MARCH 2026

Responsible AI: The VC Perspective

AUTHORS
Oliver Nixon & Elad Vromen

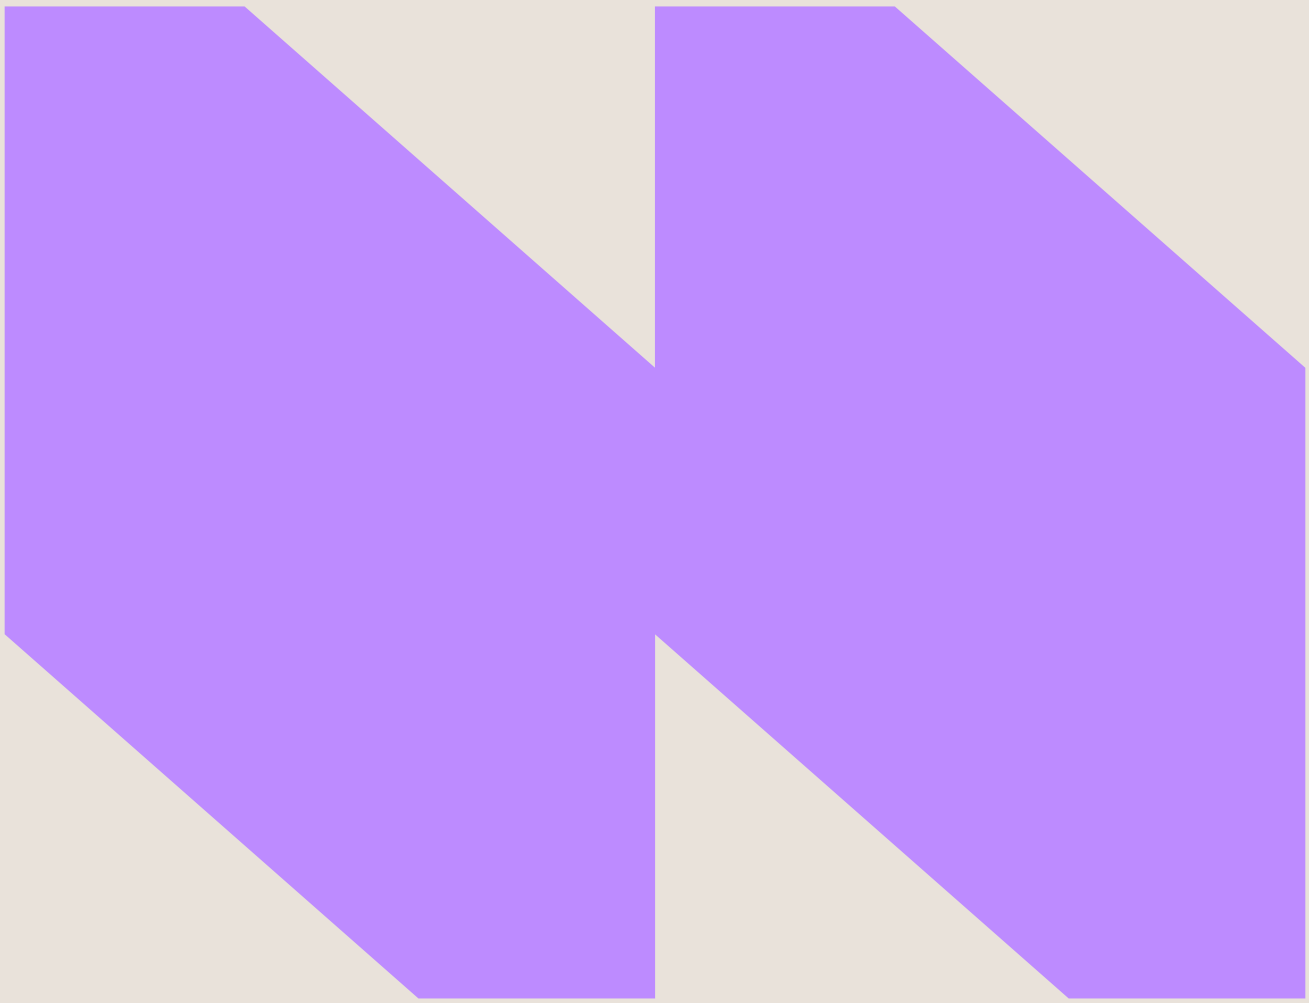

REFRAME VENTURE

Project Liberty Institute //

IMPACT VC

OVERVIEW

Contents

AUTHORS
Oliver Nixon
Elad Vromen

CO-PRODUCED BY
Project Liberty Institute //
IMPACT VC

SUPPORTED BY
zendesk

Introduction

This paper documents observations from a systematic and lengthy analysis and investigation of a powerful global group of key investors in early-stage companies who exert an outsized influence on the development and proliferation of AI technologies and applications.

Over the past 18 months, Reframe Venture, Impact VC, and Project Liberty Institute have collaborated on an extended research and engagement project on responsible and impactful AI, involving over 200 VC and growth equity funds and 80 LPs managing more than $6 trillion in AUM through conferences, panels, networking events, trainings, and workshops across Paris, Berlin, London, Tokyo, Singapore, New York, San Francisco, and Cape Cod.

