Open Future Forum

OPEN FUTURE FORUM

VC and CVC AI Investment Report

NOVEMBER 2026

November 2026 | Edition 4 | Research report | Murray Newlands | Open Future Forum

Prepared 2 November 2026. Evidence cutoff: 31 October 2026, 23:59 America/Los_Angeles. Updated investor ownership and impact measures inform investment-thesis analysis. No separate VC-versus-CVC estimate is supported by the available instrument.

About this publication: Open Future Forum publishes research for executive and investor audiences. Murray Newlands is its founder. This report separates OFF survey observations, external research, author analysis and illustrative scenarios. Vendor-authored guidance is identified as such and does not constitute an independent product assessment. The collection and supporting methods are available through the OFF research index: https://openfutureforum.com/research/

4,321 substantive body words | 5,424 total words
Tables, FAQs, methods and references are excluded from the body count.

In this report

01 The decision in brief
02 Executive summary
03 What November adds
04 Evidence and sample boundaries
05 1. Translate the product claim into a customer choice
06 2. Follow a cost claim through the operating system
07 3. Distinguish customer help from durable demand
08 4. Test the advantage after the underlying capability improves
09 5. Keep strategic and financial objectives legible
10 6. Write the disconfirming evidence into the investment paper
11 External evidence and its limits
12 An illustrative investment-committee review
13 The next investment discussion
14 Investment questions
15 Citable facts
16 Methods and limitations
17 References
18 Related research and citation
19 Editable working tool

The decision in brief

Underwrite a specific mechanism by which AI improves a product, delivery cost or customer result, then identify evidence that could disprove its durability. Adoption and favorable investor impressions are useful starting points, but neither establishes an investable advantage. Financial and strategic investment objectives should be stated separately so that one cannot silently substitute for the other.

Executive summary

The updated OFF investor instrument records reported portfolio impact, not investment returns. Of 364 respondents, 190 selected better products, 141 cutting costs, 133 helping customers and 51 nothing measurable yet. The question permits multiple selections. Eight respondents selected a positive impact and nothing measurable yet, so the positive union is 321 of 364, or 88%, rather than the complement of the 14% selecting nothing measurable. [INVIMPACT; INV_POSITIVE]

Ownership is measured separately. Of 373 respondents, 185 selected the CEO as increasingly owning the AI buying decision, 97 the CIO or CTO, 43 finance, 72 individual teams and 62 too early to say. Those answers describe respondents' views across their portfolios. They do not identify who approved a particular purchase or establish that executive involvement caused better performance. [INVOWN]

The November expansion adds a later investor event cohort and reconciles it with the published baseline. It does not create a representative monthly panel. The changed cumulative distribution reflects added respondents, overlap and the latest eligible answer selection. A difference from October should therefore not be described as a measured acceleration in portfolio returns or a change experienced by the same companies.

The practical investment question is what must remain true for an AI thesis to work. A better-product thesis requires evidence that customers value the improvement and that the company can retain part of that value. A cost thesis requires an economically relevant denominator and a credible route from operating gains to financial benefit. A customer thesis requires an accepted result in the relevant customer population. Each has different failure conditions.

This report supplies a thesis-falsification matrix and a worked investment-committee discussion. It separates evidence needed for a new investment from the post-investment remediation process covered in the companion Investor report. It also distinguishes a CVC parent's strategic objective from a financial return thesis, without pretending that the OFF data provides a reliable comparison between those investor types.

Early companies usually lack the operating history needed to demonstrate mature durability. The appropriate response is to match the commitment to what is known, state the unresolved assumptions and identify the next evidence that would change the decision. An early investment can be rational under uncertainty; describing that uncertainty clearly makes the thesis more accountable.

