The Impact of Artificial Intelligence on Venture Capital: A Critical Outlook

Thomas Hellmann, Robin Gansäuer, Junida Mulla, and Matthias Qian [1] [2]

January 2026

Submitted to "Oxford Intersections: AI in Society"

Abstract

This paper examines how the adoption of artificial intelligence (AI) technologies is transforming the venture capital industry. Leveraging interviews with industry practitioners who are actively working on the adoption of AI tools, we develop a critical outlook on the likely changes awaiting the industry. We find that AI tools accelerate the sourcing and due diligence of venture deals. However, the final authority to make investment decisions remains with humans. We identify socially grounded conviction and gut feelings as two key human proficiencies that AI systems struggle to replicate. Relationships embedded in their professional networks also allow venture capitalists to provide value-added support to their founders in ways that cannot be easily replaced by AI systems. As AI systems are adopted industry-wide, they will broaden founders' access to funding and shift market power away from investors. It is not the AI infrastructure itself, but human proficiencies, such as socially grounded conviction, gut feeling, and networks, that will allow venture capital firms to competitively differentiate themselves.

Keywords: Artificial Intelligence, Venture Capital, Investment Process, Conviction, Gut Feeling, Network Relationships.

[1] All authors are associated with Saïd Business School, University of Oxford. The corresponding author is Thomas Hellmann (thomas.hellmann@sbs.ox.ac.uk). We would like to thank Professor Mari Sako for her many helpful comments, the two referees for their valuable suggestions, and all the anonymous interviewees for their time and invaluable insights.

[2] During the preparation of this work the authors used OpenAI's ChatGPT tool to edit the manuscript for grammar and readability. After using this tool, the authors directly reviewed and edited all the content. This use of AI complies with all journal requirements. The authors take full responsibility for the content of the published article.

Introduction

The recent rapid advances in artificial intelligence (AI) have the potential to transform entire industries (Agrawal et al., 2023). However, the nature and extent of this remain uncertain at present. In this paper, we describe the recent adoption of AI technologies in the VC industry and explore its implications for how investments are made and how venture capital (VC) firms structure and differentiate themselves.

A reasonable starting premise for assessing the impact of AI is that investing in startups is a deeply human activity, so the industry should eschew any aggressive AI automation. Early-stage ventures are largely ideas in the minds of founders and require entrepreneurial talent and drive to come to fruition (Kaplan et al., 2009; Gans et al., 2019). Investors thus need to build conviction around founding teams, their evaluation being largely guided by gut feeling (Huang & Pearce, 2015). Moreover, venture capitalists (VCs) constantly leverage their personal networks, both for themselves and their portfolio companies (Gorman & Sahlman, 1989; Sahlman, 1990; Hochberg et al., 2007). There is a high level of 'Knightian' uncertainty (Townsend et al., 2025), ambiguity abounds, and there is a paucity of objective data. Hence the conjecture that AI will only have a limited impact on the VC industry.

Yet there is a marked rise in the adoption of AI tools by VC firms (Sako & Qian, 2021; Retterrath, 2025). Recent breakthroughs in large language models (LLMs) have proven particularly effective in environments rich in unstructured textual data, including VC. Moreover, VC is subject to relatively little regulatory oversight (Lerner and Tag, 2013), so unlike other parts of the financial sector, VC firms are relatively unconstrained in adopting AI tools. Due to their own investments in AI companies, VCs also gain early exposure to these rapidly advancing technologies.

Our analysis is based on 30 semi-structured qualitative interviews with VC industry practitioners who are actively engaged in, or directly exposed to, AI adoption in VC firms. Our empirical sample is drawn from VC firms based in Europe and North America, with limited representation from the Middle East and Southeast Asia. It includes participants involved in strategic leadership, technical implementation, and the practical use of AI. Based on these targeted interviews and within the exploratory nature of this paper, we identify two broad areas where AI is beginning to affect the structure of the VC industry. The first concerns the investment process and how AI affects its different stages. The second concerns competition among VC firms, including their internal organisation, and, most importantly, the nature of their competitive advantages. Throughout the analysis, we ask to what extent AI tools augment or replace human actors, and what the role of human agency is (Raisch and Krakowski, 2021; Felin et al., 2025).

The first broad area concerns the investment processes of VC firms, i.e., how deals progress through decision stages from sourcing, to screening, to due diligence, to investment decisions, and beyond. We find that AI tools mostly support the stages before investment decisions are made. LLMs are particularly useful to process heterogeneous materials such as pitch decks, internal notes, call transcripts, data-room documents, and financial models, converting them into searchable and structured knowledge bases. However, final investment decisions continue to be the sole responsibility of human actors. This partly reflects institutional arrangements where limited partners require human sign-off for capital allocations. More profoundly, there is a strong belief in the industry that investment decisions should be made on the basis of human conviction (Felin et al., 2025). While assisted by AI tools, this conviction remains grounded in VCs’ experience, the insights they learn through their professional networks, and ultimately their own gut feelings.

The importance of human conviction is particularly salient when evaluating founder teams. Huang and Pearce (2015) argue that early-stage investment decisions involve both analytic cognition, which is slow and based on data evaluation, and intuition, which is fast and based on more emotive judgment from direct founder interactions. This results in a holistic judgment that is experienced as a gut feeling. Our interviews suggest that AI is not used to replace this gut feeling, but merely to better organise the information that supports the cognitive part of this human decision process. Specifically, AI helps to standardise founder-related data, rendering previously heterogeneous founder materials into structured profiles, and possibly ranked scores. The availability of such systematic data matters in light of the pervasive evidence that VC decisions exhibit considerable biases, most notably in relation to gender and ethnicity (Coleman & Robb, 2009; Hegde & Tumlinson, 2014; Kanze et al., 2018; Ewens & Townsend, 2020; Hebert, 2025; Hu & Ma, 2025; Hellmann et al. 2025). In principle, AI tools hold the promise of monitoring, addressing, and eventually eliminating these biases, but the verdict is still out on whether and to what extent this is actually happening.

The second broad area concerns the internal structure and competitive differentiation of VC firms. The adoption of AI is not only changing the internal workflows of VC firms but also their workforce needs. Junior positions are especially vulnerable to AI, thereby eliminating a traditional entry point into the industry. Extreme cases of this are recent occurrences of VC firms run by a solo general partner without any staff. Moreover, data literacy and AI technical knowledge are becoming more important and some VCs have hybrid skills in investing and AI. Yet, how fast VC firms are changing, and how far they are willing to go, varies considerably across firms, and is largely driven by the willingness of senior partners to risk organisational changes. There is also a marked difference between traditional VC firms that are integrating AI into their existing organisational structures, versus new AI-native VC firms that place AI capabilities at the heart of their organisation.

In the short term, AI infrastructure affects the way VC firms compete with each other. Our interviews uncover an increased emphasis on execution speed, with slower VCs missing out on fast-moving deals. Over the longer term, and in the absence of further fundamental shifts in AI technology, there is likely to be a convergence of AI capabilities, especially since AI tools rely mostly on widely available information sources. Such convergence, however, does not necessarily imply convergence of investment strategies. The industry is likely to stratify into several types of VC firms that include large generalist VC firms that operate their own AI infrastructure, smaller specialist VC firms that rely on an outsourced AI infrastructure and differentiate in other dimensions, as well as hybrids and further variants.

