 31% in 2024

Note. The sources differ in definitions (generative AI versus AI, analytics and automation), samples, countries and dates. Cost and time savings are self-reported by users and are not independently measured outcomes. Trade coverage of the McKinsey survey (Alexis, 2026b) summarises timeline savings as 10–30%, which differs from McKinsey's own text; the McKinsey figures are used here. The Accenture figures are taken from trade reporting, not the original report.

2.5 Human–AI collaboration, governance and liability

Whether AI output can be relied upon depends on the information environment and the oversight arrangements around it. Kotla (2026) discusses AI-based entity resolution and schema mapping for post-merger data integration in a conceptual, untested treatment. Herbosch and Mertens (2025), in a working paper from Ghent University's Financial Law Institute, analyse several jurisdictions and argue that reliance on flawed AI-generated information can affect the validity of transaction agreements and the liability of tool providers and directors; they identify system selection, supervision and explainability as safeguards. The human-factors literature explains why verification must be designed rather than assumed: trust in automation should be calibrated to its actual reliability (Lee & See, 2004), and operators can show complacency and automation bias when automation is usually right (Parasuraman & Manzey, 2010). The field experiment of Dell'Acqua et al. (2023) shows that AI's benefits are uneven across tasks, which supports task-level screening and measurement. Institutional guidance points the same way: the NIST AI Risk Management Framework emphasises governance, mapping of context, measurement and management of AI risk (National Institute of Standards and Technology [NIST], 2023); the Financial Stability Board has discussed data, model-risk and third-party dependencies in financial uses of AI (Financial Stability Board, 2024); and the EU Artificial Intelligence Act establishes risk-based obligations that may apply to AI systems depending on use (European Union, 2024). These documents are used here as governance context, not as evidence about due-diligence performance.

2.6 Theoretical lens: task–technology fit and acceptance

Task–technology fit (TTF) holds that a technology improves individual performance when its capabilities match the demands of the task (Goodhue & Thompson, 1995). Due diligence is document-intensive and pattern-seeking, so a professional who believes AI can analyse large volumes of material, review financial and legal documents and flag anomalies would be expected also to expect gains on outcomes such as efficiency, accuracy and risk identification. Technology-acceptance research treats perceived usefulness as a core driver of adoption (Davis, 1989; Venkatesh et al., 2003). We use TTF as the primary lens and refer to acceptance research only to note that both traditions concern perceived rather than realised performance. The present study can therefore speak to attitudinal coherence and adoption readiness, not to measured outcomes.

Table 2. Synthesis of the academic and conceptual studies reviewed

Study	Type and design	Main contribution	Limitation for present purposes
Rashid et al. (2025)	Conceptual review	ML and NLP could support target identification and due diligence	Secondary sources only; no primary data
Fang et al. (2026)	Working paper; archival firm-level	AI-capable acquirers target data-rich firms and show stronger deal outcomes	AI measured through skills; not diligence tasks or perceptions
Liu et al. (2026)	Journal article; archival firm-level	AI adoption linked to M&A likelihood via transparency and reallocation	Deal propensity, not diligence tasks
Mirzayev et al. (2025)	Working paper; experiment, 109 anonymised deal announcements	Multi-agent systems showed potential in distinguishing deal value outcomes	Not peer reviewed; announcement-stage data, not live diligence
Thelisson (2024)	Journal article; single in-depth case	AI support for screening and evaluation in a takeover	Single case
Bedekar et al. (2024)	Conference paper; discussion of potential	ML/NLP could improve accuracy and reduce processing time	Limited empirical detail in the version cited
Disu et al. (2022)	Mixed methods, investment banks	Reports efficiency and accuracy gains	Inconsistencies in reported figures; treated as indicative
Herbosch & Mertens (2025)	Working paper; comparative legal analysis	Liability and validity risks when AI output is flawed	Normative; no practitioner data
Kotla (2026)	Conceptual architecture	Entity resolution and system consolidation for integration	Not empirically tested
Dell’Acqua et al. (2023)	Working paper; field experiment with consultants	AI raised performance on some tasks and lowered it on a task outside its capability frontier	Consulting tasks, not M&A diligence

Note. Publication type is stated as it could be established; “working paper” denotes a non-peer-reviewed version.

2.7 Synthesis and research gap

Three observations follow. First, the academic work reviewed examines whether AI changes deal selection and performance, how it might be used, or what legal conditions govern it, but not how practitioners connect perceived capability to diligence outcomes. Second, industry surveys agree in direction on rapid adoption and on data and security concerns but differ in definitions, report mainly marginal percentages, and do not report reliability or the stability of associations. Third, no objective, field-based comparison of manual and AI-assisted diligence performance appeared in the sources reviewed. The present study addresses the first observation through a small exploratory survey; it does not address the third, which it identifies as the priority for future work. The study’s claim to originality is therefore modest: a measurement-focused, practitioner-level examination of perceived capability and perceived outcomes, with explicit stability checks.

2.8 Search procedure

Secondary sources were identified through a targeted, non-systematic search carried out in October 2026 for surveys and studies by major advisory firms on AI in M&A and due diligence, retaining sources whose publisher, sample and date were identifiable. Key industry figures were checked

against the publishers' own pages or press releases (Table 1), and additional literature on human–AI interaction, governance and technology fit was added where it directly supported an argument. Inclusion required relevance to AI in M&A, due diligence, human–AI collaboration or governance, and an identifiable publisher. No database search protocol or PRISMA procedure was followed, and the review should be read as a structured, targeted review.