What November adds

The published baseline contained 245 ownership and 237 impact responses. November's reconciled cumulative bases are 373 and 364. The later source has 151 eligible answers to each question, with 23 ownership and 24 impact respondents overlapping the earlier source. Adding source counts without resolving those overlaps would overstate the sample. The statistics registry preserves both source-specific and cumulative results. [INVOWN; INVIMPACT]

The first analytical addition is a separation of the three positive impact categories into distinct investment claims. Better products, lower costs and help for customers can coexist, but evidence for one does not automatically establish the others. A product can become more capable while becoming more expensive to deliver. A cost saving can accrue to customers rather than the supplier. A helpful feature can improve retention without supporting a higher price.

The second addition is a falsification test for each claim. The investment paper should identify an observation that would make the team revise its thesis, rather than list only evidence consistent with enthusiasm. That observation should be measurable at the company's current stage. For an early company, it may be repeat use after a paid pilot or the ability to deliver a result without founder intervention. For a later company, it may concern retention, contribution or the durability of an observed advantage.

The third addition concerns CVC objectives. A corporate investor may value commercial learning, access to a capability or a strategic relationship alongside financial return. Those objectives are legitimate subjects for analysis, but they need their own evidence, cost and conflict assessment. A strategic rationale should not be used to conceal a weak financial case, nor should a financial metric be expected to capture every strategic objective.

The edition does not reuse the October portfolio-triage framework as a new contribution. It focuses on the investment decision and the conditions under which a thesis should be believed. The companion Investor report begins after capital is committed and addresses a company-owned evidence backlog. That separation gives each report a different practical job.

Evidence and sample boundaries

The investor questions ask about respondents' portfolios in general. A respondent may hold several companies, and several respondents may refer to overlapping companies. The unit is therefore the respondent, not the portfolio company, investment or dollar of capital. Neither a company prevalence rate nor an investment return can be calculated from these answers.

The later event source includes registrations from 20 August through 22 October. It is not an October-only cohort. The cumulative latest-answer dataset has different registration ranges because newer eligible answers replace earlier answers for overlapping respondents. Registration time is not a verified answer-edit time, so the data cannot establish when a view changed.

Both questions allow multiple selections. The impact question includes contradictory combinations, which are retained and disclosed rather than silently forced into mutually exclusive groups. The ownership and impact

bases also differ. No unsupported subtraction or cross-question ratio is used to imply that a particular owner produces a particular outcome.

The available fields do not support a clean, validated classification of VC and CVC respondents for comparative estimates. Strategic-investor discussion is therefore conceptual. External studies are used to test the reasoning of an investment claim, not to validate the OFF sample. All investment scenarios and committee tools below are illustrative recommendations, not investment advice about a named security or company.

Exhibit 1. Reported portfolio impact

Selection	Count	Base	Rounded share
Better products	190	364	52%
Cutting costs	141	364	39%
Helping customers	133	364	37%
Nothing measurable yet	51	364	14%

Exhibit note: INVIMPACT, updated cumulative; multiple selections. Positive union 321/364, 88%; eight selected both a positive category and nothing measurable. These are respondent impressions, not audited outcomes.

1. Translate the product claim into a customer choice

A claim that AI makes a product better is incomplete until the investor understands the customer choice it changes. The improvement might reduce effort, enable a previously difficult task or increase the reliability of an existing service. The commercial consequence depends on whether the relevant customer notices and values that improvement, and whether an alternative can provide a comparable result at an acceptable cost.

Diligence should begin with a specific use case and a defined customer population. A demonstration performed by the founder may establish technical possibility. It may not establish that ordinary users can achieve the same result with the available onboarding and support. The investor should distinguish demonstrated capability from repeatable customer success, then identify what evidence connects the two.

Customer behavior is stronger evidence than an undifferentiated expression of interest, but it also needs interpretation. Repeated use can reflect real value, a free trial or a temporary mandate. A paid pilot can reflect experimentation rather than an intention to renew. The investment team should ask what choice the customer made, what alternatives were available and what commitment the customer is willing to repeat under normal terms.

The supplier's ability to capture value is a separate question. A feature may save customers time while competitive pressure passes most of the benefit through in lower prices. A company may still build an attractive business through distribution, service quality or a broader product relationship, but the investment

thesis should name that mechanism. Describing the feature as transformative does not explain the company's share of the resulting value.