One important implication of AI convergence and the dramatic increase in transparency it affords to deal sourcing, is that market power is likely to tilt away from investors and toward founders. This will heighten the need for VC firms to develop more distinct capabilities. Brand and reputation, the traditional differentiators in the VC industry, will continue to matter but will have to keep pace with the changing industry dynamics. The big open question is what other types of differentiation will matter most in a fully AI-enabled VC industry.

Combining the information from our qualitative interviews with insights from the prior academic VC literature, we articulate our own outlook on an AI-enabled VC industry, focusing on how human actors will continue to play a distinct role from ever more powerful AI agents. Indeed, we submit that future competitive advantages are primarily based on human capabilities that enable VC firms to stand out from an otherwise undifferentiated crowd of generic AI-enabled investors. We emphasise three types of human capabilities that cannot be easily replicated with current AI technologies. The first one is 'socially grounded conviction.' This relates to strongly held beliefs in ambiguous environments with limited information, as well as the willingness to make decisions and bear the reputational consequences. It is socially grounded because conviction is not developed in isolation but results from an embeddedness in a diverse set of relevant expert networks. The second human capability concerns the appropriate use of 'gut feelings,' which we argued before involves a combination of human triangulation of cognition and emotions. By virtue of being a feeling, it is not available to (non-sentient) AI agents. Thirdly, we foresee considerable room for differentiation in the post-investment phase, with founders preferring investors who can provide stronger value-adding support, especially in terms of access to relevant industry networks (Sahlman, 1990; Hellmann & Puri, 2002; Hsu, 2004; Hochberg et al., 2007, 2010). These activities remain deeply rooted in the human relationships that are embedded in professional and personal networks.

Ultimately, our outlook identifies a profound irony, namely that as AI systems become more ubiquitous in the VC industry, human capabilities, such as socially grounded conviction, gut feelings, and network relationships, are poised to become VC firms' most powerful competitive advantages.

Definitions and Methodology

Venture capital is an institutional form of equity finance in which professionally managed funds make staged investments in early-stage ventures, with investment decisions, governance, and value-adding activities delegated to general partners acting on behalf of limited partners (Sahlman, 1990; Da Rin & Hellmann, 2020). The industry matters as it is a key driver of innovation and economic growth (Gompers & Lerner, 2001; Samila &

Sorenson, 2011).

Artificial intelligence denotes a heterogeneous family of computational techniques with distinct capabilities and limitations (Russell & Norvig, 2021). We use AI pragmatically to refer to data-driven, task-specific computational systems, typically grounded in machine learning, that support judgment-intensive work by performing tasks such as classification, prediction, and synthesis. The field encompasses multiple learning approaches, such as supervised and unsupervised learning (Bishop, 2006), as well as a variety of model classes, including deep neural networks (Goodfellow et al., 2016). Recent scale-driven progress has produced "foundation models," pretrained, often via self-supervised learning (Bommasani et al., 2022), including large language models (LLMs), which are commonly built on Transformer architectures (Vaswani et al., 2017). Reflecting the text-heavy nature of VC workflows, our interviews highlight LLM-enabled tools (e.g., OpenAI's ChatGPT and Google's Gemini). Unless stated otherwise, we use the single label AI throughout to cover this broader family of techniques.

This paper adopts a theoretically informed, inductive qualitative research design to elucidate the mechanisms through which AI changes the way VC firms invest and compete with each other. As our research question concerns changes in investment workflows and decision authority, semi-structured interviews are well suited to capture how these shifts are understood and enacted in practice.

We focus on broad patterns across VC firms and do not systematically differentiate by firm size, investment stage, or industry focus unless explicitly stated. Consistent with our definition of VC, we exclude alternative startup investors, such as angel investors or corporate venture investors. Our empirical material is drawn primarily from VC firms based in Europe and North America, with limited representation from the Middle East and

Southeast Asia. We focus on the use of AI by VC firms rather than their investment in companies developing new AI systems.[3] We also do not examine the use of AI by limited partners, which is addressed in Phalippou (2025).

To develop the semi-structured interview guide , we reviewed firm websites, industry newsletters, and relevant academic and practitioner publications (Flick, 2014).[4] These documentary sources also served to contextualize and complement interview accounts. Between May and December 2025, we conducted 30 semi-structured interviews with practitioners at VC firms making substantial use of AI. Interviewees spanned founding, general, managing, and senior partners, as well as senior data and AI personnel (e.g., heads of data, directors of data science, and machine-learning or analytics engineers), encompassing both investment decision-makers and technical implementers. Participants were identified through publicly available industry lists, supplemented by targeted outreach and snowball referrals. Consistent with the study's focus on technological frontier practices, the evidence describes early adoption patterns that are not necessarily representative industry-wide. All interviews were conducted via video call, lasted 45–60 minutes, and were recorded and transcribed verbatim with participants' consent.[5]

We analysed the interview material using reflexive thematic analysis following Braun & Clarke (2022), a systematic and iterative approach that involves coding the data, developing and refining themes, and synthesizing them into an analytic account.[6]

[3] For that, see the O3 tools that map the AI startup ecosystem in the UK: www.o3.ventures (accessed 14 July, 2025).

[4] The interview guide is available upon request from the lead author.

[5] To preserve confidentiality, we do not name individuals or firms. Instead, all interviewees are anonymised using numeric identifiers (I01–I30). The table in the Appendix reports non-identifying descriptors such as role, stage focus, fund size, and region to give readers a sense of the diversity of roles and contexts while maintaining privacy.

[6] Braun and Clarke set out thematic analysis as a six-phase process: (1) familiarizing yourself with your data, (2) generating initial codes, (3) searching for themes, (4) reviewing themes, (5) defining and naming themes, and (6) producing the report. In practice, we approached these phases iteratively, revisiting the dataset, coded extracts, and themes as needed.

While the study was theoretically informed, coding was conducted inductively to construct themes capturing recurring mechanisms related to AI adoption in VC firms. Although the process-focused questions in our interview guide were organised around stages of the VC investment funnel, coding was not ultimately constrained to these stages. The resulting patterns provide the interpretive foundation for the core themes developed in this paper.

The Role of AI in the Venture Investment Process

We begin by asking how VC firms adopt AI in their pre-investment activities. For this we begin by outlining a canonical structure of the VC investment process before the integration of AI-based tools. Early work by Tyebjee & Bruno (1984) provided an early systematic framework for understanding how investment opportunities move through stages of deal origination (aka. sourcing), screening, evaluation (including due diligence), the investment decision and deal structuring, and post-investment involvement. Da Rin & Hellmann (2020) build on this framework by modelling the process from both founder and investor perspectives, while Gompers et al. (2020) document its empirical regularities using large-scale survey evidence.[7]

We take the perspective of VC firms and conceptualize investment activity as a four-step process. First, there is deal sourcing and screening, where VCs identify potential investment candidates and perform an initial triage. Second, there is the due diligence stage, where select candidates are evaluated in greater depth. This includes the important challenge of evaluating the quality of founder teams. These first two stages constitute the pre-investment activities of the process. The third stage concerns the investment decision, i.e., the

[7] Further recent papers include Jang & Kaplan (2025), who examine the investment funnel of a single VC firm, and Fu & Taylor (2025), who use cell phone data to assess how the duration of founder-investor interactions affects deal selection.

formal approval by an investment committee and the ensuing deal structuring and negotiation. The fourth stage concerns post-investment activities, all the way through to exit.

