Climate Sprints

FOR FOUNDERS · AND FOR THE PEOPLE WHO BACKED THEM FIRST

What capital is actually scoring.

The twenty-three things an allocator assesses before they commit, what evidence satisfies each, and what your answer tells them when you do not have it. The companion to our fund-facing method — same questions, read from the other side of the table.

VERSION 1.2
3 SEPTEMBER 2026
REVIEW 1Q2027

CLIMATESPRINTS.COM
PRACTICE@CLIMATESPRINTS.COM
POWERED BY PURPOSE. DRIVEN BY ACTION.

01

Why this exists

We publish our assessment method so the funds who hire us can argue with it. This is the same method, rewritten for the people being assessed.

There is no advantage to us in keeping it private. A company that has organized itself against these questions is a better company and a faster engagement. A company that has not will spend three weeks discovering, expensively and in front of an investor, what it could have discovered on its own.

What already exists, and what it does not cover

There is good published guidance on raising capital. Y Combinator's Series A guide, built from work with 190 companies, covers the pitch, the memo, the metrics and the diligence materials.[1] Its due diligence checklist is widely used as a preparation roadmap.[2] Sequoia's deck template and Michael Seibel's pitch framework are the standard references, and the practical advice in them is sound — show your market math, ask directly, and prefer clarity to polish.[3] DocSend's much-quoted finding is that investors spend about three minutes and forty-four seconds on a deck.[4]

Almost all of it is shaped by software companies, and almost all of it is about the *pitch*.

None of it tells a climate founder that a named buyer with a dated trigger is worth more than a market-size slide. None of it explains that an allocator is testing whether your cost case survives the expiry of a credit that is already legislated. None of it says that the way you answer a question about a disagreement between founders is itself being scored.

THE ONE-SENTENCE VERSION

Most fundraising advice teaches you to be *persuasive*. This document is about being *assessable*. They are not the same skill, and after the first meeting only the second one matters.

Who should read it

Founders and CxOs of climate, cleantech, regenerative and circular companies between a technology that works and a first commercial deployment — and developers approaching a first-of-a-kind decision.

Angel and seed investors already holding a position. Section 3 is written for you specifically, and it is the section most likely to be uncomfortable.

Board members who need a structure for a conversation they have been avoiding.

02

What is actually being tested

Almost every climate company that failed commercially had technology that worked.

What was misjudged was not feasibility. It was whether capital, regulation and buyer urgency would arrive at the same time. A company that is early does not fail gracefully. It runs out of money while its thesis is still correct.

> *The question an allocator is answering is not whether your technology works. It is who is compelled to buy it, and when they are compelled.*

That produces four assessments rather than one, and most founders prepare for only the first.

ASSESSMENT	THE QUESTION BEHIND IT	WHAT FOUNDERS USUALLY PREPARE
Technical readiness	Does it work at the stage you claim, on evidence rather than assertion?	Thoroughly. This is usually the strongest part of the data room.
Commercial compulsion	Is a named buyer compelled to act, and on what date?	Market size, a pipeline, letters of interest. Rarely a dated trigger.
Execution capacity	Can this team do what the deck describes?	Bios. Almost never the thing actually being assessed.
Capital readiness	Can you absorb and deploy this money, and does the runway reach something fundable?	The current round. Rarely the total to breakeven or the dilution path.

Twenty-three dimensions sit under those four. Sections 4 to 7 walk each of them.

Two rules that change how you should prepare

Evidence is graded, not just counted. Every score carries a letter describing where the evidence came from — independent third party, counterparty, company-verifiable, or assertion. A dimension supported only by your own statement is **capped** and cannot score above the midpoint however plausible it is. One customer-run trial moves a score further than ten pages of internal data.

Absence of evidence is a finding, not a blank. When you cannot produce what a dimension asks for, that is recorded as what it is. A team that cannot name the budget line the purchase comes from has answered the question.

03

For the angel or seed investor

You were first in. You backed a person before there was much else to back, and that was the correct basis for the decision, because at that stage there is nothing else. This section is about what happens to that basis over time — and about a conversation that almost never takes place.

There is an old line in advertising: people buy on emotion and justify on logic. It is treated as a truth about consumers. It is at least as true about early-stage capital, and it is almost never said out loud.

Belief is the instrument at seed, and it is not neutral

The research on angel decision-making is unusually clear. In a study of 72 angels completing 1,995 evaluations across 133 live pitches, founder-displayed enthusiasm, preparedness and commitment all drove evaluations of funding potential, interacting in complicated ways.[5] One finding is worth sitting with: angels sometimes read *high enthusiasm combined with heavy personal commitment cues* as inauthentic rather than reassuring.[5]

How much weight any individual angel places on displayed passion varies with their own affective reactivity — a property of the investor as much as of the founder.[6] And angels do not, in practice, use a fully compensatory model that weights and scores many attributes; across 150 recorded entrepreneur-investor interactions the decision looked much more like elimination on a few factors.[7]

None of this is a criticism. It is a rational response to a situation with almost no data in it. But it means the seed decision rests on impressions formed early, and impressions are stickier than evidence.