3. Theoretical Framework and Hypotheses

Figure 1 shows the conceptual framework. Perceived AI capability, defined as the belief that AI can analyse large volumes of information, review financial and legal documents and identify anomalies, is expected to be associated with perceived efficiency, perceived accuracy and perceived risk identification. Under TTF, perceived fit between technology capability and task demands should accompany expectations of better task outcomes. Perceived efficiency follows because document volume and deadlines are the binding constraints that AI is believed to relax; perceived accuracy because pattern matching and consistency are believed to reduce omission and transcription errors; and perceived risk identification because anomaly flagging across large document sets is believed to surface items that conventional sampling may miss. Data quality, privacy, cost and the need for human verification are measured descriptively and are not modelled as moderators, because the sample is too small to test moderation.




Figure 1. Conceptual framework. Arrows denote hypothesised associations between perceived constructs, not causal paths.

Three research questions (Section 1) correspond to three study objectives: to describe practitioners' perceptions; to test and stress-test the associations between perceived capability and perceived outcomes; and to document perceived barriers and the perceived need for human verification. Following the TTF logic, three directional hypotheses are tested:

H1: Perceived AI capability is positively associated with perceived efficiency of M&A due diligence.

H2: Perceived AI capability is positively associated with perceived accuracy of due-diligence analysis.

H3: Perceived AI capability is positively associated with perceived effectiveness of risk identification.

Because predictor and outcomes are self-reported by the same respondents on the same instrument, support for a hypothesis is judged jointly on statistical significance, stability under sensitivity analysis and consistency of direction with independent secondary evidence. The hypotheses are directional, but all *p*-values reported are two-sided, which is the more conservative choice. Support for a hypothesis is support for an association between beliefs; it is not evidence that AI improves efficiency, accuracy or risk identification.

4. Research Methodology

4.1 Research design

The study is an exploratory, cross-sectional survey, appropriate where evidence is scarce and the aim is to describe perceptions and examine preliminary relationships (Kothari, 2004). A secondary component used published industry evidence for contextual triangulation. It is not a mixed-methods study in the sense of integrated qualitative and quantitative data collection, and the two components are not statistically combined.

4.2 Target population, sampling and recruitment

Purposive sampling targeted people with knowledge of or exposure to finance, investment banking, M&A or consulting, because the research question concerns perceptions of people who conduct or advise on due diligence and such a population has no accessible sampling frame. Respondents were recruited through LinkedIn, using a public post and direct messages to professionals. No more than about 50 people were invited, so the 20 records received correspond to a response rate of roughly 40% or higher; the exact number invited was not tracked, and the reach of the public post cannot be counted. Purposive, social-media-based recruitment does not support population inference: respondents were selected by accessibility, relevance and willingness to respond, and those who respond to an AI-themed survey may hold more favourable views than those who do not. The questionnaire was administered online in September–October 2026. Twenty records were received; one contained no substantive responses and was excluded, leaving 19 valid cases (17 investment banking or M&A professionals, one consulting professional and one undergraduate student). Eligibility was defined by exposure to finance or M&A, and 18 of the 19 respondents reported exposure to M&A due diligence.

4.3 Instrument and constructs

The questionnaire collected demographic and experience data (age, current status, educational background, familiarity with M&A and AI, exposure to due diligence and use of AI tools), followed by 20 statements rated from 1 (strongly disagree) to 5 (strongly agree). Items were developed by the authors from capability, outcome, barrier and governance themes in the literature (Bedekar et al., 2024; Fang et al., 2026; Herbosch & Mertens, 2025; Rashid et al., 2025). The questionnaire was reviewed by an academic supervisor before distribution; it was not reviewed by a practitioner panel and was not pilot-tested. The instrument is therefore a purpose-built scale that has not been formally validated. Table 3 shows the constructs and summarises the content of each item. Item 4 (self-rated knowledge) and item 16 (human verification) measure, respectively, respondents' own knowledge and a normative view of oversight, and are analysed separately from the capability and barrier scales. Construct scores are the mean of their items.

Table 3. Constructs and items

Construct	Items	Content
Perceived AI capability (CAP)	1–3	Analysing large volumes of information; reviewing financial and legal documents; identifying anomalies and red flags
Self-rated knowledge (KNOW)	4	Sufficient knowledge of AI applications in due diligence
Efficiency (EFF)	5–7	Reduced time; reduced manual work; overall process efficiency
Accuracy (ACC)	8–10	Information extraction; fewer human errors; consistency and reliability
Risk identification (RISK)	11–13	Financial risks; legal, regulatory and compliance risks; hidden red flags
Barriers (BAR)	14, 15, 17	Data quality; privacy and security; cost and complexity
Human verification (HV)	16	Findings require human verification before important decisions
Usefulness / future (USE)	18–20	Usefulness; expected growth in use; human–AI collaboration

Note. Item 20 conflates human–AI collaboration with usefulness; the USE scale should be read as a general favourability index.

4.4 Questionnaire audit

The instrument was audited after data collection. All 20 statements are positively worded, which invites acquiescence and does not allow reverse-keyed consistency checks; the barrier items are also keyed in the direction of agreement, so agreement to all items is not logically inconsistent. Several items combine concepts: item 2 pairs financial and legal documents, item 12 pairs legal, regulatory and compliance risk, and item 9 pairs reducing human errors with large data volumes. Capability and outcome items are semantically close (for example, item 3 asks whether AI can identify anomalies and red flags and item 13 whether it can identify hidden red flags), which plausibly explains part of the very high capability–accuracy and capability–risk associations. Self-rated knowledge (item 4) is a single item and is interpreted descriptively. The barrier scale mixes data quality, privacy and cost, which may not form a unitary construct despite its alpha. Social desirability may also operate, because endorsing AI may signal professional currency. These are limitations of the instrument as administered; the items were not rewritten after data collection. Improvements to the instrument for future research are outlined in Section 8.2; they do not describe the instrument the respondents answered.