We have also been running a monthly series of working groups on the topic, bringing together a core group of dedicated practitioners from VC firms to steer the development of VC-specific resources, such as our ***Responsible AI Due Diligence Tool*** and scenarios on AI's impact on ***climate*** and the ***workforce***.

Throughout the summer of 2025, we held over 30 in-depth interviews with VCs, growth equity investors, LPs, and expert practitioners from across key global markets to gather perspectives and reflections on the state of responsible AI in VC investment. In total, these interviewees came from organisations with a cumulative AUM of over $500 billion. From these interviews, we compiled a series of draft observations on this influential part of the tech ecosystem.

To test the depth of these observations across the ecosystem, and to pick out geographical and other demographic-led trends, between December 2025 and February 2026, we ran a first-of-its-kind in-depth survey with VC practitioners across different roles, investment geographies, fund sizes, and stage and sector focuses (n=56). We recognise that being a responsible investment organisation, our network is likely to select somewhat towards individuals with supportive views on responsible AI. However, as indicated by the demographic breakdown of the survey, it does not represent esoteric, small investors, instead capturing the views of some of the largest VC firms in the world. Furthermore, we do not present survey findings in isolation: the questions were explicitly and carefully designed to test the observations from our interviews and convenings. It is the results of this survey which are reported as the quantitative findings which support the observations below.