The conversation that does not happen

An angel decision is frequently made without the financial advisor, without the spouse, and without a written thesis. It is recorded, if at all, as instinct — the butterflies in the stomach that accompany a decision someone wants to make.

That instinct is often legitimate at first blush. What it is not is *defensible* eighteen months later, when milestones have slipped twice and the returns are not where they were promised. At that point there is no written record of what the decision actually rested on, which makes it impossible to test whether the original reasoning still holds or has quietly been replaced by hope.

The industry cites gut feel constantly, and only ever in the positive. Nobody publishes the instinct that cost them. Our reading of why is straightforward: in a field where competence is the currency and the last twenty-five years produced abundant capital and abundant exits, being visibly wrong about money carries a stigma that being wrong about almost anything else does not.

The behavioral literature says the same thing more precisely. Failure brings emotional pain and shame, and shame is what prevents the learning from being extracted.[15] So the losses stay private, the lessons do not circulate, and every new angel starts from the same place — which is exactly the condition under which survivorship bias does its work.

FROM THE FOUNDING PRINCIPAL

I made this mistake, and I can be specific about it.

In 1999 I took a substantial position in industrial-scale battery storage. I had believed in the solar economy since I was a teenager and I understood — correctly — that grid-scale storage was the threshold solar had to cross before its economics worked. The technology thesis was right. It is right today.

I met the principals. I saw a demonstration. I read the business plan and the pro forma. By the standards I had then, I did the work.

What I could not do was assess any of it.

On the technology. I watched a demonstration and had no way to evaluate how the system was actually engineered, or whether what I was shown was a working system or a staged result. I could not distinguish the stage claimed from the stage evidenced. Nobody independent had looked, and it did not occur to me that anybody should.

On the market. I read the plan and the projections and I saw what I wanted to see. My own conviction about where the market had to go supplied the demand those numbers assumed. I was not testing the commercial case. I was confirming it.

On the people. I was not a competent judge of the founding team's character or their capacity to deliver. Nor was I a competent judge of the person who brought me the deal — whose position, as it turned out, my money largely retired. There were material things I was not told. I did not see them because I was not looking. I was too invested in being right about the technology and the timing to examine carefully who I was actually transacting with, or why the opportunity had come to me at that moment.

It was a calculated gamble on belief, without instrumentation. It failed completely.

I lost a great deal of money. The larger cost was the reserve it drained and the positions I could not take for years afterward. The technology was not the mistake. The absence of a method was.

Marc Strauch

Founding principal, Climate Sprints

Three failures, and the assessments that exist because of them

That account is not included for candor. It is included because the three failures in it map exactly onto three of the four assessments in this document, and the fourth was created after we noticed what was still missing.

WHAT WENT WRONG	ASSESSMENT	WHAT WOULD HAVE CAUGHT IT
A demonstration accepted as evidence of engineering	Technical readiness	TR-1 and TR-5. A demo is a company-generated artifact — evidence grade C at best, and often D. The requirement for independent verification exists because a demonstration is designed to be watched, not tested.
Projections read as analysis when they were confirmation	Commercial compulsion	CC-1 and CC-3. A named buyer with a dated trigger cannot be supplied by the reader's own conviction. That is the entire point of requiring a name and a date rather than a market size.
No basis for judging the team, or the counterparty	Execution capacity	EC-1 verified by reference rather than by resume, EC-6 on the decision record, EC-8 on whether prior forecasts matched outcomes. And the provenance question below, which none of the four assessments originally asked.

The question nobody scores: where did this deal come from

The failure above has a fourth element that does not sit inside any assessment of the *company*, because it is not about the company.

Someone brought that deal. Their position was exited into the capital that came in behind them. That is not necessarily improper — secondary sales happen for legitimate reasons and early holders are entitled to liquidity. But it is material information about what is being transacted, and it was not visible from anything in the data room.

> *Diligence is pointed at the company. It is almost never pointed at the transaction, or at who benefits from it closing.*

Three questions worth asking before any position, and worth asking out loud rather than inferring:

- **Who introduced this, and what is their position in it?** Are they investing alongside you on the same terms, holding and not participating, or exiting?
- **Is any part of this round secondary?** How much of the money reaches the company's balance sheet, and how much retires an existing holder? A round that is substantially secondary is a different instrument from one that is primary, and should be priced as one.
- **Why now, and why me?** Not cynically. But an opportunity that arrives because you are known to hold a particular conviction has been selected for your belief — and that is worth noticing at the time rather than afterward.