4.5 Data cleaning and missing data

Likert responses were coded 1–5 and construct scores computed as item means. One respondent left item 11 blank, and that respondent's risk-identification score was computed from the two answered items; the reliability estimate for the risk scale uses the 18 complete cases. No other missing

data were reported. Respondents who gave identical answers to all 20 items (nine respondents: seven answering “strongly agree” and two “agree” throughout) were retained in the primary analysis and set aside only in a pre-specified sensitivity analysis.

4.6 Analytical procedure

The analysis reported the mean, standard deviation, median and percentage agreement (responses of 4 or 5). Internal consistency was assessed with Cronbach’s alpha, using .70 as a conventional reference while noting that alpha depends on the number of items and should not be applied mechanically (Nunnally, 1978; Tavakol & Dennick, 2011). Associations between perceived capability and each outcome were estimated with Pearson’s r , with 95% confidence intervals from the Fisher z transformation, and with Spearman’s ρ as a rank-based check suited to ordinal composites with ties; percentile bootstrap intervals (10,000 resamples) gave similar conclusions. A Friedman test compared the three outcome scores. Two sensitivity analyses were specified: (a) excluding respondents who gave the same answer to all 20 items, and (b) restricting the sample to investment banking and M&A professionals. Descriptive group comparisons used Hedges’ g and exact Mann–Whitney tests and are reported as exploratory because groups are small and 21 comparisons were made (Bonferroni threshold $p < .0024$). Because predictor and outcomes come from a single instrument and respondent, common-method variance is a concern (Podsakoff et al., 2003). Analyses were run in Python (SciPy).

4.7 Ethical considerations

Participation was voluntary, and consent was implied by completing the online questionnaire; no separate written consent was recorded. Responses are analysed in aggregate without identifying information. No institutional ethics approval or exemption is documented for this study, which is a limitation of the research record.

5. Results

5.1 Respondent characteristics

Table 4 describes the sample ($N = 19$). It is experienced and technologically engaged: 15 respondents described themselves as extremely familiar with M&A, 12 as extremely familiar with AI, and all 19 had used AI tools. Regular use (73.7%) is higher than the share of executives reporting AI use in Bain’s 2025 survey (45%; Bain & Company, 2026), but the populations and definitions differ, so the comparison is contextual only. The sample is likely to be favourably selected toward AI.

Table 4. Respondent profile ($N = 19$)

Characteristic	Category	n (%)
Current status	Investment banking / M&A professional	17 (89.5)
	Consulting professional	1 (5.3)
	Undergraduate student	1 (5.3)
Age (years)	20–24	4 (21.1)
	25–29	4 (21.1)
	30–39	6 (31.6)
	40–49	3 (15.8)
	50 and above	2 (10.5)
Background	Finance	16 (84.2)
	Economics, management, engineering	3 (15.8)
Familiarity with M&A	Extremely / very familiar	18 (94.7)
	Not at all familiar	1 (5.3)
Familiarity with AI	Extremely / very familiar	16 (84.2)
	Moderately familiar	3 (15.8)

Characteristic	Category	n (%)
Use of AI tools	Regularly	14 (73.7)
	Basic tasks only / occasionally	5 (26.3)
Exposure to M&A due diligence	Yes	18 (94.7)

Note. Percentages may not sum to 100 because of rounding. Country, firm type and seniority were not recorded.

5.2 Descriptive statistics and response patterns

Perceptions were uniformly favourable (Table 5; Figure 2). Construct means ranged from 4.33 (barriers) to 4.63 (usefulness/future); 60.4% of all item responses were “strongly agree” and 91.6% were “agree” or above. The three outcome means (efficiency 4.60, accuracy 4.51, risk identification 4.45) did not differ reliably (Friedman $\chi^2(2) = 5.69, p = .058$), so no ranking of outcomes is claimed. The concentration of answers at the top of the scale means there is little variance to explain, which limits what correlations can show.

Table 5. Descriptive statistics and reliability (N = 19)

Construct	k	M	SD	Mdn	% agree	α
Perceived AI capability	3	4.42	0.66	4.67	93.0	.87
Self-rated knowledge	1	4.26	0.99	5.00	84.2	–
Efficiency	3	4.60	0.53	5.00	94.7	.85
Accuracy	3	4.51	0.58	4.67	91.2	.76
Risk identification	3	4.45	0.60	4.67	92.9	.68
Barriers	3	4.33	0.90	5.00	82.5	.87
Human verification	1	4.58	0.69	5.00	89.5	–
Usefulness / future	3	4.63	0.46	5.00	98.2	.83

Note. k = number of items; % agree = share of item responses scored 4 or 5; α for risk identification uses 18 complete cases.




Figure 2. Distribution of responses by construct (SD = strongly disagree, D = disagree, N = neutral, A = agree, SA = strongly agree; neutral centred at 0).

At item level, the highest-rated statement was that effective due diligence will combine AI and human professionals ($M = 4.68$), followed by expected growth in AI use ($M = 4.63$) and the statements that AI can reduce manual work ($M = 4.63$) and reduce human errors ($M = 4.63$). Every capability and outcome item was endorsed by at least 84% of respondents. The lowest agreement was for the privacy and security barrier ($M = 4.21$, 78.9% agree) and the cost and complexity barrier ($M = 4.26$), the two items with the widest spread ($SD = 1.13$ and 1.10), showing that barriers are where respondents differ most. Poor data quality was the strongest barrier ($M = 4.53$).

5.3 Measurement reliability

Internal consistency was good for perceived capability ($\alpha = .87$), efficiency (.85), barriers (.87) and usefulness (.83), acceptable for accuracy (.76), and below the conventional benchmark for risk identification (.68). With three items and 19 cases these estimates are imprecise, and alpha for three-item scales is sensitive to a single item; the risk-identification scale is therefore retained but its results are read with caution. Self-rated knowledge (item 4) correlated $r = .93$ with the capability scale, so respondents who saw AI as more capable also described themselves as more knowledgeable; this overlap suggests a shared general attitude. Using the original four-item capability scale (including item 4; $\alpha = .92$) changed the main correlations by less than .03.