SURVEY RESPONDENTS

Geographies of Investment Teams

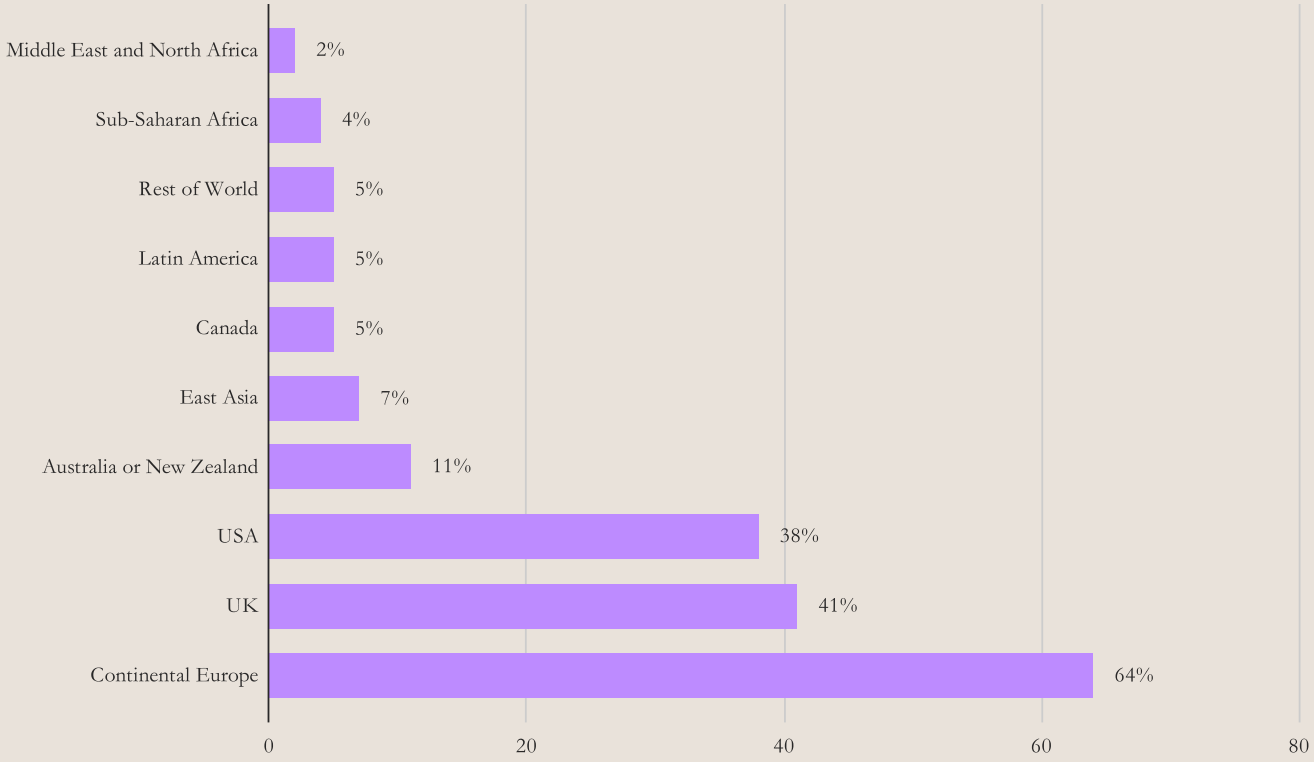

Role Distribution

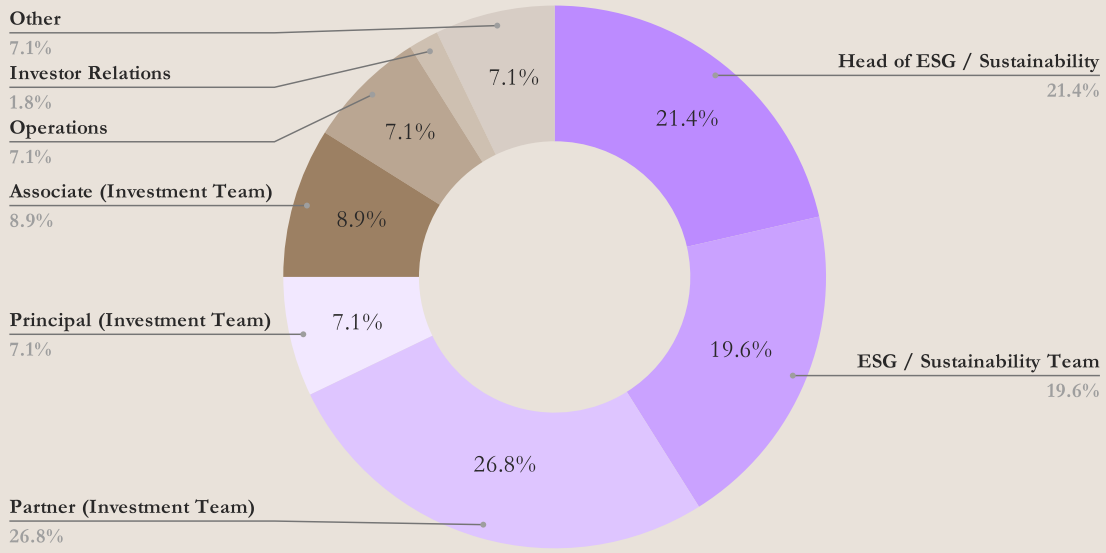

Investment Stage Focus

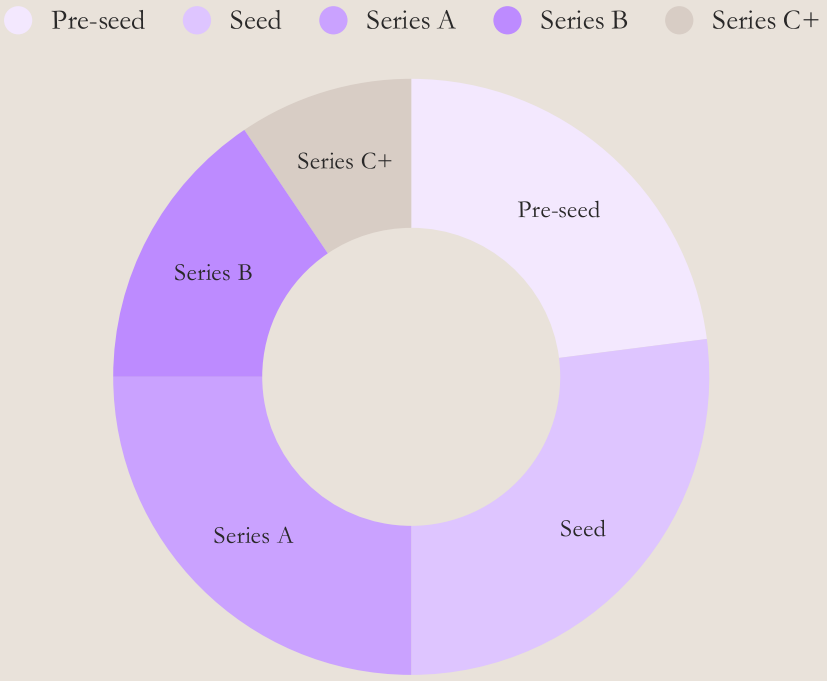

Years of VC Experience

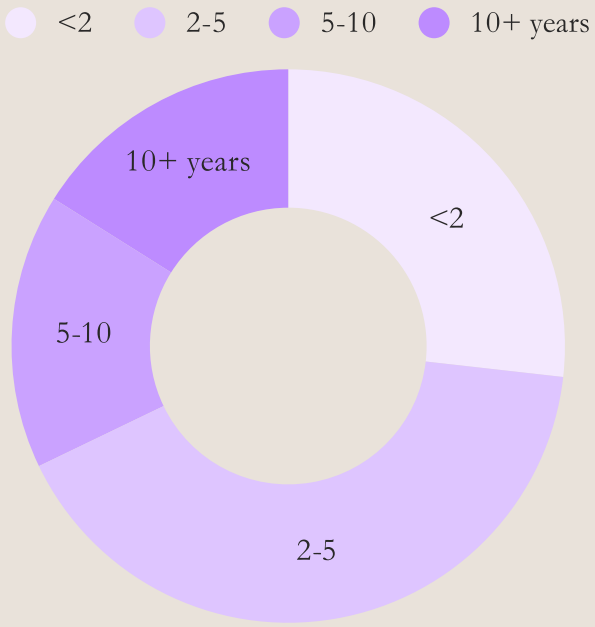

Fund Size Distribution

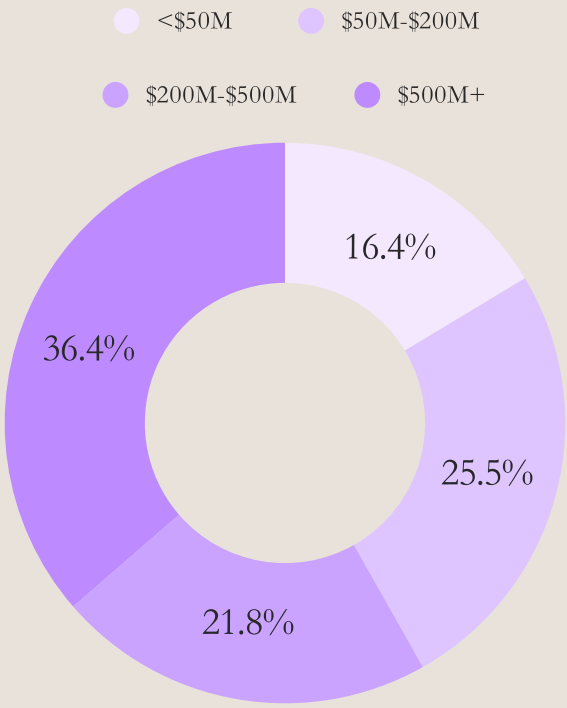

Executive Summary

This paper presents findings from 18 months of research and engagement with over 200 VC and growth equity funds and 80 LPs managing more than $6 trillion in AUM, culminating in 30+ in-depth interviews and a first-of-its-kind survey of 56 VC practitioners across global markets.

- The central finding is a strong and broadly held belief that responsible AI is good business. 73% of VCs surveyed believe companies with stronger responsible AI practices will be more financially successful, driven by two dynamics: avoiding costly incidents and earning the customer trust needed for real-world adoption. 84% see direct investment opportunity in companies that make responsibility core to their proposition, and 91% see significant opportunity in AI safety infrastructure itself, drawing a parallel to the emergence of cybersecurity as an investable market.

- This conviction is matched by serious concern about systemic risks. 84% of VCs expressed concern about AI's negative impact on human agency, 96% view malicious use of AI as a substantial societal threat, and 74% are worried about environmental impacts from AI and data centres. 88% believe VCs have a responsibility to consider these systemic risks.

- Yet a significant gap persists between belief and practice. Only 55% include responsible AI in due diligence, only 14% rate their AI risk assessment capabilities as "good," and just 27% feel they have sufficient internal expertise on the topic. The main barriers are difficulty evaluating risks and insufficient expertise, compounded by coordination challenges across ESG, legal, technical, and investment teams. These same knowledge gaps extend to portfolio companies, where only 42% of VCs believe their startups adequately understand responsible AI risks.

- Regulation adds complexity rather than clarity. While GDPR is universally respected, the EU AI Act generates ambivalence, and regulatory uncertainty is a top concern for investors. Geographic attitudes diverge sharply, with European investors treating compliance as business reality, US investors framing it as competitive threat, and Asia emerging as a surprisingly progressive region.

- LP pressure is converging globally, with US limited partners rapidly increasing engagement on responsible AI to match levels European LPs had already established. This upstream pressure, combined with growing procurement expectations from enterprise customers, is beginning to shift incentives across the investment chain.