These questions cost nothing and are almost never asked, because asking them implies a distrust that feels incompatible with the relationship that produced the introduction. That discomfort is precisely why they go unasked, and precisely why they matter.

What the concentration data says

The single largest study of angel outcomes remains Wiltbank and Boeker's work for the Kauffman Foundation — 539 angels, 3,097 investments, 1,137 exits.[16] Two findings from it matter here, and they compound.

Diligence hours track returns, sharply

Split at the median of twenty hours of due diligence per deal: angels below the median returned **1.1x over 3.4 years**, roughly a 2.2% IRR. Angels above it returned **5.9x over 4.1 years**, roughly 43% IRR. Those spending forty hours or more averaged **7.1x**.[16][17] Below-median diligence produced exits returning less than 1x in 65% of cases; above median, 45%.[18]

The honest caveat: this is correlation from self-reported data, and the causal direction has been questioned — angels who do more work may differ in other ways too.[19] But the association is large, it is consistent across studies, and it is one of very few controllable variables in the asset class.

Portfolio size is not optional either

More than half of individual angel investments return less than 1x. Across angels holding multiple positions that figure falls below 40%.[20] The Kauffman data found materially better returns for portfolios of twenty or more companies.[21] At Tech Coast Angels, four outcomes out of 247 — 1.6% — exceeded 100x, and those four drove the group's entire return.[22]

> *A concentrated position, decided quickly, is the intersection of the two worst quadrants in the data — and it is the most common shape of a first angel investment.*

Opportunity cost, which nobody counts

The loss on a position is the visible cost. It is frequently not the largest one.

An angel writing between $25,000 and $500,000 is deploying from finite reserves. An over-allocated position does not just lose its own capital — it removes the capacity to take the next twenty positions, which is precisely the number the return math requires. The power law only pays if you are still holding chips when the outlier appears.

So the real cost of a concentrated early loss is three things at once: the capital, the positions never taken, and the years spent recovering the reserve before you could take them.

None of this appears in a portfolio review, because the positions you could not take generate no records.

Escalation, and why willpower does not fix it

The literature has a name for what follows a bad early position: escalation of commitment — continuing to allocate resources in the face of negative feedback.[8] Work on persistence in under-performing firms found extrinsic motivation shaping whether a founder keeps going in ways the venture's prospects do not explain, and titled the finding a fallacy: the assumption that only the strong survive.[9]

Two mechanisms make this worse in early-stage technology specifically.

Complexity increases confidence. Introducing a technologically advanced or resource-intensive product tends to *increase* overconfidence in its future success, and overconfidence predicts escalation better when results are deteriorating than when they are improving.[23][8] A hard technology is therefore more likely, not less, to attract capital past the point at which the evidence supports it.

Experience does not reliably help. More than 80% of founders overestimate their probability of success, and cognitive biases are not reliably reduced by experience — in habitual entrepreneurs they can be strengthened.[24] The same finding is available in the parallel literature comparing venture behavior to gambling, where entrepreneurship is described as structurally similar given its uncertainty, inherent risk and high failure rate.[23]

And the finding that matters most for anyone assessing their own position:

STAW AND ROSS, 1989

Escalation of commitment is most pronounced when **the same person who made the original decision is the one evaluating whether to continue**. Personal accountability for the original call creates an ego investment that compounds the financial sunk cost, and makes objective reassessment structurally difficult.[25]

Not difficult because of weak discipline. Difficult because of who is doing the evaluating.

Which is why self-correction is unreliable. Cognitive biases operate below conscious awareness; willpower and good intentions do not address them, and decision architecture does.[24] Neuroimaging work locates sunk-cost processing in emotional circuits that can override rational evaluation, which is a physiological account of why knowing about the bias does not dissolve it.[26]

Duration of conviction

One risk factor deserves naming separately, because it is invisible from inside.

A thesis held for a long time is harder to test than one formed last year. When a belief has been part of how you understand the world for a decade or more, disconfirming evidence does not arrive as information. It arrives as a challenge to identity, and it is processed accordingly.

This is not a reason to abandon long-held convictions. Long-held convictions about technology are frequently correct — that is the whole difficulty. It is a reason to be specific about what would change your mind, and to have written it down before the evidence arrives.

The reality check

Not a recommendation to sell. Most seed positions should be held, and being early is not the same as being wrong.

But a position taken on belief should be re-examined on evidence at least once, and the research is clear that you are the wrong person to do it alone.