Deal Sourcing and Screening

We start with deal sourcing and screening, showing how AI expands market visibility and structures large inflows into manageable decision sets. Traditionally, VCs have sourced deals through networks, combining passive referrals, opportunistic search, and deliberate scouting activities. AI tools help to significantly broaden and formalise how VC firms identify potential investments, by standardising and automating the aggregation of signals across multiple sources. Interviewees explained that they are building a technology stack for sourcing deals that combines commercial vendors, such as Harmonic, Specter, PitchBook, Crunchbase, Dealroom, and similar platforms, with public sources, including company registries, and proprietary web scraping (I09; I13; I18). These inputs are integrated into internal data warehouses and CRMs. Together, these systems considerably expand VC firms' coverage across geographies, stages, and sectors (I05; I19). Proprietary models continuously update the set of companies under consideration (I10).

Interviewees identified three operational archetypes for AI-based deal sourcing among VC firms. First, some VC firms build AI-native, fully proprietary sourcing systems that combine ingestion from multiple data sources with natural language processing and predictive models. These in-house data infrastructures trigger automated alerts when they find new signals about venture traction or hiring (I05; I11; I19; I25). Second, many smaller and mid-sized VC firms rely primarily on vendor platforms that aggregate and standardise startup data. They can configure their own filters by geography, sector, and stage but do not develop proprietary pipelines (I17; I20; I24). Third, specialist automation studios design and operate AI platforms on behalf of VC firms (I21; I28).

Across all three models, VCs monitor combinations of signals to detect early signs of venture momentum. Signals include legal and organisational events (new incorporations, changes in shareholder registers), labour-market indicators (hiring spikes, job postings), digital traction (web traffic, developer activity on platforms like GitHub), community engagement (e.g., Discord), and product launches on platforms such as Product Hunt (I07; I19; I23). Some VCs also mine sector databases, university technology transfer channels, or scientific publication platforms to identify nascent technical projects (I09; I18). Interviewees report that these systems can materially shape pipelines. For example, one VC firm's system processes roughly one thousand startup leads per week, narrowing them to a few dozen qualified opportunities (I07). In this sense, AI changes which companies become visible to VCs and expands the effective scope of the market they can scan.

As sourcing volumes increase, deal screening becomes the primary mechanism for translating expanded visibility into manageable decision sets. AI-based screening systems structure high-volume inflows into a smaller set of companies that warrant deeper human attention. They standardise heterogeneous inputs from pitch decks, websites, founder profiles, application forms, and internal materials into comparable fields within CRMs and screening dashboards, enabling VCs to apply consistent evaluation criteria across expanded pipelines (I01; I06; I14; I16). On this basis, most data-driven VC firms deploy scoring and ranking models with explicit thresholds that determine which companies advance to human review. Some scores emphasise business fundamentals, some focus on venture traction, and others on founders (I07; I10; I13; I17; I23). Interviewees stress that these systems reflect deliberate design choices rather than neutral measurement. Some VC firms accept higher false-positive rates to avoid filtering out atypical founders, whereas others rely more heavily on structured performance signals, especially at later stages where richer data are available (I15; I18).

Those leads that meet defined thresholds are promoted into the CRM for human review (I07; I22), allowing some VCs to compress tens of thousands of potential companies into a shortlist of a few hundred thesis-aligned deals (I08; I25). One VC firm estimates that only the top five to seven percent of teams receive direct partner attention (I25). As a result, AI systems have effectively become a gatekeeping infrastructure between sourced deals and partner attention (I13). In practice, AI-enabled screening operates as a scalable attention infrastructure: it converts diverse, early-stage signals into routable categories and ranked queues inside the CRM, making high-volume inflows tractable and turning expanded market visibility into an operational pipeline. Similar capabilities of LLM-based evaluation have been documented in startup-competition and lab-style experimental settings (Csaszar et al., 2024; Doshi et al. 2025). Human oversight remains important in configuring and auditing this gatekeeping layer, particularly at very early stages where signals are sparse and automated filters may be overly conservative or systematically exclude non-standard profiles (I08; I18; I25).

Venture Due Diligence

For the due diligence stage, we divide our discussion into two parts. In this subsection, we discuss due diligence of the venture, and in the next subsection, we consider due diligence of the founder team. Once an opportunity advances to the due diligence stage, AI is primarily used to systematise and extend the VCs' human analytical work. Interviewees describe workflows that ingest heterogeneous materials such as pitch decks, internal notes, call transcripts, data-room documents, and financial models and convert them into searchable and structured knowledge bases (I01; I03; I06; I09). Transcription, indexing, and vector-based retrieval allow prior analyses, memos, and call notes to be reused as deals progress or re-emerge, creating a cumulative diligence memory that extends beyond individual transactions (I07; I11; I25; I29; I30).

Where companies generate meaningful operational data, AI also supports deeper quantitative analysis. Interviewees describe AI tools that transform transaction files or revenue sheets into structured metrics such as retention curves, cohort analyses, and benchmark comparisons against peer companies (I10; I19; I22). In more quantitatively oriented VC firms, diligence focuses on verification and plausibility checks, where investors cross-check information from public and private sources and conduct fraud-risk assessments. Across all these applications, AI surfaces patterns and anomalies, while interpretation remains with human investors. Sometimes investors also deliberately exclude these AI-based tools in sectors where data-driven evaluation is considered unreliable (I13).

Beyond commercial analysis, AI is also applied to due diligence on legal, risk, and governance matters. Several interviewees report using tools that parse contracts, extract clauses, and classify items into structured risk categories to support review by counsel and partners (I16; I23). Some VCs experiment with early-stage governance tools such as bias checks on interview samples, but interviewees stress that these applications remain immature and are treated as supplementary rather than decisive (I05; I10).

An important part of due diligence is seeking second opinions from other investors and experts. For this purpose, VCs can leverage AI-based relationship graphs that integrate contact histories (e.g., WhatsApp), co-investments, and second-degree links to identify potential introduction paths to other relevant investors (I04; I30).

Generative systems also support the synthesis of diligence outputs. Models are used to draft sections of investment memos or populate standardised templates with analytical outputs. Again, investors remain responsible for interpretation, narrative coherence, and final judgment (I01; I08; I14; I18; I29). Interestingly, several interviewees emphasise that memo writing remains an important learning mechanism, particularly for junior investors. AI-generated drafts are treated as starting points rather than substitutes for humans taking full responsibility for their investment memos and standing behind their content (I08; I27).

Interviewees also noted that some founders anticipate how they are evaluated by AI systems. They fine-tune websites and profiles to match the keywords and categories they believe AI-based systems prioritise, a dynamic likened to answer-engine optimisation for venture pipelines (I10; I18). Some VCs view this as resourcefulness and rely on later human evaluation to separate substance from presentation, while others add verification steps to prevent early algorithmic uplifts (I10; I12; I19; I23).

Founder Due Diligence

Founder evaluation plays a disproportionate role in early-stage investment decisions in VC (Bernstein et al., 2017), in which individual partners and investment committees assess founders' capabilities, integrity, and perceived fit under conditions of high uncertainty. Arthur Rock, widely considered the father of the US VC industry, famously argued that only three things mattered: "people, people, and people" (Bell, 1991).

Founder evaluation in early-stage VC relies heavily on experience-based judgment rather than purely analytical evaluation. The work of Huang & Pearce (2015) explains how, in complex environments with high-dimensional uncertainty, early-stage investors often rely on gut feelings to make their investment decisions. They provide a framework for conceptualising gut feeling as the output of a deliberation process that combines both cognitive and intuitive inputs. This builds on the seminal work of Kahneman (2011), which distinguishes between Type 1 'fast' thinking, driven by intuition, instinct, and affect, and Type 2 'slow' thinking, an intentional cognitive process that evaluates available evidence rationally and systematically. Huang (2019, 2025) further argues that, for experienced investors, this deliberation process generates superior results than a purely analytical decision process relying only on rational cognition.