- Additionally, responsible AI is evolving from a risk management concern into an investment thesis: encompassing both the infrastructures that enable responsibility and the broader positive impact opportunities in healthcare, education, climate, and financial inclusion that trustworthy AI can unlock.

1. Belief in The Business Case

Our main finding is the overwhelming belief from VCs globally that AI companies which prioritise building AI responsibly perform better than companies that push these considerations aside.

73%

of survey respondents believe that companies with stronger responsible data and AI practices are likely to be more financially successful.

This belief was **particularly strong amongst the more experienced VCs** with more than 5 years of experience (83%).

The broader investment opportunity signal is even stronger: **84% see investment opportunity in AI companies that make responsibility a core part of their business proposition.**

This quantitative finding reflects the very strong signal we have heard from VCs and LPs in our interviews and convenings over the past eight months. In our experience hosting events in Europe, the USA, in Singapore, and in Japan, this high-level of general belief in the business case does not have significant geographical variation, which was supported by the survey results.[1]

The main drivers of the business case are simple: responsible AI avoids incidents and gains customers.

The former has been known for years, with resources such as the AI Incident Database proving especially valuable in demonstrating the very real, and financially material, nature of AI risks. Case studies such as the collapse of ***AllHere Education***, ***Anthropic's $1.5bn IP settlement***, the ***Workday lawsuit***, and many more, have contributed to a growing library of evidence that missteps when it comes to responsibility in AI development can cause significant damage. One US-based investor we spoke to cited the ***Google and Character.AI lawsuit over teenage suicide*** as indicating *"massive liability exposure,"* adding: *"Nobody would invest in a car company making exploding vehicles. Basic due diligence is needed for the stability of investments."*

Evidence for the second driver of the business case, that responsible AI gains customers, has emerged rapidly through 2025. With evidence

1 In our survey, VCs investing in the US expressed 76% agreement, only slightly higher than the 64% in the group of VCs investing only in Europe.

that adoption rates in companies of all sizes are slowing, and of the failure of AI experiments, we're seeing a real issue with AI becoming adopted in the real economy.