The falsification test should follow the mechanism. If the thesis depends on customers switching because of a distinctive workflow, evidence that customers use the feature only as a free add-on would matter. If the thesis depends on difficult implementation knowledge, evidence that competitors reproduce the result with little effort would matter. The committee should identify these possibilities before it becomes invested in defending the original narrative.

This does not require certainty about the final market structure. It requires a clear account of what has been observed and what remains a hypothesis. An early company can have compelling evidence of a valuable problem and limited evidence of durable capture. The investment decision should reflect both, with the next evidence tied to the unresolved part of the thesis rather than another demonstration of what is already known.

2. Follow a cost claim through the operating system

An AI cost thesis should identify whose cost falls and which expense changes. Employee time saved can create capacity without reducing expenditure. A supplier's lower delivery cost can be offset by customer acquisition, review or support. A customer may use the capacity to improve service rather than reduce headcount. Each outcome can be valuable, but they support different financial claims.

The investor should ask for a consistent denominator. Cost per generated item is different from cost per accepted item. Cost per attempted resolution is different from cost per resolved customer issue. An apparent improvement can reverse when the denominator includes rework or a harder workload. The company should preserve the underlying task population and explain changes in case mix before presenting a trend.

Costs outside the model bill deserve attention. Integration maintenance, retrieval, monitoring, human review and exception handling can be economically important. Some are shared platform costs and some are customer-specific. Diligence should identify how management allocates them and whether the chosen method changes the apparent attractiveness of a customer or product. A precise-looking margin is not useful if the cost boundary is incomplete.

The route from operating improvement to financial benefit should be explicit. If the company claims avoided hiring, the investor should understand the workload forecast and the hiring plan that would otherwise have occurred. If it claims lower outsourcing expenditure, the relevant contract and service volume should be considered. If it claims higher capacity, the next question is whether demand exists to use that capacity productively.

External productivity research is useful for identifying possible mechanisms, but it cannot supply the company's missing financial bridge. The same technology can have different effects depending on task difficulty, worker experience and workflow design. The investment team should resist inserting a published average into a model as though it were a company-specific estimate. The relevant evidence is the company's own result under a defined operating scope.

A cost thesis should be falsifiable before it becomes a valuation narrative. Examples include review effort remaining above the level required for the model, savings depending on a temporary subsidy or higher volume creating proportionately more exception work. Those findings do not necessarily make the company unattractive. They require the team to revise the mechanism, price or scale assumption rather than preserve a cost advantage that the evidence no longer supports.

3. Distinguish customer help from durable demand

Helping customers is a broad and useful observation, but it does not identify the business consequence. A product may provide faster answers, more consistent service or assistance with a difficult task. The investor should ask which result matters to the customer and how the company knows that result occurred. A favorable comment can inform product discovery without establishing retention or willingness to pay.

The accepted outcome should be defined in the customer's workflow. A system may generate an answer that is fluent but unusable because it lacks the required evidence or permission. It may complete a task that the customer did not want automated. A measure of customer help should therefore connect the product's output to the customer's decision or completed work, including exceptions and rejected results.

Demand evidence should be segmented by the conditions that affect value. A product can work well for one customer type and poorly for another. Aggregating those groups can obscure both an attractive initial market and a difficult expansion. The investment team should understand which population supports the current claim and what must change for the company to serve a broader market.

Retention is especially informative when it reflects a continued choice under normal commercial terms. It should still be interpreted alongside contract length, implementation effort and switching constraints. A customer may remain because the product is valuable, because replacement is costly or because the agreement has not reached renewal. Those mechanisms have different implications for future pricing, service obligations and competitive resilience.

Customer concentration can strengthen early learning while increasing risk. A demanding anchor customer may help the company develop a valuable capability, yet shape the product around requirements that do not generalize. The investor should ask which parts of delivery are reusable and which depend on that customer's systems or staff. A reference call should explore the customer's remaining work and alternatives, not merely ask whether they like the founder.