Across our interviews, investors stressed that founder evaluation remains central to VC decision-making. This is particularly true at the pre-seed and seed stage, where companies consist of little more than a founder team with an idea (I01; I20; I21). AI does not displace the emphasis on founder teams, but changes the available information. At early stages, AI tools stitch together resumes, LinkedIn profiles, websites, and founder materials into a more coherent picture. Later, as founders accumulate richer investor interactions, systems can also draw on transcripts, reference notes, prior board materials, and signals embedded in the behaviour of other investors (I03; I06; I12; I23; I26). In this sense, AI expands and organises founder information rather than substituting for judgment.

Many VCs convert heterogeneous founder information into structured profiles, tags, and scores, presented as founder cards or dashboard views for internal discussion (I03; I14; I25). Some teams use founder-centred graphs that link individuals to past companies, funding events, and relevant news, to flag patterns associated with reputational concerns or fraud risk (I05). Others rely on point-based scoring frameworks that weight signals such as seniority, prior exits, accelerator participation, and backing by highly regarded investors (I07; I23; I26).

Interviewees emphasised that building these systems involves making critical choices. Several VCs place less emphasis on educational credentials and greater emphasis on professional experience, role progression, and evidence of execution abilities (I11; I18; I23). Others remove founder-centric metrics altogether. One quantitatively oriented vehicle relies on the historical efficiency of earlier stage investors as a proxy for founder quality, viewing direct founder features as too noisy (I02). Even before any direct interaction has occurred, different VCs can thus form meaningfully different views of the same founder.

A small but notable group of VC firms in our sample experimented with using AI to derive behavioural and psychological founder signals from language and interaction data. Interviewees describe tools that analyse meeting transcripts, emails, and internal notes to examine how founders communicate, explain trade-offs, and respond to pushback. They then translate these observations into structured cues that can inform further discussion (I03; I07; I30). Some investors use text-based models to extract traits such as persistence, goal orientation, or product discovery skill (I07; I25). Others build founder graphs linking individuals to prior companies, funding histories, and media signals to flag reputational concerns or risk patterns (I05). These efforts are described as more viable at slightly later stages, when repeated interactions generate enough textual and relational data to analyse (I06; I12).

Interviewees stressed that these signals were still fragile. Several respondents reject computer vision or voice analytics tools that claim to infer affect, facial expression, or vocal tone, citing weak correlation with entrepreneurial quality, risk of penalising non-expressive founders, and ethical discomfort with surveillance-style evaluation (I06; I13; I22; I29). Interviewees also report that they tried more ambitious psychological approaches that were ultimately rejected. Experiments with personality frameworks and psychometric assessments, such as the Myers–Briggs Type Indicator, as well as simple trait scores derived from pitch transcripts, were described as too noisy to be of use (I22; I26). In this sense, the frontier is defined as much by what VCs refuse to adopt as by what they retain.

Across our interviews, investors emphasise that direct interaction with founders remains indispensable, particularly at the earliest stages (I01; I08; I09; I17; I21). Judgments hinge on how founders show up in conversation, including authenticity, curiosity, and how they handle uncomfortable, probing questions. Interviewees describe gut feeling less as emotion than as integrated pattern recognition drawing on experience, contextual cues, and relational information that rarely appears in formal datasets (I09; I11; I20; I22).

AI can also be leveraged to provide feedback about human conviction. Interviewees describe using analytical tools, benchmarks, and red flag summaries to challenge overly optimistic first impressions, identify risks intuition might underweight, and force partners to articulate why they still back a founder despite adverse signals (I01; I11; I12; I20; I23; I25; I26; I30). Interviewees also describe how AI can assess proposed decisions against past decisions and realised outcomes, treating systematic disagreement as a cue either to refine the model or to conclude that their judgment relies on information the model cannot observe (I10; I11; I12; I13; I26). Interviewees also draw clear boundaries around what AI can contribute in founder assessment. Several argue that models cannot reliably interpret humour, cultural nuance, or situational behaviour in high-stakes meetings. Nor can they observe how investors and founders build trust over time (I08; I11; I19).

Investment Decision and Deal Structuring

In VC, investment committees hold formal decision authority over all investment decisions. There, the role of AI remains severely constrained. In most VC firms we interviewed, AI is confined to a supporting role within investment committees, mostly in terms of organising information and sharpening arguments. However, all investment decision rights remain firmly in the hands of human partners. Committee members are expected to understand the firm’s AI tools and interpret their outputs critically. They treat all AI outputs as inputs to human deliberation rather than determinants of outcomes (I01; I14; I23). Some firms deploy AI-based critique agents as a devil's advocate, reviewing memos before they reach the committee, and flagging weak arguments or missing evidence (I04; I28).

For the vast majority of firms we interviewed, investment committees are intentionally reserved for humans only. Interviewees emphasise that limited partners back general partners, not models. They expect traceable reasoning and accountability in venture investing (I04; I08; I09; I18; I19). Several VCs mentioned that they had explored using AI for indicative votes, but encountered interpretability challenges. The concern was that the AI systems were too agreeable and/or fragile to provide independent opinions (I05; I09; I15; I27). Interviewees also noted that some regulated settings require documented human sign-off for investment decisions (I19; I24). Short of granting AI models a vote, some VCs use AI to compare model recommendations against past committee outcomes and track partner rationales and confidence over time to improve process discipline and calibration (I10; I20; I21; I26). Across these practices, AI remains limited to supporting preparation and consistency, while the act of deciding and owning consequences remains with humans.

Notably, a very small number of frontier VCs depart from this pattern. One quantitative VC firm defines the committee as two general partners plus the AI model, tabling a deal only when the model signals positively (I13). Another specialised VC firm operates fully algorithmically, triggering investments only when internal signals from multiple AI models align. This VC firm has no conventional investment committee at all. Importantly, it never leads any deals but only invests as a non-lead investor on predefined terms (I02). These two VC firms are clear outliers that mark the outer boundary of AI-automation we encountered in our interviews.

Once the investment committee approves an investment, the process moves on to deal structuring and negotiation. Interviewees describe internal systems that store and index historical term sheets, cap tables, and deal outcomes, which partners use to benchmark valuations and terms. These systems function primarily as data infrastructure rather than automated negotiators (I10; I21; I23). In some firms, historical portfolio analysis informs target check-size ranges and allocation rules that guide, rather than determine, capital deployment across approved deals (I13).

AI is also used to review investment documents. Interviewees describe tools that extract clauses, summarise term sheets and investment agreements, and flag inconsistencies or missing items for lawyers and partners to examine (I14). Legal staff still manage term sheets, contracts, and final signatures, while partners retain responsibility for valuation, control rights, and governance choices (I13; I14).