Much of this can be attributed to a lack of trust in AI systems (see ***MIT NANDA 2025***), and we're increasingly seeing ***stricter procurement expectations from larger companies***. Whilst AI models are improving in capabilities with time, they lack the security, accountability, and transparency to be reliably integrated into business workflows.

In addition, whilst human-interaction is an excellent way to mitigate these drawbacks, AI systems are rarely developed to augment human work, instead prioritising the complete automation of it – better product design is a major competitive advantage here.

There is a need from VCs for guidance and valuation signals on the complexity of the problems enterprises face, and their strict requirements for trust in AI.[2]

2 It is of particular concern that whilst ***founders are getting younger***, they are also ***focussing more on enterprise software***.

2. Significant Concern of Systemic Risks

With 84% of surveyed VCs concerned about AI's negative impact on human agency.

AI responsibility is not just about enhancing business success, we also observed high levels of concern over AI's systemic risks. We asked about several of these in our survey and interviews. In one of our most significant survey results, 84% of surveyed VCs expressed concern about AI's negative impact on human agency. As AI takes on more roles in our daily lives, there is growing evidence that reliance can reduce individuals' ability to make value judgments, form beliefs, and pursue actions in line with their own values. ***A 2026 study by Anthropic*** found that mild disempowerment on these lines already appears in 1 in 50 conversations with their chatbot, Claude.

In our ***2026 scenario analysis of AI's impact on the workforce***, we wrote about the opportunities in designing products to augment rather than automate human capabilities, and that this should be reflected in the ***benchmarks used for AI models***. Such products are likely to achieve greater trust in their systems, which is ***increasingly required by corporate customers***, in addition to allow AI to function as a true productivity enhancement tool, with far less risk of systematic degradation of human agency.

Such opportunities are also likely to lessen the impact of AI on the labour market, which 66% of survey respondents believe is a risk threatening democracies. Agency is not the only social systemic risk which garnered concern, with **96%** of VCs surveyed believing that malicious use of AI is a substantial societal threat, which is reflected in the proliferation of such incidents as documented by ***MIT's AI Incident Tracker***.

On the environmental impact, which we also ***published a series of scenarios for in 2026***, there was significant concern in the short term, and uncertainty about the net impact. **74%** of our survey respondents stated that they are concerned about the environmental impact of AI and related data centres, with only 4% disagreeing.

Later in the survey, we wanted to find out whether, given this high level of concern for systemic societal risks, VCs thought it was their responsibility as investors to think about them.

Our survey respondents overwhelmingly thought it was

88%

agreed that VCs have responsibility to consider systemic societal AI risks.

3. Barriers to Responsible AI in VC

Despite this high level of belief in the business case and concern over systemic risks, implementation by VCs has been less comprehensive.

Only 14% rated their firm's current AI risk assessment capabilities as 'good', and only 17% rate their ability to help post-investment as 'good' or better. This appears to scale slightly with fund size (26% for funds over $200m, vs 5% for funds under $200m).