The falsification test can be simple: does the claimed value persist when special support, discounts or executive sponsorship are reduced to normal operating conditions? The answer may require time. The investment paper should state that limitation rather than treating initial enthusiasm as durable demand. A company can deserve funding to resolve the question, but the unresolved question should remain visible in the investment case.

4. Test the advantage after the underlying capability improves

A thesis about durability should consider what happens when the underlying technology becomes more capable or cheaper. An improvement in a shared model can help the company deliver better results, but it can also make a previously distinctive feature easier for others to build. The investor should examine both effects instead of treating every model improvement as unambiguously favorable.

The relevant advantage may sit in the workflow, distribution, integration, data rights, customer relationship or operational knowledge. Each proposed source of advantage needs its own evidence. A large volume of data is not automatically useful if the company lacks the rights or quality required for the intended use. An integration is not automatically defensible if customers can replace it easily. A familiar brand is not automatically sufficient if the product's value is weak.

The investment team should ask what the company learns from delivery that a new entrant would struggle to reproduce. The answer might involve domain-specific exceptions, a difficult operating process or feedback tied to accepted outcomes. It should be possible to explain the learning mechanism without claiming that every interaction creates a proprietary moat. Some interactions add little information or cannot lawfully be reused for the proposed purpose.

Dependence on a supplier should be assessed as a business relationship. The company may benefit from access to a powerful platform while remaining exposed to changes in price, terms or product scope. Diligence should identify which assumptions depend on that relationship and what alternatives are credible. A theoretical ability to change models is less informative than evidence about the work required to preserve the product's accepted outcomes.

The committee can use a counterfactual discussion: if a broadly available model became materially better at this task, what would customers still need from this company? The purpose is not to forecast a specific release. It is to test whether the investment thesis depends entirely on a temporary capability gap. A strong answer should identify a customer problem and a delivery mechanism that remain valuable under plausible changes.

The result may be a narrower thesis than the initial pitch. The company might have a compelling advantage in one regulated workflow rather than a general platform opportunity. That can still support an attractive investment at appropriate terms and expectations. Diligence should improve the precision of the thesis, not force every promising company into the broadest possible market narrative.

5. Keep strategic and financial objectives legible

A corporate investor may pursue objectives that differ from a conventional financial fund. It may want access to a capability, learning about a market, a commercial relationship or an option on future collaboration. Those objectives should be stated before evaluating the investment, because they change what evidence matters and who inside the parent organization must participate.

The strategic objective should identify the mechanism of benefit. Access to technology is not useful if no operating team can deploy it. A commercial relationship may require procurement, integration and legal work

that the investment team cannot authorize. The CVC paper should therefore identify the parent sponsor, the operating owner and the commitments needed to realize the proposed strategic value.

Financial and strategic cases should be evaluated separately before they are considered together. A weak financial return expectation should not disappear behind a general claim of strategic relevance. Equally, an objective such as learning may have value that a short-term revenue model does not capture. The committee should see the cost, uncertainty and evidence for each case so that it can make an explicit tradeoff.

Potential conflicts deserve practical treatment. The parent may want exclusivity while the startup needs a broad market. Information rights may create concerns for other customers or investors. A commercial pilot may be slow because of the parent's internal process, even when the startup performs well. These issues should be examined in the actual transaction rather than assumed to be advantages of corporate participation.

The strategic thesis also needs a failure condition. If the parent cannot assign an operating owner, secure the necessary data access or use the capability within a relevant period, the strategic rationale may weaken. If the startup's product direction diverges from the parent's needs, the relationship may need to change. Recording those possibilities helps prevent a strategic investment from becoming an indefinite commitment without an accountable objective.

OFF's instrument does not provide a validated VC-versus-CVC split. This section is therefore a decision framework, not a claim that one investor type behaves differently or produces better outcomes. The common principle is clarity about what the investment is intended to achieve. Different objectives can coexist, but each should remain visible enough that the committee can later judge whether the rationale still holds.

6. Write the disconfirming evidence into the investment paper

An investment paper often becomes more persuasive as supporting examples accumulate. The harder task is to identify which observation would change the team's view. A falsification statement should be specific enough to guide diligence and proportionate to the company's stage. It should address the central mechanism of value rather than a minor operational issue that would not alter the investment decision.