Biases in Investment Decisions

Prior academic research has identified systematic biases in VC investment decisions.[8] This issue has become an important consideration in the development of AI-based evaluation models. Interviewees identified three layers of bias in founder evaluation: conscious preferences, unconscious stereotypes, and new algorithmic biases introduced by AI (I01; I05; I06; I11; I20; I23). Models trained on historical venture data may favour founders who resemble those previously backed (I06; I11; I12; I26). This affects not only how founders are scored but who is scored in the first place (I10; I14; I25). Several interviewees warn that

[8] Much of this literature focuses on the large gender gap in venture capital, both on the investor and founder side. Coleman & Robb (2009) document the gender gaps in a broad population of startup companies. Kanze et al. (2018) show how VCs pose different questions to male vs female founders. Hebert (2025) shows how gender stereotypes vary across various industry segments. Hu and Ma (2025) provide evidence based on video analysis that shows that female founders are judged more harshly than their male counterparts. Ewens & Townsend (2020) and Hellmann et al. (2025) show that gender-assortative matching between founders and investors accounts for a significant portion of the gender gaps. This result is primarily driven by male investors' preference for male founders. Beyond gender, there are other important biases, most notably those related to ethnicity. Hegde & Tumlinson (2014) document assortative matching by ethnicity. To what extent these biases are conscious or unconscious remains a matter of ongoing debate.

founder filters, if left unchecked, risk hardwiring past preferences into future portfolios and systematically excluding atypical or less well-networked founders (I11; I18; I22; I26).

Some VC firms explicitly design systems to mitigate bias rather than amplify it. Interviewees describe reweighting founder scoring so that prestige variables such as elite education or brand-name employers do not dominate (I12; I23). Some firms anonymise or suppress selected background variables, such as gender, while others avoid founder models altogether to reduce bias. They may exclude founder metrics and rely on alternative proxies. They may also reject tools such as computer vision and covert scraping that could encode sensitive attributes in opaque ways (I02; I13; I29). In these cases, bias mitigation is pursued not through debiasing but through restricting what information the models can observe. However, as suggested by the work of Fuster et al. (2022) and Howell et al. (2024), one should also note that debiasing capital allocation remains a significant challenge that cannot be fixed by technology alone.

Post-Investment Activities and Exit

After the initial investment, many VCs use AI to systematise portfolio monitoring, reporting, and portfolio support. Interviewees describe data pipelines and dashboards that aggregate financial and operational inputs from board documents, founder updates, investor reports, and external news (I01; I19; I27). Models parse emails and slide decks, extract key metrics such as revenue growth, hiring, and digital traction, tracking them over time (I17; I20). AI models clean and normalise incoming reports, flag missing submissions, and trigger reminders. They detect anomalies and provide benchmarks against comparable companies (I16; I22; I25). Several VCs also use generative tools to initially draft limited partner updates based on structured portfolio data, though humans always refine these drafts before releasing them (I04; I12; I16).

AI agents also support post-investment value creation by helping VCs match portfolio needs to external resources. Platform and talent teams use search and recommendation tools to identify candidates for key roles and return ranked shortlists (I06; I21; I23). Network-mapping tools can surface prospective customers, technical experts, or co-investors whose profiles, and prior activity align with a company's stage and sector (I07; I24; I29). In parallel, some VCs gather qualitative material, such as board documents, founder emails, customer-call transcripts, and internal notes, to compare best practices and recurring risks across companies, and, where appropriate, share them with founders (I18; I26; I30).

Post-investment data increasingly feeds back into the top-of-the-funnel decision engines discussed above. Several VC firms maintain structured datasets on portfolio performance and founder engagement and use these to refine quantitative models and qualitative heuristics for evaluating new deals (I11; I26; I27). Portfolio metrics and customer-call analyses can also function as live comparables when evaluating new opportunities or reassessing existing ones. Cohort views and index-like constructs help benchmark companies against peer groups and inform discussions when performance diverges (I01; I02; I19; I27). In this way, AI-enabled portfolio tooling acts as a monitoring layer and a learning mechanism that informs sourcing priorities, screening thresholds, and diligence questions in subsequent cycles (I18; I23; I29).

Compared with earlier stages of the investment process, the role of AI in supporting exits remains nascent in our sample. Many data-driven VCs have not yet reached a lifecycle stage where many exits occur. Several interviewees said they had no AI-based exit processes to date (I13; I19; I24). Where AI tools are used, they primarily support preparatory analysis, but do not extend to exit strategy or decision-making. Interviewees describe experiments that identify potential acquirers or follow-on investors by analysing historical acquisition and coinvestment patterns (I10; I17; I22; I23). These efforts are often framed as exit analytics projects, but remain experimental rather than embedded in day-to-day practice (I13; I19).

AI and the Competitiveness of VC Firms

In this section, we examine how AI is changing the internal organisation of VC firms and how it affects their competitive strategies.

Internal organisation of VC firms

VC firms are characterised by lean organisational structures in which a small number of partners exercise disproportionate influence over investment decisions. Prior research emphasises that competitive advantage in VC derives primarily from partner-level human capital rather than firm-level organisational routines (Sahlman, 1990; Ewens & Rhodes-Kropf, 2015). AI adoption affects the internal organisation of VC firms, modifying their workflows and changing their workforce needs. AI platforms absorb many repetitive tasks, resulting in less administrative and information-gathering work. They even assign leads and recommend who should join meetings (I04; I22; I24; I30). However, the extent to which workflows and workforces change varies considerably, with a clear divide between new AI-native VC firms, versus traditional VC firms that are in the process of bringing AI into their existing organisations.

The new AI-native VC firms like to describe themselves as “technology companies whose product happens to be VC” (I03). Consequently, they also have a distinct internal structure with different talent needs. Specifically, they rely on in-house engineering, data, and AI teams that build proprietary platforms mediating workflow across the firm (I03; I05; I10; I21; I30). Technical roles are now at the core of the investment organisation. They employ in-house data engineers, data scientists, and machine learning specialists to build internal platforms that support the full investment lifecycle, either in dedicated teams or embedded in partner workflows (I03; I10; I11; I13; I15; I19; I22). Interestingly, their workflows are not always built only by central engineering teams. In some VC firms, investors and operations staff prototype micro-workflows before involving engineers, accelerating experimentation and reducing dependence on small technical teams (I01; I06; I14; I15; I22). The boundaries of traditional roles also begin to blur. Investors with technical backgrounds increasingly co-design internal tooling. At the same time, low-code platforms and vibe coding approaches enable non-technical staff to build their own automations (I01; I14; I15; I21; I23).

In established VC firms, any adoption of AI systems requires integration with existing workflows and internal hierarchies. AI threatens junior positions the most, directly impacting traditional analysts' work, such as sourcing support, data preparation, and first-pass screening. Many VC firms are therefore dispensing with (or not replenishing) their analyst ranks, the traditional entry point into the industry (I23). This risks disrupting the industry's natural apprenticeship structure. VC firms that want to preserve junior layers only use AI to remove low-value coordination and shift the work of analysts and associates toward higher-value qualitative work and relationship building (I11; I14; I24).

In terms of hiring preferences, some VC firms emphasise AI literacy and quantitative reasoning. Others, however, still focus on domain expertise and relationship-building (I01; I10; I13; I17; I18; I20; I22; I23; I30). Overall, AI increases the diversity of internal structures, with some firms compressing junior layers and becoming more reliant on technology, while others make more incremental adjustments to the traditional apprenticeship model.

Alternative Strategic Positions

Despite their common quest for status, VC firms remain a diverse set of investors. AI does not push them toward a single optimal operating model. Instead, we observe an increasingly segmented landscape in which different archetypes of AI-use coexist. At one end are large, multi-stage platforms that treat AI as core infrastructure, supported by dedicated data and engineering teams, strong brands, and broad platform services (I05; I11; I21; I23; I27). At the other end are small, highly specialised VC firms that use automation to run lean operations (I03; I07; I14; I25; I28; I29). In between those two ends, VC firms are searching for coherent strategies of how best to integrate AI into their operations. Some interviewees talk of a "barbell" structure that leaves limited room for undifferentiated AI models in the middle (I21; I22; I26; I28).