5.4 Correlations between perceived capability and perceived outcomes

Table 6. Correlations of perceived AI capability with perceived outcomes

Outcome	Sample	n	r [95% CI]	p	Spearman p
Efficiency	All	19	.68 [.32, .87]	.001	.69
	Excl. uniform	10	.46 [-.24, .84]	.182	.44
	Professionals	17	.65 [.24, .86]	.005	.63
Accuracy	All	19	.96 [.90, .99]	< .001	.93
	Excl. uniform	10	.94 [.75, .99]	< .001	.91
	Professionals	17	.93 [.82, .98]	< .001	.91
Risk identification	All	19	.82 [.59, .93]	< .001	.81
	Excl. uniform	10	.64 [.02, .91]	.045	.61
	Professionals	17	.74 [.40, .90]	< .001	.72

Note. All = full sample; Excl. uniform = excluding nine respondents who gave identical answers to all 20 items; Professionals = investment banking / M&A professionals only. p -values are two-sided Pearson tests. The Spearman p -value for risk identification excluding uniform responders is .060. Confidence intervals are Fisher- z intervals.

In the full sample, perceived capability was positively and significantly associated with efficiency ($r = .68$, $p = .001$), accuracy ($r = .96$, $p < .001$) and risk identification ($r = .82$, $p < .001$), and Spearman coefficients were similar (Table 6). The associations were not equally stable. Nine of the 19 respondents gave identical answers to every item, including the barrier items. These respondents may genuinely hold uniform views, so their data are not discarded as invalid, but straight-line responding carries little item-level information and removes the within-person variation that distinguishes constructs. With these respondents set aside ($n = 10$), the accuracy association remained very strong ($r = .94$), the risk-identification association weakened ($r = .64$, 95% CI [.02, .91]; Spearman $p = .060$) and the efficiency association was no longer significant ($r = .46$, 95% CI [-.24, .84], $p = .182$; Figure 3). The interpretation of this attenuation requires care: removing respondents who sit at the top of every scale also restricts the range of the variables, which mechanically lowers correlations, and the sub-sample of 10 yields very wide intervals. The sensitivity analysis therefore shows that two of the three associations are not robust in this sample; it does not show that the true associations are weak or absent. Restricting the sample to professionals ($n = 17$) left the estimates close to the full-sample values, but this restriction removed only two respondents and is a weak test of robustness.




Figure 3. Correlations of perceived AI capability with perceived outcomes in the full sample, excluding uniform responders, and among professionals only (Pearson r, 95% CI). Plotted from the values in Table 6.

Figure 4 shows the correlation matrix of construct scores. The three outcomes were highly intercorrelated ($r = .82-.89$), usefulness correlated strongly with all of them ($r = .71-.87$), and barriers were only weakly related to every other construct ($r = .18-.33$). This pattern is consistent with respondents holding a single favourable view of AI's benefits and a separate, more varied view of its constraints; it cannot be taken to show that the benefit constructs are empirically distinct.




Figure 4. Correlation matrix of construct scores (Pearson r, N = 19). CAP = capability; EFF = efficiency; ACC = accuracy; RISK = risk identification; BAR = barriers; USE = usefulness/future.

5.5 Exploratory group comparisons

Table 7. Mean construct scores by respondent group (descriptive)

Construct	Regular (n=14)	Occasional (n=5)	<i>g</i>	Under 30 (n=8)	30+ (n=11)	<i>g</i>	Extremely familiar (n=12)	Very/mod. (n=7)	<i>g</i>
Capability	4.62	3.87	1.25	4.50	4.36	0.19	4.67	4.00	1.09
Efficiency	4.67	4.40	0.48	4.63	4.58	0.09	4.61	4.57	0.07
Accuracy	4.67	4.07	1.08	4.58	4.45	0.21	4.67	4.24	0.74
Risk identification	4.62	3.97	1.16	4.40	4.48	-0.14	4.60	4.19	0.67
Barriers	4.38	4.20	0.19	3.71	4.79	-1.40	4.33	4.33	0.00
Human verification	4.71	4.20	0.73	4.25	4.82	-0.84	4.50	4.71	-0.29
Usefulness / future	4.71	4.40	0.67	4.63	4.64	-0.02	4.69	4.52	0.35

Note. *g* = Hedges' *g* (first group minus second group within each pair). Comparisons are exploratory: groups are small, 21 comparisons were made, and no difference is significant after Bonferroni correction (threshold $p < .0024$). Unadjusted exact Mann-Whitney *p*-values for the largest differences: capability by AI use $p = .059$; capability by AI familiarity $p = .038$; barriers by age $p = .014$.

Three patterns are hypothesis-generating only. First, regular AI users rated capability (4.62 vs 3.87), accuracy (4.67 vs 4.07) and risk identification (4.62 vs 3.97) higher than occasional or basic users ($g = 1.08$ to 1.25), whereas efficiency differed little (4.67 vs 4.40); respondents who were extremely familiar with AI likewise rated capability higher (4.67 vs 4.00, $g = 1.09$). These patterns are consistent with experience raising perceived fit but equally with self-selection, since people who favour AI are more likely to use it. Second, respondents under 30 rated barriers lower than those aged 30 and above (3.71 vs 4.79, $g = -1.40$) and endorsed human verification less strongly (4.25 vs 4.82), while their views on capability and outcomes were almost identical. This could reflect greater comfort with AI limitations, less exposure to the consequences of errors, or chance. Third, barrier scores differed little by AI use and familiarity, consistent with data, privacy and cost concerns being shared across levels of experience. With groups of 5 to 14, these estimates are very imprecise, and the unadjusted *p*-values should not be interpreted as evidence of differences.