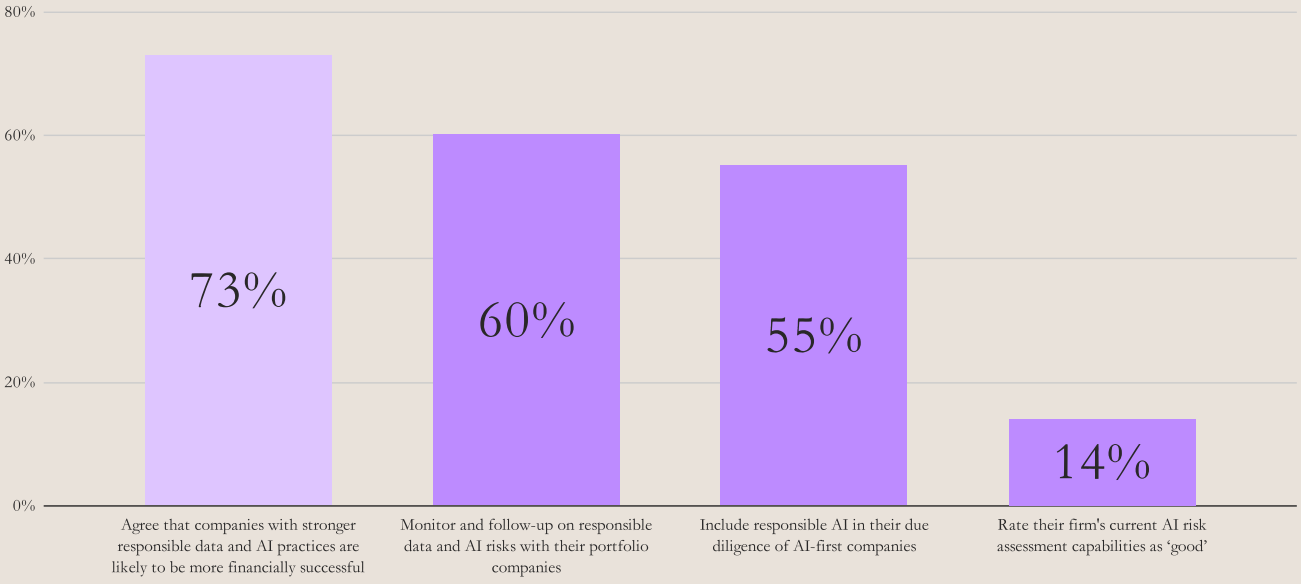

Lack of knowledge

We asked survey respondents and interviewees what the main barriers to better responsible AI due diligence are, and found that difficulty evaluating risks (65% of respondents indicating this as a top factor) and insufficient expertise (57%) are the top barriers, followed closely by lack of time (39%).

'Lack of time' was especially a factor for VCs investing in the USA, cited by 72% of US-exposed VCs compared to just 30% of VCs investing only in the EU.

This echoes what we have been told repeatedly in interviews and during events: there is a significant lack of knowledge and expertise on responsible AI in VC.

Whilst 74% of respondents felt they had sufficient internal expertise on ESG generally, only 27% felt that they had sufficient expertise on responsible AI.

Furthermore, only 26% of respondents believe that there is consensus in their investment team on what responsible AI is.

Many people we spoke to emphasised the need for venture-specific practical guidance, and this was backed up in our survey, with only 12% believing that existing guidelines are available and usable, a number that dropped to 4% for investment team members. It was this strong demand which led to our development of the ***Responsible AI Due Diligence Tool*** in December 2025.

Coordination difficulties

A secondary reason for the lack of implementation is the coordination difficulties caused by the multi-disciplinary nature of responsible AI in VC. Responsible AI factors fit somewhere between the:

1. ESG/Sustainability teams (with governance expectations, incidents having significant social implications, and a lack of transparency around energy use)

2. Legal teams (with confusing, untested legislation of varying focus, scope, and enforcement across Europe, US federal and state levels, and APAC; and unclear interactions to existing copyright, data privacy and security, and sector-specific regulations)

3. Technical due diligence (integration of software engineering best practices such as data minimisation, efficiency, privacy, security)

4. Investment teams (with factors such as reliability, accountability, security, and bias mattering to different degrees with different business models and to different customers, over different time-scales).

Only

27%

of respondents felt they had sufficient internal expertise on responsible AI

This identity crisis has consequences deeper than just organisational coordination issues. Some VCs and LPs we spoke to worried that responsible AI's positioning under 'ESG' or 'compliance' in some firms could actively deter strategic engagement from technical and commercial teams who might otherwise champion it.

Such coordination issues can be further confused by outsourcing elements of due diligence, which 69% of survey respondents said that they sometimes do with responsible AI. And furthermore, as noted by a VC investor we spoke to, outsourced providers often *"lack sufficient expertise and robust methodologies for responsible AI assessment,"* with *"nobody yet asking difficult enough questions to raise the game."* One partner at a global PE and VC firm with over $100bn AUM, who had spent six months building enough expertise to *"challenge provider outputs"*, said that they had recently *"disagreed with peer benchmarking results from our tech due diligence provider"*, indicating that quality concerns could be widespread but undetected.

4. VCs See Knowledge Gaps in Startups

The knowledge-gap is an issue in startups too, VCs believe.

Without the pressure from VCs, we've seen implementation of responsible AI take a backseat as startups develop products, despite the growing evidence of it as a requirement for widespread adoption.

We found that only 40% of VCs surveyed agreed that their portfolio companies have adequately assessed AI risks, reflecting the lack of pressure on this topic from VCs at due diligence and post-investment.

Startups are also subject to the expertise problem: only 42% of VCs surveyed thought that their portfolio companies actually understood what responsible AI risks were. This number was even lower, at 24%, for growth-stage investors surveyed – indicating that AI companies are not maturing in their knowledge of responsibility as they grow.

That said, expectations of when in their growth companies should begin to think about formal responsible AI governance vary. Series A is the most common expectation (41% of survey respondents), but a reasonable portion expect formal governance earlier, at Seed (22%), or later, at Series B (22%).

One one general counsel at a global VC firm with over $10bn AUM observed that Series B companies *"behave as if they are seed"* when it comes to RAI, citing founder intentions and mission rather than implementing concrete governance mechanisms.

Founder intentions are a poor proxy

In our interviews and discussions, we witnessed a chain of responsibility: LPs stating that it was VC's responsibility to think about RAI in their investments, but without adequate ways of assessing this, and VCs, without knowledge and clear expectations, using founder intentions as a proxy for responsible AI development.

Our survey results reveal that founder intent is relied upon frequently as a proxy for responsible AI: 38% rate it as a good proxy and 47% as an adequate proxy.

However, the sense we have from interviews is that we are operating in unknown waters here, with 42% of those surveyed saying it was 'too-early-to-tell' whether this faith in the founder was effective. It's noteworthy that more experienced VCs with over 5 years of experience were more than twice as likely to say that they have experienced founder intentions which didn't materialise, than their less experienced counterparts.