AI-native VC firms deploy automation as a core element of their business model rather than as an incremental enhancement. Interviewees describe very small teams that automate large parts of sourcing, screening, memo preparation, and portfolio analytics, positioning themselves as technology-driven investment firms rather than traditional partnerships with auxiliary tools (I03; I11; I13; I25; I28). Some of these VC firms explicitly avoid leading rounds, instead merely co-investing (I02; I21; I25). Others still aim to lead early-stage rounds. They maintain partner-led founder engagement but use AI to widen their funnel and compress research time (I03; I07; I11; I14).

Another group of VC firms competes by narrowing scope and using AI to deepen specialization rather than maximise coverage. Sector- and region-focused managers describe custom classifiers, scoring systems, and local data pipelines that encode years of domain expertise and informal network knowledge into repeatable workflows (I07; I09; I14; I17;

I29). For these VC firms, AI supports systematically identifying "their" founders, often in niches that generic tools miss (I07; I09; I29).

AI is clearly poised to be a disruptive force in the VC industry, so that VC firms that fail to adapt to these changes may eventually be forced to exit the industry. In this context, it is also worth noting that the adoption of AI reflects a cultural self-understanding of VC firms, and especially their senior partners. The speed and extent of AI adoption depend more on internal dynamics than on technical feasibility. Adoption tends to accelerate when senior partners visibly champion it, pull data teams into their own workflows, and treat AI outputs as relevant input for investment decisions. However, if partners remain more sceptical and stick to existing routines, data teams remain peripheral, producing dashboards that rarely enter the firm's core debates (I11; I15; I19; I21; I23; I26).

A VC firm's openness to experimentation further affects the pace and depth of AI adoption. VC firms with bottom-up enthusiasm, top-down permission to iterate, and tolerance for prototyping integrate AI faster, often diffusing practices informally through daily collaboration and raising AI literacy expectations in hiring. More hierarchical VC firms, however, may constrain AI to more harmless uses such as transcription, bounded research support, and templated reporting. In this sense, AI acts as a stress test of culture. The same tools either remain peripheral and reinforce existing hierarchies, or they become central by reconfiguring how collaboration and authority are enacted inside the firm (I10; I12; I14; I15; I19; I21; I23; I26; I28; I29).

Differentiation in VC markets

We now consider how the diffusion of AI tools alters the bases of differentiation amongst VC firms. Prior to the AI revolution, early-stage VC markets were characterised by imperfect competition, with differentiation sustained by limited informational transparency and heterogeneous capabilities.[9] The advent of AI tools is changing this competitive logic, especially because of the commoditization of deal sourcing. As AI-enabled data providers and sourcing platforms are diffusing across the industry, potential deal candidates now appear simultaneously on the radar of many VC firms. Tools such as Harmonic, Spectre, Dealroom, PitchBook, and Crunchbase are increasingly treated as "table stakes" for operating VC firms (I06; I14; I17; I18; I20; I22). Once enough firms subscribe to the same vendors, signals and alerts converge, pipelines overlap, and any discovery edge becomes short-lived (I06; I13; I14; I17; I18; I21; I22). This greater transparency therefore increases competition for deals, shifting power away from investors and to founders. It might also increase the importance of 'reverse due diligence' where founders carefully evaluate which investors they want to bring into their company (I04; I12; I18).

To counteract this changing power balance, VC firms need to reconsider their competitive differentiation. AI-native VC firms seek durable advantages from proprietary overlays and ongoing infrastructure improvements. They build architectures on top of commoditised inputs, scraping and fusing vendors, adding internal labels, and embedding outputs directly into daily workflows (I03; I10; I11; I15; I19; I23; I25). Their competitive edge lies in how quickly models can be adapted and operationalised before practices diffuse (I03; I10; I11; I19). Traditional VC firms respond differently to these competitive pressures. Some develop proprietary data that rivals cannot access, even when external data feeds are shared. Others move even earlier to build relationships before companies appear in widely visible deal flow (I06; I07; I14; I18; I19; I23; I24; I27; I30). Overall, however, deal sourcing is likely to become a relatively minor part of how VC firms can differentiate themselves.

[9] This characterization is consistent with prior theoretical work: for example, Inderst & Müller (2004) model imperfect competition in VC using a search framework, while Hong et al. (2020) develop a theory of competitive dynamics among heterogeneous actors. Other studies further emphasise persistent differences in specialization as a source of differentiation (Gompers et al., 2009; Hochberg et al., 2015) and the importance of networks for deal sourcing, information flow, and performance (Sorenson & Stuart, 2001; Hochberg et al., 2007, 2010).

When many investors all receive the same alerts, who has a good chance at getting a meeting, let alone the deal? It is tempting to say the fastest and/or the best. This leads us to two important types of competitive differentiation: speed and brand. We now discuss those in greater detail.

AI is changing the speed at which investments are made. There is a race to be the first to identify and get to the deal. On the discovery side, we already noted that VC firms have built sourcing engines that trigger alerts (I03; I07; I10; I11; I13; I19; I23). Moreover, there has been a significant shift from reactive, inbound sourcing to proactive engagement, contacting founders weeks or months earlier than competitors, relying on manual deal flow and informal networks (I07; I19; I23; I25).

AI also speeds up the work after identification. Powerful AI platforms shorten the path from first signal to term sheet. Leads can be triaged quickly, founders approached earlier, and first-pass diligence conducted at a pace that slower organisations struggle to match (I07; I19; I25). Faster VC firms are more likely to anchor syndicates, and speed up approvals before others can compete (I07; I18; I19; I25; I26).[10]

Yet several interviews stressed that speed alone rarely decides outcomes. When multiple VC firms discover the same company via shared data providers or widely used sourcing platforms, founders may receive several offers in quick succession (I12; I18; I20; I22; I24). In these settings, speed has to be combined with other distinct advantages, related to brand, reputation, and networks.

[10] Not everyone agrees with the new emphasis on speed. Some VC firms avoid “hot” or heavily intermediated rounds where AI-enabled signalling drives rapid convergence. They still use AI to accelerate analysis, but move more slowly on pricing and papering, arguing that overpaying in competitive processes destroys fund returns. They see their competitive advantage as being disciplined and not being swayed by hasty market dynamics (I09; I16; I22; I23; I27).

There is a strong status ranking in the VC industry, where a small number of brand name VC firms attract the best dealflow and consistently outperform the rest of the industry (Podolny, 2001). This is partly due to self-fulling expectations, with founders paying considerable attention to status (I01; I02; I18). One interviewee estimated that over two thirds of all unicorns are backed by a tightly clustered group of VC firms, in part because many top founders decide early who they want on the cap table (I18). While AI-driven sourcing has equalised visibility, it has not equalised founders' preferences for attracting high-status investors (I12; I20; I22; I24).

Top brands also function as quality signals to other investors. When globally recognised firms lead early rounds, their participation is treated as a proxy for diligence and company quality (I01). Smaller VC firms often respond to these signals, effectively outsourcing conviction and hoping to get residual allocations after the lead investors have set their terms (I01; I09). Algorithmic VC firms might further reinforce this logic by allocating capital based on the historical performance of investors, which could further entrench high-status VC firms (I02; I18).