5.6 Comparison with secondary evidence

Table 8 sets each hypothesis and the barrier findings beside published evidence. The comparison is one of direction and context, not of magnitude. The survey populations, definitions and wording differ, and the secondary figures are self-reported by advisory-firm respondents. Direction is broadly consistent, but the secondary sources report moderate savings and substantial caveats, which tempers the near-ceiling agreement in this sample.

Table 8. Contextual comparison of survey findings with secondary evidence

Item	This survey	Contextual secondary evidence	Caution	Assessment
H1 Efficiency	M = 4.60 (94.7% agree); $r = .68$ (all); $r = .46$, n.s. (excl. uniform)	Users reporting average cost reduction of about 20% and faster deal cycles (McKinsey & Company, 2026); rising adoption (Bain & Company, 2026)	Self-reported savings are moderate; a Bain user noted checking output can take as long as the time saved (Bain & Company, 2024)	Direction consistent; association not robust here
H2 Accuracy	M = 4.51; $r = .96$ (.94 excl. uniform)	Multi-agent systems showed potential in stylised deal evaluation (Mirzayev et al., 2025)	Data inaccuracy a frequent risk (Bain & Company, 2024); flawed output creates liability (Herbosh & Mertens, 2025); survey measures overlap	Supported statistically; read as overlapping measures
H3 Risk identification	M = 4.45; $r = .82$ (.64 excl. uniform)	AI evaluators added predictive information in an experiment (Mirzayev et al., 2025)	Risk scale reliability below benchmark ($\alpha = .68$)	Tentative

Item	This survey	Contextual secondary evidence	Caution	Assessment
Barriers	Data quality $M = 4.53$; privacy $M = 4.21$; cost $M = 4.26$	Data security (67%) and data quality and availability (65%) among top concerns (Deloitte, 2025a; PE Professional, 2025)	Different item wording and populations	Consistent in direction
Human verification	$M = 4.58$; collaboration $M = 4.68$	Supervision and explainability as safeguards (Herbosch & Mertens, 2025); calibrated trust (Lee & See, 2004)	Verification time can erode gains	Consistent in direction
AI use	73.7% regular users	86% integration (Deloitte, 2025a); 45% use (Bain & Company, 2026)	Definitions, samples and dates differ; sample here is AI-engaged	Consistent with a high-adoption sample

Note. n.s. = not significant. Secondary figures are as published and are not pooled with survey results.

5.7 Hypothesis assessment

H1 is supported in the full sample but is not statistically robust: the association loses significance and its interval spans zero once uniform responders are excluded. H2 is statistically supported and stable, but $r = .96$ (about 92% shared variance) indicates that capability and accuracy items come close to measuring the same belief, so it should not be read as a large independent effect. H3 receives tentative support: the association is large in the full sample and remains positive after exclusions, but it is imprecise and the risk scale has low reliability. Secondary evidence is consistent in direction with all three hypotheses but does not corroborate the magnitude of survey agreement or the specific correlations.

6. Discussion

Favourable sentiment. Respondents agreed almost uniformly that AI can analyse large volumes of information, review documents and surface red flags, and that it improves efficiency, accuracy and risk identification. This is consistent in direction with the rapid adoption reported by advisory firms (Bain & Company, 2026; Deloitte, 2025a). Two features of the sample probably amplify optimism: most respondents were experienced, regular AI users, and people with favourable views are more likely to use AI and to answer an AI survey. The group comparisons (exploratory) point in the same direction.

Capability–outcome associations. TTF would predict that practitioners who see AI as better matched to document-heavy diligence tasks also expect better outcomes, and the full-sample correlations are in that direction. The associations are, however, partly a product of a shared general attitude. The accuracy correlation approaches unity; the outcome scores are highly intercorrelated; self-rated knowledge correlates .93 with capability; and the efficiency and risk-identification associations depend on respondents who answered uniformly. Because predictor and outcomes were measured by the same instrument and respondents, common-method variance, acquiescence and consistency motives are plausible contributors (Podsakoff et al., 2003). Conceptual overlap between items, for example items 3 and 13, is a further plausible contributor. The data do not allow these explanations to be separated, and this paper does not claim to have established the cause of the overlap. The accurate reading is that the data show attitudinal coherence, not that perceived capability produces better outcomes, and certainly not that AI delivers the gains perceived. The attenuation after excluding uniform responders is informative but ambiguous: it may reflect range restriction, lower information content in straight-lined answers, or a genuinely weaker association among respondents with differentiated views. A study with more varied responses and mixed-keyed items could distinguish them.

Perceived versus reported gains. Respondents agreed almost unanimously that AI improves efficiency (94.7% agree), whereas the best available consultancy evidence reports average cost reductions of roughly 20% among users and a minority reporting 30–50% faster cycles (McKinsey & Company, 2026), and practitioners in Bain's 2024 study noted that checking AI output can consume the time saved (Bain & Company, 2024). These measures differ in kind, so they do not contradict each other, but they suggest that near-ceiling agreement does not translate into near-ceiling realised gains and that verification effort is a hidden cost. Reported savings are also self-reported and commercially sourced. The practical implication is that gains depend on how verification is organised and measured, not only on tool capability, which is consistent with evidence that AI's benefit varies by task (Dell'Acqua et al., 2023).