Pressure from market starting to come through

The "speed versus responsibility" trade-off, repeatedly cited as justification for delaying responsible AI adoption, is increasingly recognised as a false framing by investors we spoke to. In regulated industries in particular – healthcare,

defense, financial services – moving fast simply isn't viable. A London investor working with mission-critical applications noted: *"[startups] can't move fast in these industries, but they can certainly scale."*

The real moat to adoption in the real economy is building systems that are trustworthy. Multiple investors emphasised that *"doing things slow in the beginning allows you to build well and move fast"* later, particularly as technical debt compounds and regulatory scrutiny intensifies. As a New York-based partner of a >$1bn VC firm argued, *"AI is more suited to 'move slow to move fast' than the classic 'move fast and break things'."*

A couple of VC and LPs that we spoke to indicated that in some forums, there has been a failure to communicate responsible AI in terms that founders understand: as software quality, data science rigor, and engineering best practices rather than standalone ethical obligations. As one UK-based investor noted: *"The name is contaminated – Responsible AI – too much ethics. It should be approached as a technical issue, like cybersecurity."*

In many industries, and particularly as high-profile AI incidents hit the headlines, pressure from corporate clients through stricter ***procurement expectations*** is moving the needle in this direction.

As one medical-tech expert practitioner explained: *"Compliance is a constraint, but consumer value is the primary driver for the standards being applied. RAI is not an external, value-laden add-on, it [needs to be] explicitly embedded into design choices."*

Many VCs described to us how they see their role in providing this commercial guidance to their portfolio companies. As one pre-seed impact investor stated: *"Educating is VC's role! VCs should provide resources for responsible AI to founders."*

> *AI is more suited to 'move slow to move fast' than the classic 'move fast and break things*
>
> **A New York-based partner at a >$1bn AUM generalist VC firm.**

5. Regulation: not guardrails or speedbumps, but a blind corner?

In our interviews, survey, and convenings, there is a high level of variation in how regulation is perceived and operationalised. GDPR emerged as the lone universally respected framework – described by one investor as having "from the start" implementation expectations, thought about proactively and included in roadmaps from seed stage. In contrast, the EU AI Act has generated profound ambivalence, with just over half of survey respondents saying they had a 'favourable view of the EU AI Act', but this differed significantly between ESG/Sustainability team members (74%) and investment team members (35%). VCs who only invest in Europe were more likely to hold favourable views of the Act (65%), compared to US-based or US-exposed investors (44%). Interestingly, there isn't much of a relationship between respondents who believe that responsibility in AI is financially material and favourability towards the EU AI Act, i.e., the regulation designed to promote that responsibility.

The core criticism wasn't regulatory burden per se, but instability and speculation. 'Regulatory uncertainty' was a top-3 concern about regulation for 67% of survey respondents, outweighing competitive disadvantage, liability risks, technical requirements, and compliance costs. As one sustainability manager at a large ($1bn+) UK-based VC firm explained: *"the EU AI act is problematic: the approach wastes money on changes in regulation, and then companies waste money and time on speculation – this actually delays implementation."*

Geographic patterns revealed fundamentally different relationships with regulation. In our interviews, European investors treated compliance as normal business reality – embedded in early-stage thinking and sometimes viewed as a competitive advantage for building trust. American investors often framed regulation primarily as a competitive threat a company has to overcome in order to expand geographically, rather than an enabler. Within the domestic context, the US approach centers on litigation risk management rather than proactive compliance, with sector-and risk-specific regulations used rather than horizontal AI governance legislation. With a complex web of US state-level AI regulations emerging, and the looming legal uncertainty of the executive order banning such regulation, there are significant uncertainties in the US when it comes to responsible AI regulation.

The UK has emerged as a potential middle ground, with interviewed investors highlighting its "innovator-regulator mindset" and the possibility of serving as a sandbox for EU-aligned practices with lower initial friction. In addition, a number of interviewees identified Asia as "the most progressive region" for responsible AI maturity, with relatively sophisticated regulatory frameworks in South Korea, China, Japan, and Singapore, plus higher levels of portfolio company curiosity and stakeholder engagement quality.

6. LP pressure growing, especially in the US

The chain of responsibility identified in Observation 4 – LPs to VCs, VCs to founders – raises the question of whether upstream pressure from limited partners is beginning to shift incentives. The short answer: it's starting, but the picture is more nuanced than it first appears.

When asked whether LP pressure on responsible AI is likely to affect their next fundraise, the results are similar across geographies: 56% of EU-only VCs, 53% of US-exposed VCs, and 50% of others think it will at least "come up."

But when asked whether pressure has increased over the past two years, the picture diverges dramatically: 41% of US-exposed VCs report increased pressure, compared to just 5% of EU-only VCs. Taken together, these two findings suggest that US LPs have recently increased their engagement on responsible AI up to a level that European LPs had already established. The convergence in forward expectations, combined with the divergence in recent change, indicates a catching-up rather than a gap.