Another key dimension of competitive differentiation for VC firms is their reputation among limited partners. Some interviewees were not convinced that AI has fundamentally changed their brand perceptions. While limited partners increasingly question VC firms how they use AI, their allocation decisions remain anchored in realised returns, team reputation, governance quality, and brand familiarity (I09; I10; I12; I18; I27; I28).

Overall, the interviews suggest that AI has not flattened VC's status structure. Top brands still attract the strongest founders and partners. They are also in a good position to build bigger platforms and accumulate more data (I02; I08; I14; I18; I23). AI-native VC firms may operate more efficiently, but have to face the fact that reputation and trust ultimately govern status and access to the most valuable deals, and take a long time to acquire (I04; I14; I23; I24; I27).

Concluding Outlook

We are finally in a position to articulate our outlook on the likely direction of change that AI will bring to the VC industry. For this, we specifically focus on the role of human agency and the challenges that VC firms face in differentiating themselves in a fully AI-enabled investment environment. We identify three types of proficiencies that, in the foreseeable future, will continue to distinguish human actors from AI agents: (i) socially grounded conviction, (ii) gut feeling, and (iii) network relationships.

To fully appreciate the role of conviction, consider the so-called power laws which predict that most venture returns are concentrated in a few outsised winners (Cochrane, 2005; Kaplan & Schoar, 2005; Harris et al, 2014). These outlier ventures are more challenging to identify at the outset because they pursue high-risk, unconventional opportunities under opaque and ambiguous conditions. Investing in them requires high conviction and a willingness to be contrarian.

The natural question is whether AI tools are well-suited to identify such positive outliers. Several of our interviewees expressed concern that AI systems might overlook them, precisely because they don’t fit into the standard patterns (I13; I18; I23). A common practical response is to keep model parameters relatively broad, accepting more false positives in the hope of not eliminating those coveted outliers (I15; I23; I24). However, the very nature of these outlier ventures makes them essentially indistinguishable from outright bad ventures. Moreover, machine learning techniques are premised on success data. Yet these outliers often require a long time to generate the required success signals, making AI predictions less accurate.

An influential study by Bonelli (2025) shows that VC firms adopting AI technologies shift their investments toward startups that resemble past successes. While this improves the prediction of common outcomes, such as follow-on funding, it reduces funding for innovative startups that hold the potential to achieve rare, transformative success. Indeed, Bonelli's study shows that AI-powered VC firms were significantly worse at identifying ventures with breakthrough innovations.

It is tempting to dismiss this evidence by faulting the current implementation of AI systems for not paying enough attention to the outlier problem. However, LLM models actually excel at analysing data from multiple perspectives and can generate many more contrarian scenarios than any human could. The deeper problem is that AI agents cannot develop any true conviction for any of the perspectives they generate. In the VC context, they can identify numerous outliers, but they cannot meaningfully discriminate among them.

Human conviction goes beyond contrarian reasoning. It requires the formation of strong beliefs in the absence of credible data. Moreover, it requires a willingness to bear the consequences of any proposed decision. We argue that the problem is not that AI is unable to take contrarian positions; the problem is that, by itself, it cannot develop any true 'conviction' behind any of these positions. This limitation is related to what Felin et al. (2025) call "actor-specific decisions." These are decisions that are inherently linked to individuals. They are forward-looking, idiosyncratic, and require a combination of both experimentation and reasoning.

VCs' convictions are "socially grounded" because they are not merely the stubborn opinions of individualistic decision-makers. Instead, they are the outcome of social interactions that occur through professional networks. VCs constantly seek information and opinions from a vast array of people with different expertise. These human-to-human interactions convey nuances and emotions that current AI systems struggle to capture.

Socially grounded convictions do not need to be conformist. Through their networks, VCs learn how specific opportunities are viewed by others, but they may still disagree with the consensus view. Related to this, the work of Malenko et al. (2026) shows how some VC firms use ‘championing rules’ where individual partners can proceed with making early-stage investments, even if they don’t get majority support in the investment committee. A recent study from MIT venture competitions further suggests that ventures with more disagreements achieve superior performance (Gius, 2024).

Closely related to human conviction is the concept of gut feeling that we have already discussed extensively in this paper. For example, we noted how the evaluation of founder teams involves an interplay between cognition and intuition, resulting in judgment calls that are largely driven by gut feeling. The argument that AI systems cannot replicate gut feelings is simply that they are not sentient and therefore cannot possess any such feelings.

While current AI systems cannot replace human conviction or gut feelings, they can still support them in two important ways. First, they can organise the underlying information that human conviction and gut feelings are based on. Most of the AI tools used in the VC industry fall into that category. Second, they can also be used to refine human proficiencies by tracking the consistency of human decisions over time. Comparing current investment opportunities to past investment decisions, especially rejections which are more easily forgotten, VCs learn about the intertemporal consistency of their decision criteria. This kind of feedback can improve their own understanding of what drives their convictions and gut feelings. Put differently, AI systems cannot replace human decisions, but they can teach human decision-makers to become more consistent.

Another area where AI tools cannot replace humans concerns the relationships that VCs maintain across their industry networks. Whereas the first two human proficiencies, conviction and gut feelings, pertain to the investment stage, this third one pertains to value-adding support in the post-investment phase. Of particular value are warm introductions to key industry players, including customers, strategic partners, and other key stakeholders. VCs can also play a supporting role in assembling management teams and boards of directors. These networks can be supported by AI tools that map out networks (I12; I15; I19; I24; I30) or track relationship interactions (I09; I24; I29; I30). However, the trust that underpins them is based on human-to-human relationships that AI systems cannot replace. Put differently, Al can map the networks, but it’s people who activate them.

We conclude by noting that, precisely because current AI technologies cannot replace human proficiencies related to forming convictions, utilising gut feelings, and nurturing network relationships, these aspects are likely to become key sources of competitive advantage in the VC industry.

References

Agrawal, A., Gans, J. S., & Goldfarb, A. (2023). Artificial intelligence adoption and system-wide change. Journal of Economics & Management Strategy, 32(3), 557–576.

Bell, C. G. (1991). High-tech ventures: The guide for entrepreneurial success. Addison-Wesley.

Bernstein, S., Korteweg, A., & Laws, K. (2017). Attracting early-stage investors: Evidence from a randomized field experiment. Journal of Finance, 72(2), 509–538.

Bishop, C. M. (2006). Pattern recognition and machine learning. Springer.

Bommasani, R., Hudson, D. A., Adeli, E., et al. (2022). On the opportunities and risks of foundation models. arXiv Working Paper.

Bonelli, M. (2025). Data-driven investors. Review of Financial Studies, 34(7), 3226–3264.

Braun, V., & Clarke, V. (2022). Thematic analysis: A practical guide. Sage.

Cochrane, J. H. (2005). The risk and return of venture capital. Journal of Financial Economics, 75(1), 3–52.

Coleman, S., & Robb, A. (2009). A comparison of new firm financing by gender: Evidence from the Kauffman Firm Survey data. Small Business Economics, 33(4), 397–411.

Csaszar, F. A., Ketkar, H., & Kim, H. (2024). Artificial intelligence and strategic decision-making: Evidence from entrepreneurs and investors. Strategy Science, 9(4), 322–345.

Da Rin, M., & Hellmann, T. (2020). Fundamentals of entrepreneurial finance. Oxford University Press.

Doshi, A., Bell, G., Mirzayev, S., & Vanneste, B. (2025). Generative artificial intelligence and evaluating strategic decisions. Strategic Management Journal, 46(3), 583–610.