Score the company against the twenty-three dimensions in Sections 4 to 7. Not as the founder would answer. As you can evidence. Then do the same exercise with someone who has no position in the outcome — a co-investor, an advisor, a spouse. The point of the second person is not their expertise. It is that Staw and Ross do not apply to them.

Four things to notice.

- **How many dimensions you cannot answer at all.** Not answer badly — cannot answer. That count is the most useful number in the exercise, and for most seed positions it is higher than the holder expects.
- **Whether your evidence is grade D.** If your confidence rests on what the founder told you, it rests on assertion. Normal at seed, not disqualifying, worth knowing.
- **Whether CC-3 has a date.** If the trigger that makes a buyer act cannot be dated, you hold a timing bet. Timing bets can be excellent positions. They should be sized as timing bets, and most seed positions were not.
- **What this position is costing you in capacity.** Not what it has lost — what it prevents you from doing next.

WHY AN ANGEL SHOULD WANT THEIR FOUNDER TO READ THIS

The next investor will run some version of this assessment whether or not anyone has prepared for it. A founder who has already found the gaps can address them before a term sheet depends on it.

The uncomfortable version: if a company cannot answer these questions, the seed investor will discover that at the same moment as the Series A lead — in the room, without warning, with no time to fix anything.

04

Technical readiness

Whether it works at the stage claimed rather than the stage hoped for. Usually a founder's strongest ground — and the place where a small number of specific gaps do the most damage.

DIMENSION	WHAT SATISFIES IT	WHAT YOUR ANSWER REVEALS
TR-1 Stage integrity	Third-party test data at the stage you claim, under representative rather than ideal conditions.	A stage given as a range, or a readiness number with no artifact behind it, reads as a claim rather than a fact. It caps the score regardless of how good the technology is.
TR-2 Scale-up discontinuity	Named engineering discontinuities between bench, pilot and first unit, each with a stated approach.	Describing scale-up as "more of the same" is the single most common technical tell. Assessors read it as either not knowing what breaks or not wanting to say.
TR-3 Dependency exposure	Feedstock and input agreements, permits with dates, interconnection position, second sources where a dependency is critical.	Calling a single-sourced component a commodity is checkable in an afternoon, and getting caught on it damages every other technical claim.
TR-4 Performance durability	Degradation and uptime from operating hours, not from models.	Life claims supported only by accelerated testing are scored as models. If you have field hours, they are worth more than you think; if you do not, say so plainly.
TR-5 Independent verification	Certification, customer-run trials, published review, prior diligence by an unrelated party.	Advisors, investors and paid consultants are not independent verification. Naming them as such is read as not understanding the distinction.

The highest-leverage preparation in this section is a single piece of third-party verification. One customer-run trial, one certification, one independent test report moves TR-1 and TR-5 together and lifts the evidence grade on everything adjacent to them.

05

Commercial compulsion

Not whether a market exists. Whether a specific buyer has to act, and when. This is where most climate companies lose the room, and it is almost never because the market is small.

DIMENSION	WHAT SATISFIES IT	WHAT YOUR ANSWER REVEALS
CC-1 Named buyer	A person, with a title, at a named organization, who would sign. Ideally an LOI, paid pilot or contract.	A logo wall is not a buyer. A job function without a name is not a buyer. If the answer arrives as a category, the assessor concludes nobody has been asked.
CC-2 Budget line	The budget this comes out of, who approves at that size, and evidence they have spent against it before.	"They'll find the money" is scored as a two. Using capex and opex interchangeably signals you have not had the procurement conversation.
CC-3 Trigger	What happens to that buyer if they do nothing, **with a date**. A compliance deadline, a replacement cycle, a contract expiry, a threshold already crossed.	This is the dimension the whole assessment turns on. A trigger described in terms of the buyer's stated ambition or values is a preference, not a trigger.
CC-4 Unsubsidized cost case	Your cost stack with and without each credit, grant or mandate, tested against the expiry dates actually legislated.	Not knowing the expiry dates is the tell. If the case fails unsubsidized, say so and name the volume or input price at which it stops failing — that answer scores far better than a claim that it holds.
CC-5 Displacement friction	The incumbent, its remaining book value, the switching cost including requalification and downtime, and a measured procurement cycle.	Saying the incumbent is "nothing" or "spreadsheets" is read as not having looked. Every buyer is doing something today, and stopping costs them.

IF YOU FIX ONE THING BEFORE YOUR NEXT RAISE

Date the trigger. Find the one buyer whose obligation, replacement cycle or threshold puts a date on the decision, and be able to say what that date is and where it comes from.

A company that can do this converts from a timing bet into a position. A company that cannot is asking an investor to fund a wait of unknown length, whatever else is true about the technology.