The statement should include an alternative explanation for the current evidence. Strong pilot use might reflect free access, founder involvement or a temporary customer initiative. A reported cost reduction might reflect a change in workload or an omitted expense. The team should identify what evidence would distinguish the preferred explanation from the alternative. This makes the next diligence step more useful than collecting additional positive anecdotes.

Evidence quality should be described without turning the process into a rigid score. A verified customer record, a management estimate and an external benchmark provide different kinds of support. The committee should know which claim depends on each. A missing record can be acceptable in an early company if the investment is explicitly funding the learning required to obtain it; it is more concerning when management presents the claim as already established.

The team should also define the consequence of an adverse finding. Would it change the price, the amount committed, the scope of the thesis or the decision to proceed? An observation that is labelled important but

cannot affect any decision is not a useful diligence condition. The committee should avoid demanding evidence merely because it is easy to request.

After investment, the company needs an operating process for resolving the remaining assumptions. That is the subject of the companion Investor report. The investment paper should hand over a concise record of the thesis, the unresolved questions and the evidence that would change the view. It should not become a permanent narrative that operating teams are expected to confirm regardless of what they learn.

This discipline improves disagreement within the committee. Members can debate the importance of an assumption, the strength of the evidence or the consequence of failure. Those are more productive discussions than competing expressions of confidence in AI as a category. The aim is not to eliminate judgment. It is to make the judgment specific enough that future evidence can improve it.

Exhibit 2. Thesis and falsification matrix

Thesis	Evidence that supports the mechanism	Finding that should prompt revision
Better product	Accepted results and repeated customer choice in a defined use case	Use disappears when special support or free access ends
Lower delivery cost	Complete cost per accepted outcome under normal workload	Review, exceptions or subsidies account for the apparent saving
Better customer result	Relevant customer outcome with a credible comparison	Improvement is confined to an unrepresentative easy subgroup
Durable advantage	Reusable capability, rights or relationship customers continue to value	A shared capability improvement removes the reason to choose the company
CVC strategic value	Named parent owner and a feasible route to the strategic objective	Parent cannot provide the commitments required to realize the benefit

Exhibit note: Proposed diligence framework, not an OFF respondent result or a universal investment scoring model.

External evidence and its limits

The QJE study Generative AI at Work provides evidence from a particular customer-support setting and reports heterogeneous effects. It supports asking which workers and tasks benefit, rather than assuming a uniform effect. It does not establish the economics of an AI startup or the durability of an investment advantage. Those require company-specific evidence. [E-QJE]

METR's February 2026 update explains difficulties in obtaining a reliable current estimate of AI's effect on developer productivity, including selection and measurement issues. That is relevant to diligence because a credible research organization can find that a simple headline estimate is not supported by its newer design. The appropriate lesson is to inspect the measurement, not to repeat an older result as a universal current conclusion. [E-METR]

Stanford HAI's October 2026 article describes research involving 184 employees across three firms and emphasizes the role of framing and organizational support in adoption. The intervention included support such as training and time, so the account should not be read as a clean estimate of wording alone. For an investor, this raises a question about the organizational work required to realize a product's value. [E-STANFORD]

The SEC's March 2024 announcement describes settled charges involving misleading AI-use statements by two investment advisers. Its direct legal scope concerns those adviser cases. It is relevant here as a concrete reason to verify how an AI claim is substantiated, not as proof that every optimistic product statement violates the same rule. [E-SEC-AIWASH]

Taken together, these sources challenge four shortcuts: applying a single productivity average, overlooking changes in measurement, ignoring the organizational conditions of adoption and treating AI language as self-validating. They do not determine which investment to make. The committee still needs to judge the company's customer evidence, delivery economics, competitive position and terms.

Demanding mature evidence would exclude the earliest opportunities. The committee needs to distinguish a measured result from a hypothesis worth funding. The amount and form of commitment should reflect that distinction, together with the investor's mandate and risk tolerance. A thesis can be ambitious while remaining precise about the evidence it lacks.