Ewens, M., & Rhodes-Kropf, M. (2015). Is a VC partnership greater than the sum of its partners? Journal of Finance, 70(3), 1081–1113.

Ewens, M., & Townsend, R. (2020). Are early-stage investors biased against women? Journal of Financial Economics, 135(3), 653–677.

Felin, T., Sako, M., & Hullman, J. (2025). Artificial intelligence and actor-specific decisions. SSRN Working Paper.

Flick, U. (2014). An introduction to qualitative research (5th ed.). Sage.

Fu, X., & Taylor, L. (2025). Due diligence and the allocation of venture capital. NBER Working Paper.

Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. Journal of Finance, 77(1), 5–47.

Gans, J. S., Stern, S., & Wu, J. (2019). Foundations of entrepreneurial strategy. Strategic Management Journal, 40(5), 736–756.

Gius, L. (2024). Disagreement predicts startup success: Evidence from venture competitions. Strategy Science, 10(2), 93–108

Gompers, P. A., Gornall, W., Kaplan, S. N., & Strebulaev, I. A. (2020). How do venture capitalists make decisions? Journal of Financial Economics, 135(1), 169–190.

Gompers, P., & Lerner, J. (2001). The venture capital revolution. Journal of Economic Perspectives, 15(2), 145–168.

Gompers, P., Kovner, A., & Lerner, J. (2009). Specialization and success: Evidence from venture capital. Journal of Economics and Management Strategy, 18(3), 817–844.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.

Gorman, M., & Sahlman, W. A. (1989). What do venture capitalists do? Journal of Business Venturing, 4(4), 231–248.

Harris, R. S., Jenkinson, T., & Kaplan, S. N. (2014). Private equity performance: What do we know? The Journal of Finance, 69(5), 1851–1882.

Hebert, C. (2025). Gender stereotypes and entrepreneur financing. Review of Financial Studies, 38(2), 1–48.

Hegde, D., & Tumlinson, J. (2014). Does social proximity enhance business partnerships? Management Science, 60(9), 2355–2380.

Hellmann, T., & Puri, M. (2002). Venture capital and the professionalization of start-up firms. Journal of Finance, 57(1), 169–197.

Hellmann, T., Mostipan, I., & Vulkan, N. (2025). Gender in start-up financing. Management Science.

Hochberg, Y. V., Ljungqvist, A., & Lu, Y. (2007). Whom you know matters. Journal of Finance, 62(1), 251–301.

Hochberg, Y. V., Ljungqvist, A., & Lu, Y. (2010). Networks as a barrier to entry. Journal of Finance, 65(3), 829–859.

Huang, L., & Pearce, J. L. (2015). Managing the unknowable. Administrative Science Quarterly, 60(4), 634–670.

Kahneman, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux.

Kanze, D., Hung, L., Conley, M. A., & Higgins, E. T. (2018). We ask men to win and women not to lose. Academy of Management Journal, 61(2), 586–614.

Kaplan, S. N., & Schoar, A. (2005). Private equity performance: Returns, persistence, and capital flows. The Journal of Finance, 60(4), 1791–1823.

Kaplan, S. N., Sensoy, B. A., & Stromberg, P. (2009). Should investors bet on the jockey or the horse? Journal of Finance, 64(1), 75–115.

Malenko, A., Nanda, R., Rhodes-Kropf, M., & Sundaresan, S. (2026). Catching outliers: Committee voting and the limits of consensus when financing innovation. Forthcoming in Journal of Finance.

Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management. Academy of Management Review, 46(1), 192–210.

Russell, S. J., & Norvig, P. (2021). Artificial intelligence: A modern approach (4th ed.). Pearson.

Sahlman, W. A. (1990). The structure and governance of venture-capital organisations. Journal of Financial Economics, 27(2), 473–521.

Samila, S., & Sorenson, O. (2011). Venture capital, entrepreneurship, and economic growth. Review of Economics and Statistics, 93(1), 338–349.

Sorenson, O., & Stuart, T. E. (2001). Syndication networks and the spatial distribution of venture capital investments. American Journal of Sociology, 106(6), 1546–1588.

Tyebjee, T. T., & Bruno, A. V. (1984). A model of venture capitalist investment activity. Management Science, 30, 1051–1066.

Appendix: Interview Participants (Anonymized; Non-Identifying Descriptors)

ID	Role	Stage Focus	Fund Size	Region
I01	Investment Partner	Growth / Late-Stage	Small	North America
I02	Senior Executive	Growth / Late-Stage	Small	North America
I03	Investment Partner	Early-Stage	Small	Europe
I04	Data/AI Leader	Multi-Stage	Large	Europe
I05	Data/AI Leader	Multi-Stage	Large	Europe
I06	Data/AI Leader	Multi-Stage	Large	Europe
I07	Investment Partner	Early-Stage	Small	Europe
I08	Data/AI Leader	Multi-Stage	Large	North America
I09	Investment Partner	Multi-Stage	Medium	Europe
I10	Data/AI Leader	Growth / Late-Stage	Large	Europe
I11	Data/AI Professional	Multi-Stage	Large	Europe
I12	Senior Executive	Multi-Stage	Medium	North America
I13	Data/AI Leader	Growth / Late-Stage	Small	Europe
I14	Investment Professional	Early-Stage	Small	Europe
I15	Data/AI Leader	Multi-Stage	Medium	North America
I16	Senior Executive	Early-Stage	Small	Europe
I17	Investment Partner	Growth / Late-Stage	Small	Middle East
I18	Investment Professional	Early-Stage	Small	Europe
I19	Data/AI Leader	Growth / Late-Stage	Large	North America
I20	Senior Executive	Early-Stage	Large	Europe
I21	Data/AI Professional	Early-Stage	Small	Europe
I22	Investment Professional	Early-Stage	Medium	Europe
I23	Data/AI Professional	Multi-Stage	Large	North America
I24	Investment Partner	Early-Stage	Small	Europe
I25	Investment Partner	Early-Stage	Small	North America
I26	Investment	Early-Stage	Small	North America

	Professional			
I27	Investment Partner	Multi-Stage	Large	Europe
I28	Investment Partner	Early-Stage	Small	Asia-Pacific
I29	Data/AI Professional	Early-Stage	Small	Europe
I30	Senior Executive	Multi-Stage	Medium	Europe

Note. Descriptors reported in this table were coded from participants' self-descriptions at the outset of each interview and verified against publicly available sources. Where details were missing, those sources were also used to supplement the record and ensure consistent categorization. Role is classified into five categories: Senior Executive (C-suite or equivalent senior management with firm-level operating oversight); Investment Partner (partner-level investing roles with ultimate or shared authority over investment decisions); Investment Professional (non-partner investing roles involved in sourcing, diligence, execution, and portfolio work); Data/AI Leader (senior functional leadership responsible for data, analytics, and AI strategy, capability, and governance); and Data/AI Professional (hands-on technical roles executing data and AI work). Stage focus is coded as Early-Stage (Pre-Seed–Series A), Growth/Late-Stage (Series B+), or Multi-Stage (Pre-Seed–Series B+). Fund size is coded into assets under management (AUM) bands (USD-equivalent) based on disclosed AUM. Where AUM is not disclosed, the latest publicly reported fund size or total capital raised across funds is used as a proxy. Bands are Small (less than $500M), Medium ($500M to less than $2B), and Large ($2B or more). Region denotes the interviewee's primary office location (i.e., where the participant is based) and may differ from the investment firm's headquarters and/or the fund's legal domicile.