06

Execution capacity

Whether this team can do what the deck describes. Eight dimensions, and the one section where *how* you answer is scored alongside what you answer.

This is the part founders find hardest to read, so it is worth being direct about the premise. It is not a judgment about anyone's ability. It is an observation about a role change.

The orientation problem

A research organization and a commercial organization are not the same organization running a different play. They select for opposite temperaments.

Inward-facing work rewards depth, patience with ambiguity, tolerance for repeated failure against a fixed question, and a high threshold for calling something finished. Outward-facing work rewards speed, comfort with rejection, willingness to put something imperfect in front of a buyer, and a low threshold for calling something good enough to sell.

A founder who is excellent at the first is not thereby competent at the second. The belief that the transition is a matter of effort rather than of capability is the most common form this failure takes — and the founders most confident that they have already made the turn are frequently the ones who have not started.

It is addressable. By a hire, by fractional support, or by the founder doing real work on it. What is not addressable is a founder who will not look.

DIMENSION	WHAT SATISFIES IT	WHAT YOUR ANSWER REVEALS
EC-1 Demonstrated execution	What you have personally delivered, at what scale, verified by reference rather than by resume.	Presenting advisory and operating experience as equivalent is checked in the first reference call.
EC-2 Orientation shift	Who owns the buyer relationship and what they have closed. The senior team's actual time allocation last quarter. Commercial hires with start dates.	If the founder still owns every buyer conversation by preference rather than by necessity, the shift has not happened.
EC-3 Blind-spot permeability	Naming what the company cannot do, unprompted. Evidence of a belief revised under challenge.	Scored on the response itself. Reframing gaps as strengths, or defensiveness when a gap is named, scores lower than the gap would have.
EC-4 Conflict metabolism	A stated mechanism for who decides when founders disagree. One traced example of a real disagreement and how it ended.	"We've never disagreed" is the highest-signal answer in the entire assessment, and not in your favor. It reads as avoidance or as a dispute nobody will name.
EC-5 Cohesion and hiring integrity	What the last three hires were selected against. Voluntary departures over eighteen months, with reasons.	Hiring for capability while treating fit as a formality shows up two years later as a team that cannot decide together.
EC-6 Decision behavior	Traced examples of expensive decisions — what was decided, how fast, who dissented, what happened. A decision reversed on new information.	A decision record in which every decision was correct is not read as a strong record. It is read as an incomplete one.
EC-7 Governance permeability	Whether bad news reaches the board before it reaches the numbers. Whether prior investors were used as a resource.	A board described as a reporting obligation rather than a resource tells an incoming investor what their own seat will be worth.
EC-8 Disclosure integrity	The forecast you gave your last investor, against what actually happened.	Unavailable prior forecasts, or forecasts revised after the fact, are the fastest way to lose an assessment that was otherwise going well.

What you cannot prepare for, and should not try to

Some of this is assessed under pressure — a challenged assumption, a compressed decision, a question the team has not rehearsed. What is being observed is not the answer. It is who speaks, who defers, whether disagreement surfaces or is suppressed, and whether the founder can be told they are wrong in front of their team.

Rehearsing this makes it worse, and experienced assessors are looking for the rehearsal. The only real preparation is to have done the underlying work: resolve the dispute, name the gap, make the hire, tell the board the bad thing early.

07

Capital readiness

Whether you can absorb and deploy what you are asking for, over how many rounds, at what dilution, toward what liquidity event. The most conventional of the four and the one most often prepared only one round deep.

It matters more now than it did three years ago. Capital has concentrated sharply — in Q4 2025, roughly 60% of venture funding went to the top ten percent of rounds, and the recipients were disproportionately capital-efficient.[10] Inefficient companies now run out of runway *before* the next round rather than after it.

DIMENSION	WHAT SATISFIES IT	WHAT YOUR ANSWER REVEALS
CR-1 Absorptive capacity	Your actual time-to-fill on the last five hires, against the headcount ramp the plan assumes.	A plan assuming hiring velocity you have never achieved is the most common capital finding and the easiest to check.
CR-2 Capital plan integrity	Total capital to breakeven, built bottom-up. Each round tied to a milestone rather than a period. Non-dilutive sources identified and pursued.	Stating the requirement as a round size rather than a total says you have not modeled past the next raise.
CR-3 Runway and milestone	Months of runway, the specific milestone it reaches, and why a later-stage investor treats that milestone as de-risking. Eighteen to twenty-four months to the next raise is the working benchmark.[11]	Runway expressed in months to a date rather than to a milestone is the tell. Investors fund milestones, not calendars.
CR-4 Dilution and cap table	Modeled dilution across all planned rounds. Founder and employee ownership at exit. Preference stack. Option pool adequacy.	Unmodeled dilution beyond the current round means nobody has checked whether the team is still motivated at the end of the plan.
CR-5 Liquidity horizon	Named acquirers who have bought comparable assets, with dates and multiples. Time to that event against the fund's remaining life.	Acquirers named as a category, or comparables more than three years old, reads as a route nobody has tested.