Oversight, data quality and governance. Optimism coexisted with strong endorsement of oversight and awareness of constraints. Human verification ($M = 4.58$) and human–AI collaboration ($M = 4.68$) were among the most strongly endorsed items, and poor data quality was the most strongly endorsed barrier. This is consistent in direction with legal analysis treating supervision and explainability as safeguards (Herbosch & Mertens, 2025), with human-factors research on calibrated trust and automation bias (Lee & See, 2004; Parasuraman & Manzey, 2010), and with industry surveys naming data security and quality as leading concerns (Deloitte, 2025a; PE Professional, 2025). It does not show that respondents' organisations practise verification. The age pattern (exploratory) suggests that training and review protocols, not only tools, will determine how safely AI is used as junior staff take on more AI-assisted work, but that pattern is too imprecise to carry weight on its own.

What remains uncertain. The study cannot show whether AI-assisted diligence is faster, more accurate or better at identifying material risks than manual diligence; whether the observed associations would persist in a larger and more varied sample; or whether the group differences are real. Statistical significance in a sample of 19 does not remove selection bias or measurement overlap, and a convergence of direction with industry surveys is contextual triangulation, not independent validation of the correlations obtained.

7. Proposed Framework and Implications

7.1 A five-stage framework for governed AI-assisted due diligence

Drawing on the survey, the literature and the secondary evidence, we propose a five-stage framework (Figure 5; Table 9). It is a proposed framework: it is derived from the study's findings and the literature, it has not been tested in the field, and the present survey does not validate it. Its purpose is to translate the study's central caution, that perceived capability should not be mistaken for measured performance, into a sequence of design and measurement steps that future studies can test.




Figure 5. Proposed framework for governed, human-verified AI-assisted due diligence (not empirically validated).

Stage 1 screens tasks for fit, favouring document-heavy, repeatable tasks with verifiable reference answers (for example clause extraction and cross-document comparison), which respondents rated as areas of high capability. Stage 2 tests data readiness, reflecting the finding that data quality was the strongest barrier and industry reports of data security and quality concerns. Stage 3 uses AI for a logged first pass. Stage 4 requires documented human verification by qualified professionals of any finding that affects value or deal terms. Stage 5 closes the loop with governance and measurement against a manual baseline, which is the step that converts perceptions into evidence. Table 9 sets out, for each stage, the purpose, activities, risks, safeguards, responsible personnel, indicators, evidence base and limitations.

Table 9. Framework stages (proposed): purpose, safeguards, indicators and evidence

Stage	Purpose and key activities	Risks and safeguards	Responsible; indicators	Evidence and limits
1. Task-fit screening	Choose tasks where AI plausibly fits: high volume, repeatable, verifiable. Classify tasks by fit.	Risk: applying AI to judgement-heavy tasks outside its capability. Safeguard: written task inclusion criteria; pilots before scale.	Deal lead and technology lead; share of tasks classified high-fit; pilot size.	Survey: highest agreement on volume, document review, anomaly flagging; TTF; task-dependent benefits (Dell'Acqua et al., 2023). Perceived, not measured, fit.

Stage	Purpose and key activities	Risks and safeguards	Responsible; indicators	Evidence and limits
2. Data readiness	Check completeness, formats, privacy, access control and vendor terms before use.	Risk: poor or confidential data. Safeguard: data checklist; confidentiality and vendor review; seller NDA terms.	Data owner, legal and security leads; data completeness score; checklist pass rate.	Survey: data quality strongest barrier; industry concerns on security and quality (Deloitte, 2025a). Thresholds not empirically established.
3. AI first pass	Run extraction, comparison and anomaly flagging with logged prompts, versions and outputs.	Risk: errors, omissions, unreliable output. Safeguard: version control; sampling of outputs; confidence flags.	Analyst team; time per document or workstream; extraction precision and recall.	Efficiency perceptions high; reported savings moderate (McKinsey & Company, 2026). No field benchmark.
4. Human verification	Qualified professionals verify findings that affect value or terms; document decisions and overrides.	Risk: over-reliance and automation bias; verification cost. Safeguard: defined review protocol; independent spot checks; training.	Senior reviewers and counsel; share of findings verified; false-positive rate; missed material findings; verification time.	Survey: strong endorsement of verification; literature on calibrated trust and liability. Survey does not show current practice.
5. Governance and measurement	Maintain an audit trail; compare against a manual baseline; review indicators and update stages 1–2.	Risk: untracked drift and liability. Safeguard: periodic review; escalation rules; alignment with AI risk-management guidance.	Risk committee and deal leadership; audit-trail completeness; performance relative to manual baseline; periodic review.	Governance context (NIST, 2023; Financial Stability Board, 2024; European Union, 2024). Indicators are suggestions, not validated measures.

Note. Indicators are suggestions for future field testing. Future studies could test the framework through matched manual versus AI-assisted workstreams on verified reference documents, measuring time, precision, recall, verification time and missed material findings, and comparing teams that do and do not apply the stages.

7.2 Practical implications

The strongest practical recommendation supported by the study is not to automate everything but to select appropriate tasks, establish data readiness, verify consequential findings and measure performance. For deal teams and investment banks, this suggests piloting AI on high-volume, document-centred tasks with indicators such as time, issues found against a manual baseline and false-positive and miss rates, because practitioner optimism is untested against outcomes. Data governance should precede scale-up, and verification should be a documented stage. Because younger respondents (exploratory) reported fewer barriers and less need for verification, firms may wish to build verification training and review protocols for junior staff. For financial analysts and advisers, the findings suggest recording AI use and verification in the work file. For technology vendors, the emphasis respondents place on data quality and verification points to explainability, audit trails and data-handling assurances as relevant features. For business educators, curricula could teach verification and governance of AI output alongside its use, since respondents consistently expected AI to complement rather than replace professionals. These implications are derived from perceptions and secondary evidence and should be treated as suggestions for testing.