The committee can make this source comparison useful by assigning a different diligence request to each uncertainty. For a productivity claim, request task and worker coverage together with the comparison design. For an adoption claim, ask what support the customer had to supply and whether that burden recurs at renewal. For an advantage claim, identify the alternative available to the customer if underlying models improve. For a public AI statement, obtain the operating facts supporting the description. These requests do not guarantee a successful investment. They expose the assumptions that the investor is choosing to finance and allow later evidence to challenge the original thesis without rewriting its meaning after the event.

Evidence comparison

Source	Contribution to this decision	Limit on interpretation
QJE study [E-QJE]	Performance varies across workers and tasks	One support setting; not startup investment performance
METR update [E-METR]	Selection and measurement affect estimates	No reliable new universal productivity estimate asserted
Stanford HAI [E-STANFORD]	Adoption involves support as well as framing	184 employees, three firms; bundled intervention
SEC cases [E-SEC-AIWASH]	AI-use descriptions require substantiation	Specific adviser settlements; not a finding about OFF companies

An illustrative investment-committee review

Consider a hypothetical company selling an AI service that prepares compliance-related customer responses for business software providers. Management's pitch combines three claims: the product improves response quality, reduces delivery effort and creates a durable advantage through workflow knowledge. The committee should separate those claims before deciding whether the available evidence supports the proposed investment.

For product quality, the team asks how responses are accepted and which cases were included in the pilot. It discovers that experienced customer staff reviewed every response and that complex exceptions were excluded. This can still be useful evidence of an assisted workflow. It does not yet support a claim that the product can independently handle the full workload. The next test should address the intended operating scope rather than repeat the same easy demonstration.

For cost, the team asks whether the model includes customer review and the supplier's support effort. If those costs are omitted, the apparent saving may represent transferred work. The company should show a complete delivery view for the pilot and identify which costs are expected to fall through reusable product development. The committee can then distinguish a measured current result from a forecast that depends on future engineering.

For durability, the team asks what remains valuable if a broadly available model improves at drafting similar responses. The company points to customer-specific integrations and a record of accepted exceptions. Diligence should examine whether those assets are reusable, whether the company has the necessary rights and whether customers value the resulting service enough to renew. The answer may support a focused workflow business even if the initial general-platform claim is too broad.

Exhibit 3. Illustrative committee decision record

Question	Current evidence	Next decision-relevant evidence
Does the product improve accepted responses?	Restricted pilot with experienced review	Defined broader workload and acceptance results
Does delivery create economic benefit?	Initial model omits some review and support	Reconciled cost per accepted response
Is the advantage durable?	Integrations and claimed exception knowledge	Reuse, rights and customer renewal evidence
Can a corporate investor realize strategic value?	Parent interest without an operating commitment	Named owner, agreed pilot scope and feasible integration route
What changes the investment view?	Thesis depends on repeatable paid delivery	Evidence that normal delivery cannot meet the required economics

Exhibit note: Hypothetical company and proposed decision record. No valuation, investment return or observed portfolio performance is asserted.

If a corporate investor participates, its strategic case should be considered alongside these financial questions. The parent may want to use the service internally, but interest from an executive is not the same as an authorized deployment. The committee should identify the operating team, data access and procurement commitments required. If those are unavailable, the strategic value should be treated as uncertain rather than included as an automatic benefit of the investment.

The final paper should state which parts of the thesis are supported, which remain hypotheses and what would cause the team to revise its view. It may recommend proceeding because the opportunity justifies funding the next stage of learning. It may recommend a narrower commitment or further diligence. The useful outcome is a decision connected to the actual evidence, with a clear handoff of unresolved assumptions.

The next investment discussion

Take one active AI opportunity and separate its product, cost and customer claims. For each, identify the relevant population, the observed result, the alternative explanation and the evidence that could change the investment view. If the same favorable anecdote is carrying all three claims, the paper probably needs more precise reasoning before it needs more slides.