THE COMPARISON THAT CATCHES GOOD COMPANIES

Put the date from CC-3 next to the date your runway ends in CR-3.

If the buyer becomes compelled in thirty months and your cash runs out in eleven, you do not have a commercial problem. You have a financing problem, and every investor who runs this assessment will find it in the first week. Finding it yourself first is the difference between a plan and a surprise.

08

What automation changed

Diligence has automated fast and unevenly, and it has moved where you are actually being judged.

Adoption is no longer emerging. A 2026 survey of nearly three hundred private capital dealmakers found 85% now using AI in their workflow.[12] Automated document and contract review is credited with accelerating that work by 70–80%.[13] Practitioner accounts put the data-intensive share of a diligence team's time at 60 to 70% — and that is the share being compressed.[14]

Two consequences for you.

Your documents will be read completely, and quickly

The era in which an inconsistency between the deck, the model and the data room might go unnoticed is over. Every number will be reconciled against every other number, and discrepancies surface in hours rather than weeks. Internal consistency is no longer diligence hygiene; it is a precondition.

The judgment has moved to the part no tool can read

When most investors run comparable tools over the same data room, they reach comparable findings. The technical and financial picture is becoming a commodity input. What differentiates one investor's decision from another's is increasingly the assessment no software performs: how your team behaves, decides and discloses.

> *The parts of your company that leave no document trail are now the parts carrying the most weight. A deferred decision produces no artifact. An unresolved dispute appears in no board pack. That is what unresolved means.*

Which is why Section 6 is the longest in this document, and why the preparation it asks for cannot be done in the week before a raise.

09

How to prepare

In order, because the sequence matters more than the effort.

	DO THIS	WHY FIRST
1	Score yourself against all twenty-three dimensions, honestly, with grades.	The count of dimensions you cannot answer is the real output. Most teams are surprised, and the surprise is the useful part.
2	Date the CC-3 trigger, or establish that you cannot.	Highest-leverage single item. It determines whether you are a position or a timing bet, and everything else is read through it.
3	Put CC-3 and CR-3 side by side.	If compulsion arrives after the cash runs out, that changes what you are raising and how much. Better to know now.
4	Get one piece of grade A evidence.	A single independent verification lifts a whole assessment. Usually the cheapest score improvement available.
5	Name your gaps in writing, and what you are doing about each.	EC-3 is scored on this. A named gap with an action beats an unnamed gap every time, and beats a denied gap by a wide margin.
6	Resolve the founder disagreement you have been managing around.	It will surface under pressure. Sooner is cheaper, and the resolution itself becomes evidence for EC-4.
7	Reconcile the deck, the model and the data room.	Automated review finds every discrepancy. Fix them before they are found.
8	Model dilution to exit and look at the number.	If the team is not meaningfully held at the end of the plan, an investor will ask why they should expect you to stay.

Items 1, 5 and 6 cost nothing but honesty and are the ones most often skipped. Items 2 and 4 cost time and money and are worth both.

10

What this document is not

Not a promise. Preparing against these dimensions improves how you are assessed. It does not guarantee an outcome, and no honest document about fundraising could.

Not a script. The execution dimensions are scored on behavior under conditions you cannot control. Rehearsed answers are a known tell and are looked for.

Not investment advice. Climate Sprints is not a registered investment adviser or broker-dealer. Nothing here is a recommendation to any investor to buy, sell or hold any position, or to any founder to raise or not raise.

Not every investor's method. This is ours, published so it can be argued with. Other allocators weight differently and some will not assess execution capacity at all. We think that is a mistake; you should not assume it is universal.

Not a substitute for advisors. Legal, tax, technical and financial specialists do work this does not replace.

WHY WE PUBLISH THIS

We are hired by capital, not by companies. That is stated plainly on our site and it is worth stating here.

We publish the founder-facing version because a prepared company is a better company and a better engagement, because the questions are better questions when everyone can see them, and because a method that only works by surprise is not a method.

11

References

Publisher, title, date and URL. Accessed 3 September 2026. The full fund-facing method, with its own reference set, is published at climatesprints.com.