7.3 Theoretical implications

The study illustrates a measurement problem for perception-based applications of TTF and technology-acceptance models: when capability and outcomes are both elicited as beliefs from the same respondents, they can collapse into a single construct, so a high correlation is weak evidence of a distinct fit-to-performance link. TTF's performance claim is about realised performance, so tests need objective performance data. Future work should measure capability through observable features (tools used, tasks automated, years of exposure), measure outcomes from separate sources, and include mixed-keyed items and attention checks to separate substantive agreement from response style. Experience with AI and age are candidate moderators but are not tested here.

7.4 Implications for regulators and professional bodies

The convergence in direction between practitioner concerns and the legal and governance literature suggests that guidance on supervision, explainability and record-keeping in AI-assisted transactions may be relevant to professional bodies. This is a suggestion drawn from convergent concerns, not a finding about regulatory effectiveness.

8. Limitations and Future Research

8.1 Limitations

The sample is small ($n = 19$), purposive and concentrated among investment banking and M&A professionals, so estimates are imprecise and cannot be generalised; selection towards favourable views is likely because respondents were heavy AI users and the topic is AI. The confidence intervals for the efficiency and risk correlations are wide, and exploratory group comparisons involve groups of 5 to 14 with 21 comparisons. All data are self-reported perceptions from a single instrument, so common-method variance and acquiescence are plausible; all items were positively worded, answers cluster at the ceiling, and nearly half of respondents (9 of 19) gave identical answers throughout, which compresses variance. Constructs may overlap, notably capability and accuracy, and capability and self-rated knowledge. The instrument was developed by the authors and was not formally validated; the risk-identification scale fell below the conventional reliability benchmark, although alpha estimates from three items and 19 cases are themselves uncertain. The design is cross-sectional, so association cannot be interpreted as causation. The secondary evidence comes largely from advisory firms with commercial interests, differs in definitions and samples, and is partly reported through trade media, so it supports convergence in direction only. Some secondary sources are working papers or conceptual contributions. The study measures perceptions of AI in general, does not distinguish tools, tasks or deal types, does not measure AI performance objectively, and did not record respondents' country or firm context. Recruitment relied on LinkedIn and the exact number invited was not tracked; consent was implied rather than written; the questionnaire was reviewed by a supervisor but not pilot-tested; and no institutional ethics approval or exemption is documented.

8.2 Future research

Future research should: (i) recruit a larger and more varied practitioner sample, with country, firm type and seniority recorded; (ii) improve the questionnaire by adding mixed-keyed items, attention checks, distinct operational definitions and separate measures of experience and expected benefit; (iii) use validated or formally developed scales and report expert review, pretesting and confirmatory factor analysis when the sample permits; (iv) replace belief-based predictors with measures of actual AI use; (v) compare manual and AI-assisted analysis in controlled studies, benchmarked against verified reference documents, for time, precision, recall, error rates and verification time, for example on public contract datasets; (vi) conduct field studies with live due-diligence teams to test the proposed framework; (vii) combine surveys with interviews or case studies to examine how verification and governance work in practice; and (viii) replicate across countries, firm types and transaction types.

As a precision benchmark, detecting a population correlation of .30 with 80% power at a two-sided 5% level requires roughly 85 respondents under the Fisher- z approximation, and about 113 if a Bonferroni-adjusted α of .0167 is applied across three outcome tests (assumptions: single predictor–outcome correlation, approximately bivariate normal composites, independent observations). A larger sample improves precision and power only; it does not resolve measurement validity, common-method variance or selection bias.

9. Conclusion

This study examined how finance and M&A professionals connect perceived AI capability with the outcomes that matter most in due diligence: efficiency, accuracy and risk identification. The findings point in an encouraging direction. Respondents were strongly positive about AI's usefulness in due diligence, with construct means between 4.33 and 4.63 on a five-point scale. Perceived AI capability was positively associated with all three perceived outcomes, and the association with accuracy stayed very strong even in the stricter sensitivity analysis. The associations with efficiency and risk identification were also positive in every analysis, although they were less stable once respondents who gave identical answers throughout were set aside. Respondents paired their optimism with a clear preference for human verification and named data quality as the main barrier. This matches the direction of published industry evidence, which suggests that practitioners see AI as a partner to professional judgement and not a replacement for it.

The study's main contribution is to link perceived AI capability to specific due-diligence outcomes and to test how stable those links are. It also offers a five-stage framework that covers task-fit screening, data readiness, an AI first pass, human verification, and governance and measurement. The framework is proposed, not yet validated, and gives practitioners a practical way to adopt AI responsibly. These results describe perceptions, not measured performance. The most valuable next step is therefore to test the framework in live due-diligence settings, comparing manual and AI-assisted work on time, accuracy and verification effort.

Declarations

Ethics and consent. Participation was voluntary and consent was implied by completing the anonymous online questionnaire. No institutional ethics approval or exemption is documented for this study.

Conflict of interest. The authors reported no conflict of interest.

Funding. No funding source was reported for this study.

Author contributions. Both authors are responsible for the manuscript. Ajay Kumar Varshney reviewed the questionnaire before distribution.