This aligns with what we heard in interviews. In Europe, responsible AI questions have been somewhat, albeit sometimes naively, a part of the LP-GP conversation for some time, embedded in broader ESG, sustainability, and cybersecurity expectations. In the US, where federal AI regulation remains fragmented and ESG-specific language carries political baggage, LPs appear to be stepping into a governance vacuum – asking direct questions about AI risk management in ways that bypass the ESG framing entirely.

As one sustainability lead of a global asset manager (over $100bn AUM) said: *"it's the LP's job to remind VCs to communicate issues and risks, and right now, that includes responsible AI".*

Whether this convergence in LP expectations proves sufficient to shift VC behaviour at scale remains to be seen. But the direction of travel is clear, and suggests that LPs – particularly in the US – may prove a more potent catalyst for responsible AI adoption than regulation has been to date.

56%
of EU-only VCs

53%
of US-exposed VCs

50%
of others

think responsible AI will at least 'come-up' in the next fund-raise

7. Investing in the 'Responsibility Stack'

91% of our survey respondents agreed that there is significant investment and positive impact opportunity in companies developing AI infrastructure that prioritises and enables responsibility – such as models, assurance, and certification. This was the single highest agreement item in our entire survey, reaching 100% among VCs with more than five years of experience.

91%

agree 'Good AI' is a significant investable opportunity

The broader investment opportunity signal is similarly strong: 84% see investment opportunity in AI companies that make responsibility a core part of their business proposition.

This finding moves responsible AI from a risk management concern – the focus of Observations 1 through 6 – into an investment thesis.

As described in Observation 3, many VCs lack the internal capacity to evaluate and support responsible AI practices. As described in Observation 4, many startups lack the expertise to implement them. As described in Observation 5, regulation provides principles but uncertainty over operationalisation.

Our interviews revealed that the current landscape – with big tech producing some standards that mainly serve competitive interests, while nonprofits focus on frontier risks rather than deployment guidance – leaves a vacuum in practical, commercially viable safety infrastructure. **This vacuum is an investable gap.** The companies that fill it serve a market with structural demand from every AI-deploying company that needs to earn customer trust, meet regulatory requirements, or simply build products that work reliably.

VCs are uniquely positioned to catalyse this market: through direct investment, and through aggregating demand for these products across their portfolio. When VCs systematically assess whether companies have engaged appropriate safety infrastructure, it shifts founder incentives in a way that voluntary frameworks alone cannot. The parallel to cybersecurity is instructive. In the earlier days of IT adoption, cybersecurity was perceived as a cost centre and compliance burden. Today, it is understood as essential infrastructure, with a thriving commercial ecosystem serving companies of every size and stage. AI safety infrastructure is on a similar trajectory – with the added tailwind that VCs' own portfolio interests are directly aligned with its development. Our survey results suggest the investment community already recognises this opportunity; the question is how quickly it materialises into deployed capital.

8. AI's Positive Potential

Beyond the opportunity in responsible AI infrastructure, our engagements and survey revealed broad enthusiasm for AI's positive social and environmental impact – and critically, this enthusiasm extends well beyond the impact investing community.

82%

of respondents agree there is a substantial positive social impact opportunity in AI technology.

Environmental impact opportunity drew a lower but still highly significant **67% agreement**, with only 11% actively disagreeing. The lower figure likely reflects a more nuanced assessment of AI's environmental net benefit, given the energy consumption concerns identified in Observation 2 (where 74% expressed concern about data centre environmental impact). VCs are simultaneously excited about AI's potential to address environmental challenges and worried about its own environmental footprint – we've attempted to address this uncertainty through a workshop with VCs, LPs, and experts, and the resulting scenario analysis we published in ***our 2026 write-up on the topic***.

The environmental impact question showed more demographic variation than the social impact question. Growth-stage investors were notably more sceptical (57%) than early-stage investors (74%), likely reflecting closer proximity to the energy and infrastructure realities of scaling AI systems. US-exposed VCs were also more sceptical (57%) than EU-only VCs (71%), consistent with the geographic patterns on environmental concern observed throughout our survey and interviews.

In our interviews and convenings, we heard consistent interest in how responsible AI practices can be enabling of impact, not merely constraining of harm. In healthcare, responsible AI means systems that clinicians trust enough to adopt, extending diagnostic capabilities to underserved populations. In education, it means age-appropriate design that prevents cognitive over-reliance and protects sensitive data, while unlocking personalised learning. In climate, it means efficient low-cost models and transparent data practices that make AI-driven solutions credible to policymakers and customers.

Tackling these huge societal challenges represents a significant business opportunity, and one that is supported by the advantages in hiring, sales, funding, and regulation that companies pursuing positive impact benefit from, as described in detail in ImpactVC's ***How impact drives business value creation*** report.

Realising these opportunities for impact requires building the trust that in Observation 1 was identified as a key driver of business value. This reframing – from responsible AI as risk mitigation to responsible AI as an impact enabler – represents a significant investment opportunity. We will be exploring this in detail through upcoming sector-specific investment memos examining the intersection of responsible AI and impact across healthcare, education, climate, and financial inclusion.

Join our community

REFRAME
VENTURE

reframeventure.com