[1] Y Combinator Series A Guide — 70-page guide built from work with 190 YC companies, covering the investor pitch, investment memos, valuations and metrics. Summarized via TEH CPA.
https://tehcpa.net/tax-and-finance-due-diligence-checklist-to-prepare-for-raising-vc-investment-funding-and-series-a-rounds/

[2] Founders CPA. *Due Diligence Checklist for Startups* — built on Y Combinator's Series A Due Diligence Checklist. 13 Nov 2025.
https://founderscpa.com/blog/due-diligence-checklist-for-startups

[3] SaaStr. *The Y Combinator Guide to Perfectly Pitching Your Seed Stage Startup, with Michael Seibel* — on transparent market math, asking directly with 18–24 month milestones, and clarity over polish. 9 Mar 2025.
https://www.saastr.com/the-y-combinator-guide-to-perfectly-pitching-your-seed-stage-startup-with-michael-seibel/

[4] Deckary. *Pitch Deck Template: The 12-Slide Framework That Gets Funded* — citing DocSend research that investors spend an average of 3 minutes 44 seconds reviewing a deck; combines Y Combinator and Sequoia practice. 7 Oct 2025.
https://deckary.com/blog/pitch-deck-template

[5] Cardon, M. S., Mitteness, C. and Sudek, R. *Motivational Cues and Angel Investing: Interactions Among Enthusiasm, Preparedness, and Commitment.* Entrepreneurship Theory and Practice, 2017. Empirical study of 72 angels completing 1,995 evaluations across 133 live pitches.
https://journals.sagepub.com/doi/abs/10.1111/etap.12255

[6] Mitteness, C., Sudek, R. and Cardon, M. S. *Angel Investor Characteristics That Determine Whether Perceived Passion Leads to Higher Evaluations of Funding Potential* — on affective reactivity as an investor-side variable.
https://github.com/francisjervis/venture-capital-research/blob/master/angels.md

[7] Maxwell, A. L., Jeffrey, S. A. and Lévesque, M. *Business angel early stage decision making* — 150 entrepreneur-investor interactions; angels found not to use a fully compensatory decision model. Cited via [5] and the venture-capital-research collection.
https://github.com/francisjervis/venture-capital-research/blob/master/angels.md

[8] McCarthy, A. D., Schoorman, F. D. and Cooper, A. C. (1993) — escalation of commitment in new ventures, continuing to allocate funds despite negative feedback. Cited in [5].

[9] DeTienne, D. R., Shepherd, D. A. and De Castro, J. O. *The fallacy of "only the strong survive": The effects of extrinsic motivation on the persistence decisions for under-performing firms.* Journal of Business Venturing, 23(5), 2008, 528–546.

[10] Carta, *State of Private Markets Q4 2025*, cited via Waveup, *Capital Efficiency in 2026*, 22 Apr 2026 — 60% of Q4 2025 capital to the top 10% of rounds.
https://waveup.com/blog/capital-efficiency/

[11] Forecastr. *Burn multiple and beyond: The startup metrics founders need for capital efficiency* — burn multiple, developed by David Sacks at Craft Ventures; 18–24 months runway to next raise as the working benchmark. 15 Jun 2026.
https://www.forecastr.co/blog/burn-multiple

[12] Affinity. *10 AI tools transforming venture capital in 2026* — survey of nearly 300 private capital dealmakers, 85% now using AI. Published by a software vendor with a commercial interest in the finding. 3 Aug 2026.
https://www.affinity.co/guides/vc-ai-tools

[13] Qubit Capital. *AI Startup Due Diligence Documents & Metrics Checklist Guide* — automated review accelerating contract and data analysis by 70–80%. 23 Jan 2026.

https://qubit.capital/blog/ai-startup-due-diligence-documents-metrics

[14] Axion Lab. *AI Due Diligence in 2026* — AI handling the data-intensive analysis consuming 60–70% of a diligence team's time. 24 Mar 2026.
https://axionlab.ai/blog/ai-due-diligence-2026

[15] B.BIAS. *Failing to fail: behavioral insights behind startup failure* — on the ostrich effect, survivorship bias, and shame following failure hindering learning from it. 19 Apr 2021.
https://bbiasblog.com/2021/04/09/failing-to-fail-behavioral-insights-behind-startup-failure/

[16] Wiltbank, R. and Boeker, W. *Returns to Angel Investors in Groups.* Ewing Marion Kauffman Foundation, November 2007. 539 angels, 3,097 investments, 1,137 exits. Below-median diligence 1.1x over 3.4 years; above-median 5.9x over 4.1 years; 40+ hours averaging 7.1x.
https://www.researchgate.net/publication/228264869_Returns_to_Angel_Investors_in_Groups

[17] Crowdwise. *How Much Time to Spend on Due Diligence* — detailed restatement of the Wiltbank and Boeker figures including IRR of 2.2% below median and 43.3% above; median 20 hours, mean 60 hours. 9 Dec 2023.
https://crowdwise.org/crowd-investing-101/due-diligence-crowdfunding-time/

[18] Angel Investors Network. *Why Most Angel Investors Lose Money (And What Works)* — split of the 1,137 exits at the 20-hour median: 65% below 1x for below-median diligence, 45% above median. 17 Jun 2026.
https://angelinvestorsnetwork.com/angel-investing/angel-investor-returns-data

[19] DeGennaro, R. P. and Dwyer, G. P. (2014) — questioning whether higher returns are a direct result of angel diligence activity or of correlated investor characteristics. Cited via [16].