REFERENCES

- Alexis, A. (2026a, March 17). GenAI gains ground in M&A, but post-deal adoption lags. *CFO Dive*. <https://www.cfodive.com/news/generative-ai-gains-ground-mergers-postdeal-adoption-lags/814895/>
- Alexis, A. (2026b, February 18). Generative AI reduces M&A costs by 20%, McKinsey says. *CFO Dive*. <https://www.cfodive.com/news/generative-ai-reduces-merger-acquisition-costs-20percent-mckinsey/812514/>
- Angwin, D. (2001). Mergers and acquisitions across European borders: National perspectives on preacquisition due diligence and the use of professional advisers. *Journal of World Business*, 36(1), 32–57. [https://doi.org/10.1016/S1090-9516\(00\)00053-5](https://doi.org/10.1016/S1090-9516(00)00053-5)
- Bain & Company. (2024, April 29). *Generative AI in M&A: Where hope meets hype*. <https://www.bain.com/insights/generative-ai-m-and-a-report-2024/>
- Bain & Company. (2026, January 27). *Global M&A poised to sustain momentum in 2026 after great rebound, finds Bain & Company* [Press release]. <https://www.bain.com/about/media-center/press-releases/2026/global-ma-poised-to-sustain-momentum-in-2026-after-great-rebound-finds-bain--company/>
- Bedekar, M. A., Pareek, M., Choudhuri, S. S., Abhishek, P. F., Jhurani, J., & Shah, R. (2024). AI in mergers and acquisitions: Analyzing the effectiveness of artificial intelligence in due diligence. In *2024 International Conference on Knowledge Engineering and Communication Systems (ICKECS)*. IEEE. <https://doi.org/10.1109/ICKECS61492.2024.10616599>
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
- Dell’Acqua, F., McFowland, E., III, Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality* (Harvard Business School Working Paper No. 24-013). <https://doi.org/10.2139/ssrn.4573321>
- Deloitte. (2025a). *2025 M&A generative AI study*. Deloitte US. <https://www.deloitte.com/us/en/what-we-do/capabilities/mergers-acquisitions-restructuring/articles/m-and-a-generative-ai-study.html>
- Deloitte. (2025b, October 9). *New Deloitte survey: 86% of corporate and private equity leaders now use generative AI* [Press release]. <https://www.deloitte.com/us/en/about/press-room/deloitte-survey-genai-in-mna.html>
- Disu, A. O., Makinde, H., Ajide, O. A., Ojo, A., & Mbonu, M. (2022). Artificial intelligence in investment banking: Automating deal structuring, market intelligence, and client’s insights through machine learning. *Communication in Physical Sciences*, 8(4), 808–825.
- European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union. <http://data.europa.eu/eli/reg/2024/1689/oj>
- Fang, Y., Goyal, V. K., He, Y., & Zhang, Z. (2026). *When AI acquires data: Strategic complementarities in M&A* (HKU Jockey Club Enterprise Sustainability Global Research Institute Paper No. 2026/010). <https://doi.org/10.2139/ssrn.6100206>
- Financial Stability Board. (2024, November 14). *The financial stability implications of artificial intelligence*. <https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/>
- Gillespie, T. (2025, December 10). Banks mine decades of deal history to feed AI banking tools. *Bloomberg*. <https://news.bgov.com/artificial-intelligence/wall-street-builds-deal-datasets-for-ai-investment-banking-era>
- Goodhue, D. L., & Thompson, R. L. (1995). Task–technology fit and individual performance. *MIS Quarterly*, 19(2), 213–236. <https://doi.org/10.2307/249689>
- Haspeslagh, P. C., & Jemison, D. B. (1991). *Managing acquisitions: Creating value through corporate renewal*. Free Press.
- Herbosch, M., & Mertens, F. (2025). *The future of mergers & acquisitions? Risk allocation in AI-guided transactions* (Working Paper 2025-03). Financial Law Institute, Ghent University. <https://doi.org/10.2139/ssrn.5275610>
- Kothari, C. R. (2004). *Research methodology: Methods and techniques* (2nd rev. ed.). New Age International.
- Kotla, M. R. T. (2026). AI-driven data integration for mergers and acquisitions: Automating entity resolution and system consolidation. *International Journal of Innovative Research in Science, Engineering and Technology*, 15(1), 603–618. <https://doi.org/10.15680/IJRSET.2026.1501083>
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392
- Liu, D., Zou, Y., Wang, B., & He, X. (2026). Artificial intelligence as a merger and acquisition catalyst: A resource-based view. *International Review of Financial Analysis*, 114, Article 105163. <https://doi.org/10.1016/j.irfa.2026.105163>

- McKinsey & Company. (2026). *Gen AI in M&A: From theory to practice to high performance*. <https://www.mckinsey.com/capabilities/m-and-a/our-insights/gen-ai-in-m-and-a-from-theory-to-practice-to-high-performance>
- Mirzayev, E., Testoni, M., & Vanneste, B. S. (2025). *AI agents and the evaluation of M&As* [Working paper]. SSRN. <https://doi.org/10.2139/ssrn.5119314>
- National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
- Nunnally, J. C. (1978). *Psychometric theory* (2nd ed.). McGraw-Hill.
- Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors*, 52(3), 381–410. <https://doi.org/10.1177/0018720810376055>
- PE Professional. (2025, October 15). *Deloitte: PE loves AI*. <https://peprofessional.com/2025/10/deloitte-pe-loves-ai/>
- Podsakoff, P. M., MacKenzie, S. B., Lee, J.-Y., & Podsakoff, N. P. (2003). Common method biases in behavioral research: A critical review of the literature and recommended remedies. *Journal of Applied Psychology*, 88(5), 879–903. <https://doi.org/10.1037/0021-9010.88.5.879>
- Rashid, M. M., Ullah, N., Uddin, M., & Rahman, M. A. (2025). Artificial intelligence on merger and acquisition processes: Observation from the target identification and due diligence perspective. *International Journal of Innovative Research in Multidisciplinary Education*, 4(1), 21–27. <https://doi.org/10.58806/ijirme.2025.v4i1n03>
- Stewart, J. (2024). AI and investment banking. *American Banker Magazine*, 134(1), 36.
- Tavakol, M., & Dennick, R. (2011). Making sense of Cronbach's alpha. *International Journal of Medical Education*, 2, 53–55. <https://doi.org/10.5116/ijme.4dfb.8dfd>
- Thelisson, A.-S. (2024). How can artificial intelligence impact the evaluation phase of a company takeover? Contribution of an in-depth case. *Strategy & Leadership*, 52(5/6), 57–63. <https://doi.org/10.1108/SL-05-2024-0047>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>