Use the updated OFF results as a source of questions about what investors are seeing, not as a substitute for company diligence. The reported impact categories are broad, overlapping and self-reported. They are useful because they reveal several mechanisms worth examining. They do not establish returns, valuation effects or a preferred investor strategy.

November's recommendation is to make the thesis accountable to evidence that can arrive after the decision. A clear falsification condition does not weaken conviction. It tells the company and the investor what must be learned, what deserves continued attention and which finding would justify changing course. That is a practical basis for underwriting an uncertain market without turning adoption into proof of durable value.

Investment questions

Does 88% positive impact mean 88% of portfolio companies earned a return?

No. It is 321 of 364 respondents selecting at least one positive category. The unit is the respondent, not the company or investment. [INV_POSITIVE]

Why is the positive share not 100% minus 14%?

Eight respondents selected both a positive category and nothing measurable yet. The positive union is calculated directly. [INVIMPACT]

Are November changes a panel trend?

No. The cumulative sample adds a later event cohort and resolves overlap. It does not track the same companies through time.

Does CEO ownership cause better outcomes?

The data does not establish that relationship. Ownership and impact are separate self-reports with different bases. [INVOWN; INVIMPACT]

Does the report compare VC and CVC performance?

No. The available instrument does not support a validated respondent-type comparison.

Must early companies prove mature durability before investment?

No. The recommendation is to state what is known, what remains a hypothesis and what evidence would change the decision.

Citable facts

- Better products: 190/364, 52%; cutting costs: 141/364, 39%; helping customers: 133/364, 37%. [INVIMPACT]

- CEO ownership selected: 185/373, 50%; multiple selections permitted. [INVOWN]

- Positive impact union: 321/364, 88%; eight contradictory positive-plus-none combinations retained. [INV_POSITIVE]

Methods and limitations

Calculation-v1.1 excludes byte-identical exports and invited-only records and selects the latest eligible nonblank answer per normalized respondent identifier per question. The published October-compatible baseline is reproduced. Ownership reconciles as 245 plus 151 minus 23 equals 373; impact as 237 plus 151 minus 24 equals 364. Counts within categories also reflect replacement of overlapping respondents' answers.

The selective registration sample is self-reported and nonrepresentative. Portfolio-company overlap is unknown. No investment return, causal relationship, valuation estimate or VC/CVC subgroup difference is claimed. The later event source is not an October-only cohort. Investment tools and the worked company are author recommendations and illustrations.

References

[OFF-DATA] Open Future Forum. Executive AI Statistics, November 2026. INVOWN, INVIMPACT and INV_POSITIVE.

[E-QJE] Brynjolfsson, E., Li, D. and Raymond, L. Generative AI at Work. Quarterly Journal of Economics, 140(2), 889-942, 2025. https://academic.oup.com/qje/article/140/2/889/7990658

[E-METR] METR. We are Changing our Developer Productivity Experiment Design. 24 February 2026. https://metr.org/blog/2026-02-24-uplift-update/

[E-STANFORD] Stanford HAI. Want Employees to Embrace AI? Stop Selling It as a Productivity Tool. 7 October 2026. https://hai.stanford.edu/news/want-employees-to-embrace-ai-stop-selling-it-as-a-productivity-tool

[E-SEC-AIWASH] US Securities and Exchange Commission. SEC Charges Two Investment Advisers with Making False and Misleading Statements About Their Use of Artificial Intelligence. 18 March 2024. https://www.sec.gov/newsroom/press-releases/2024-36

Related research and citation

The November Investor report covers post-investment evidence completion, YC Founders addresses pricing and buyer authority, CFO qualifies financial benefits, and Private Equity examines transferability through a transaction.

Suggested citation: Newlands, Murray. VC and CVC AI Investment Report, November 2026, Edition 4. Open Future Forum. Evidence cutoff 31 October 2026.

Open Future Forum research: https://openfutureforum.com/research/

Editable working tool

Investment thesis and falsification: download the editable CSV worksheet. The worksheet implements the report's recommendation; it is not a validated score, mandatory standard or completed assessment.