[20] Innovestor. *What returns to expect as an angel investor?* — over 50% of single investments return under 1x, falling below 40% across angels holding multiple positions. Referencing Wiltbank & Brooks, Kauffman Foundation, *Tracking Angel Returns*, May 2016. 12 Jun 2024.
https://innovestorgroup.com/blog-what-returns-to-expect-as-angel-investor/

[21] VC Beast. *Angel Investing Returns: What the Data Actually Shows* — Kauffman finding of materially higher returns for portfolios of 20 or more companies. 3 Apr 2026.
https://vcbeast.com/angel-investing-returns-what-the-data-actually-shows

[22] Play Money. *How Angel Investors Make Money: Returns & Exits* — Tech Coast Angels, 4 of 247 outcomes exceeding 100x, 25% IRR across 1997–2022; Angel Capital Association 2023 benchmarks. 13 Jun 2026.
https://letsplaymoney.com/blog/how-angel-investors-make-money

[23] Invernizzi, A. C., Menozzi, A., Passarani, D. A., Patton, D. and Viglia, G. *Overconfidence: A common psychological attribute of entrepreneurs which leads to firm failure.* New England Journal of Entrepreneurship, 2019 — on entrepreneurship as structurally comparable to gambling; on technologically advanced or resource-intensive products increasing overconfidence (citing Simon and Shrader, 2012); on overconfidence predicting escalation better when results deteriorate (citing McCarthy et al., 1993).
https://www.emerald.com/neje/article/23/1/25/308836/OverconfidenceA-common-psychological-attribute-of

[24] Commerce Catalyst. *Types of Founder Cognitive Biases That Kill Startups* — over 80% of founders overestimating probability of success; experience not reliably reducing bias and potentially strengthening it in habitual entrepreneurs; structural countermeasures including the pre-mortem developed by Gary Klein. 9 Jun 2026.
https://commercecatalyst.ai/blog/types-of-founder-cognitive-biases-that-kill-startups/

[25] Staw, B. M. and Ross, J. (1989) — escalation of commitment most pronounced when the person who made the original decision evaluates whether to continue; personal accountability creating ego investment that compounds sunk cost. Summarized via Cognitive Train, *The Sunk Cost Fallacy*.
https://cognitivetrain.com/sunk-cost-fallacy/

[26] Leadership IQ. *The Sunk Cost Fallacy* — synthesis of four decades of research across behavioral economics, psychology and neuroscience; on sunk-cost processing localized to emotional circuits that can override rational evaluation. 26 Jan 2026.
https://www.leadershipiq.com/blogs/leadershipiq/the-sunk-cost-fallacy

A note on source quality

Peer-reviewed or foundation research: [5], [7], [9], [16], [20] and [23]. [8], [19] and [25] are cited within those sources rather than consulted in the original, and are identified as such.

The Wiltbank data. References [16] to [18] all restate the same 2007 Kauffman study, which remains the largest angel-returns dataset available and is the basis for most figures quoted in the field. It is self-reported and its causal direction has been questioned in the literature [19]. We cite it because it is the best evidence available, not because it is conclusive, and Section 3 says so where the figures appear.

Weaker sources. The AI adoption figures in [12], [13] and [14] are directionally consistent with practice but should not be treated as established; [12] is published by a software vendor with a commercial interest in the result. References [15], [21], [22], [24] and [26] are practitioner syntheses of academic work; where a specific finding from them is load-bearing in Section 3, the underlying study is named in the citation so it can be checked.

Revision history

v1.0 3 Sep 2026. Four assessments, twenty-three dimensions, written for founders.

v1.1 Section 3 expanded for angel and seed investors: concentration and diligence-hours evidence, opportunity cost, escalation of commitment, duration of conviction. Twelve references added.

v1.2 The 1999 account published in full under the founding principal's name, with the three failures mapped to the assessments that exist because of them. Deal provenance added as the question no assessment scores.

CLIMATE SPRINTS · WHAT CAPITAL IS ACTUALLY SCORING · V1.2

CLIMATESPRINTS.COM · PRACTICE@CLIMATESPRINTS.COM